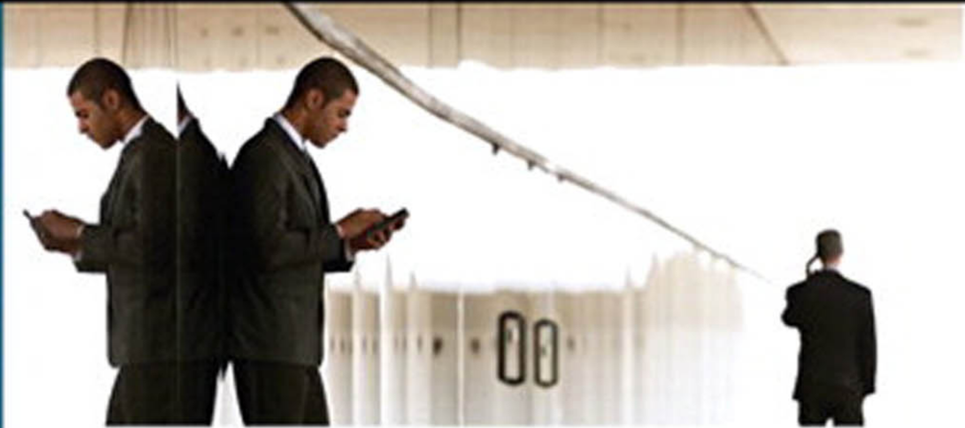

Designing and Deploying 802.11 Wireless Networks

Second Edition

A Practical Guide to Implementing 802.11n and 802.11ac Wireless Networks For Enterprise-Based Applications

ciscopress.com

Jim Geier

Designing and Deploying 802.11 Wireless Networks

A Practical Guide to Implementing 802.11n and 802.11ac Wireless Networks For Enterprise-Based Applications

Second Edition

Jim Geier

Cisco Press

800 East 96th Street

Indianapolis, IN 46240

Note 802.11 exchanges data between stations in frames, which have a specific format for carrying information from higher layers, as well as 802.11 management and control data. For a description of different types of 802.11 frames, see the sections "MAC Frame Structures" and "MAC Frame Types," later in this chapter.

In 802.11 terminology (and in relation to the MAC layer), the payload of an 802.11 data frame is referred to as a *MAC service data unit* (*MSDU*), which is the information handed down to the MAC layer from the logical link control (LLC) layer. This is part of the layering process that takes place with communications protocols. For example, the MSDU may be an LLC information frame carrying a voice packet. The complete MAC layer frame, which includes header, frame body (payload), and FCS (frame check sequence) fields, is what 802.11 refers to as a *MAC protocol data unit* (*MPDU*). Figure 6-2 shows the distinction between the MSDU and MPDU. This terminology might not seem important to know, but it is necessary as the basis for understanding other functions, such as frame aggregation and 802.11 physical layer operations. The 802.11 physical layer, for instance, wraps the MPDU with an additional header (see Chapter 7 for more information).

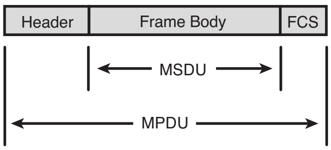

Figure 6-2 *Distinction Between MSDUs and MPDUs*

Medium Access

Because multiple stations share the wireless medium, 802.11 has strict rules for stations accessing the medium. These rules apply to the transmission of all types of 802.11 frames, including data frames, control frames, and management frames.

The following are different types of medium access that 802.11 specifies:

- **Distributed coordination function (DCF):** The DCF access method is part of the original 802.11 standard and uses a carrier sensing access method similar to Ethernet that provides distributed asynchronous (unpredictable) access to the medium.
- **Point coordination function (PCF):** The PCF access method is also part of the original 802.11 standard, but it implements a polling function that determines when a specific client station can transmit. As a result, PCF provides synchronous (predictable) access to the medium, which supports time-bounded delivery of information. Very few vendors implemented PCF, so it is not described in this book.
- **Hybrid coordination function (HCF):** The HCF access method was introduced to the 802.11 standard through the 802.11e amendment as an enhancement to the original DCF and PCF in order to support QoS needs.

Distributed Coordination Function

The DCF access method implements carrier sense multiple access with collision avoidance (CSMA/CA), a contention-based protocol similar to IEEE 802.3 Ethernet that requires stations to decide on their own when to access the medium, based on the presence or absence of traffic. Think of this process of accessing the medium as a meeting where everyone is polite and each person speaks only when no one else is talking. The DCF is a mandatory medium access protocol that was introduced in the original 802.11 specification. As a result, it is part of all 802.11 networks.

The DCF uses a combination of physical and virtual carrier-sense mechanisms to determine whether the medium is busy or idle. If both physical and virtual mechanisms indicate an idle medium, the station can transmit data. If not, the station must wait. Figure 6-3 illustrates the DCF process. The PHY layer provides a physical means of sensing the channel. The result of the physical channel assessment from the PHY layer is sent to the MAC layer as part of the information in deciding the status of the channel.

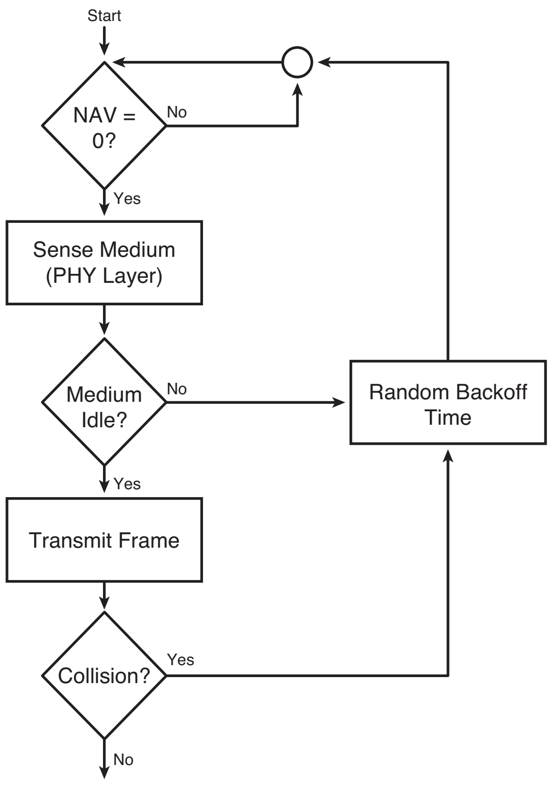

Figure 6-3 *Flowchart Illustrating the Operation of the Distributed Coordination Function*

The MAC layer carries out the virtual carrier sense protocol based on reservation information found in the Duration field of all frames. This information announces (to all other stations) a station's impending use of the medium. The MAC layer monitors the Duration field in all MAC frames and places this information in the station's NAV (network allocation vector) if the value is greater than the current NAV value. The NAV operates like a timer, starting with a value equal to the Duration field value of the last frame transmission sensed on the medium and counting down to zero.

As a condition to accessing the medium, the MAC layer checks the value of its NAV. If the NAV equals zero and the PHY layer indicates a clear channel (that is, received signal strength indication [RSSI] is below a specific threshold), the station can transmit a frame. Just before sending a frame, a station calculates the amount of time necessary to send the frame, based on the frame's length and data rate. The station places a value representing this time in the Duration field in the header of the frame. This process reserves the medium for the sending station because the Duration field causes the MAC layers in other stations to hold off transmissions until the sending station is done sending its frame.

An important aspect of the DCF is a random backoff time for which a station must wait if it detects a busy medium. If the channel is in use, the station must wait a random period of time before attempting to access the medium. Thus, if the PHY layer indicates that the medium is not clear (that is, the RSSI threshold is above a specific threshold), the MAC layer implements a backoff algorithm, regardless of the status of the NAV. This avoids the probability of collisions among stations waiting to transmit. The period of time immediately following a busy medium is when the highest probability of collisions occurs, especially under high utilization. Multiple stations may be waiting for the medium to become idle and will attempt to transmit at the same time. Once the medium is idle, a random backoff time defers each station a different amount of time, causing stations to wait different periods of time before transmitting a frame, minimizing the chance that stations will collide.

In addition to implementing a backoff for busy medium indications, stations use the same backoff mechanism when retransmitting frames due to transmission errors. Under low utilization, stations are not forced to wait very long before transmitting their frames. If the utilization of the network is high, the protocol holds back stations for longer period of times to avoid the probability of multiple stations transmitting at the same time.

Hybrid Coordination Function

The HCF includes two medium access methods: Enhanced Distributed Channel Access (EDCA), which is an extension of DCF for accommodating multiple priority levels, and HCF Controlled Channel Access (HCCA), which is based on PCF for providing time-bounded delivery of information.

Enhanced Distributed Channel Access (EDCA)

The 802.11e amendment extended the original 802.11 DCF by defining multiple queues, with each queue pertaining to a different access category/priority (see Figure 6-4). Each queue has an EDCA function (EDCAF) that determines when a frame in a respective queue can be transmitted (in addition to the access methods defined for DCF). This time is represented as an interval referred to in the 802.11 standard as the Arbitration Interframe Space (AIFS). Each queue has a different AIFS, with the higher-priority queues having a shorter AIFS and the lower-priority queues having a longer AIFS. As a result, the higher-priority traffic is permitted to transmit before lower-priority traffic.

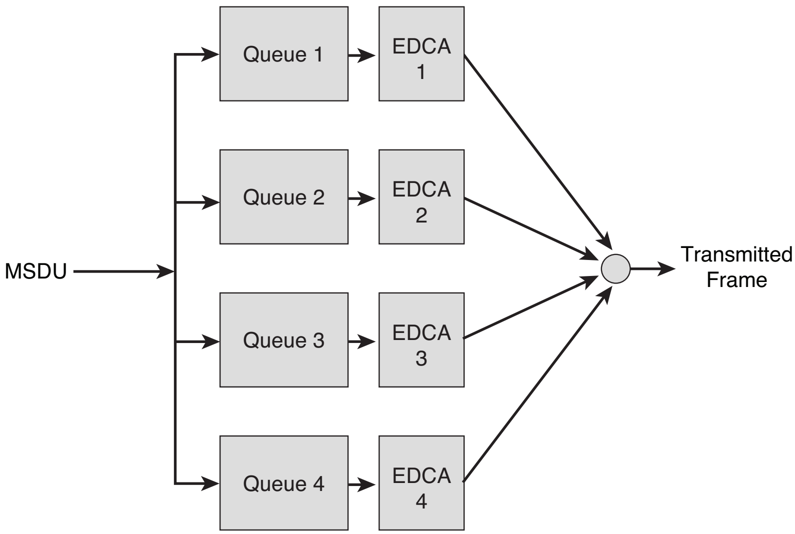

Figure 6-4 *Multiple Queues with Separate EDCA Functions*

HCF Controlled Channel Access (HCCA)

HCCA is a method that 802.11 specifies for priority-based, contention-free access to the medium. The 802.11e amendment to the 802.11 standard introduced HCCA, which includes enhancements to the PCF that was part of the original 802.11 standard (but very few vendors implemented actual products based on PCF).

With HCCA, the access point includes a hybrid coordinator (HC) that controls the transmission of frames from stations (see Figure 6-5). Instead of competing for access to the medium, all stations obey the HC during what is referred to as the *contention-free period*, which is when the HCCA is valid. A client station must receive a poll from the HC before the client station can send a data frame. For example, in Figure 6-5, Client Station A can send a data frame after receiving the poll from the HC. All other stations will refrain from sending frames until they receive a poll from the HC. For WLANs implementing HCCA, there is an alternation of the contention-free period and a contention period. Stations use DCF during the contention-free period and the HCCA during the contention period. This allows some stations to have asynchronous access to the medium while minimizing the overhead of polls from the HC (that is, contention period) and priority access to the medium at the expense of additional overhead to implement the HCCA protocol (that is, contention-free period). This is essentially how PCF works.

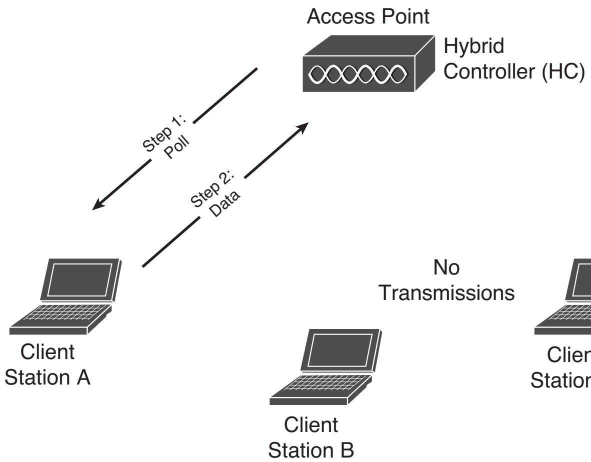

Figure 6-5 *Operation of the HCCA Protocol*

The HC, however, can also poll stations during both the contention-free period and the contention period to reduce latency of higher-priority data transfers. This is different from the original PCF, which allowed polling to occur only during the contention-free period. Also, each station with HCCA can send multiple data frames for each HC poll instead of only sending a single frame as with PCF.

Error Recovery

Because of transmission impairments, 802.11 frames can become corrupted. For example, a station may send a data frame and, while it is en route, strong signals from a nearby microwave oven may alter the bits in the frame. Or two stations may transmit at the same time, despite 802.11's attempt to keep collisions from happening via the random backoff time. Because of these problems, the MAC layer includes error recovery mechanisms.

Data Frame Acknowledgements

To recover from these situations, the 802.11 protocols incorporate the use of an Acknowledgement (ACK) frame and Block ACK frame, as shown in Figure 6-6. If a destination station receives a directed (unicast) data frame without errors, the destination station always sends an ACK frame to the station that sent the data frame. The receiving stations do not acknowledge multicast frames, though. If the sending station does not receive an ACK frame for a directed frame within a specific period of time, the sending station will assume that the corresponding data frame was corrupted and attempt to retransmit the data frame. Retransmissions will take place several times (the actual number depends on the vendor) before the sending station gives up. If the sending station is not able to successfully send the data frame and receive an ACK frame, higher-level protocols (such as TCP) may provide an additional level of error recovery.

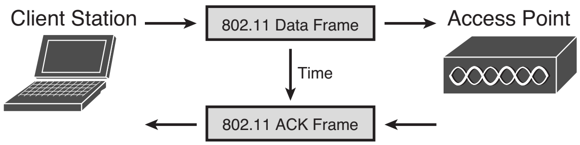

Figure 6-6 *The 802.11 ACK Frame Indicates Successful Reception of a Data Frame*

Dynamic Rate Switching

Something else that relates to data frame error control is 802.11's dynamic rate shifting. If excessive retransmissions occur, stations can operate at a lower data rate to increase the range boundary, mainly for the purpose of decreasing the signal-to-noise requirements at the receiver. The inability to successfully send a data frame (that is, no ACK frame in response) automatically prompts the sending station to lower its data rate. The receiving station will know that the data rate is lower when analyzing the 802.11 PHY layer header of the data frame, which includes data rate information.

An important point, though, is that data rate is not always a good measure of performance. The data rate is the number of data bits sent per second, but it only applies during the time that the frame is actually sent. There are lots of times when other stations are using the shared medium and possibly radio signal interference is causing delays. In addition, the 802.11 protocol periodically sends control and management frames that do not carry any information from higher layers. As a result, the actual throughput varies widely, depending on utilization, and is much lower than the data rate (often less than 50 percent lower).

Hands-on Exercise: Observing 802.11 Dynamic Rate Shifting

This exercise allows you to observe 802.11 dynamic rate shifting in action by operating a client device and moving away from an access point (where signal level decreases, causing transmission errors and corresponding reduced data rates). This will help you better understand various data rates that a client radio can support throughout the coverage area, which provides a measure of how well the client radio will perform under varying signal levels.

Perform the following steps:

Step 1. Obtain test equipment. You need a wireless client device that indicates the association data rate. For example, you could use an 802.11n-equipped laptop. You can use the Windows wireless configuration utility or the client radio configuration software supplied by the radio vendor to observe the association data rate (sometimes referred to as "speed"). You also need an access point.