University of Science and Technology

Faculty of Computer Science and Information Technology

Information Technology Department

Computer Networking: Lecture (8)

Instructor: Prof. Dr. Noureldien Abdelrahman   Date 2/4/2016

Wireless and Mobile Networks

6.3 Wi-Fi: 802.11 Wireless LANs

There are several 802.11 standards for wireless LAN technology, including 802.11a, 802.11b, and 802.11g. **Table 6.1 summarizes the main characteristics of these standards**. 802.11g is by far the most popular technology. A number of dual mode (802.11a/g) and tri-mode (802.11a/b/g) devices are also available.

Standard	Frequency Range (United States)	Data Rate
802.11b	2.4–2.485 GHz	up to 11 Mbps
802.11a	5.1–5.8 GHz	up to 54 Mbps
802.11g	2.4–2.485 GHz	up to 54 Mbps

Table 6.1 ♦ Summary of IEEE 802.11 standards

## 6.3.1 The 802.11 Architecture

Figure 6.7 illustrates the principal components of the 802.11 wireless LAN architecture.

The fundamental building block of the 802.11 architecture is the **basic service set (BSS)**.

A BSS contains one or more wireless stations and a central base station, known as an access point (AP).

Figure 6.7 shows the AP in each of two BSSs connecting to an interconnection device (such as a switch or router), which in turn leads to the Internet.

In a typical home network, there is one AP and one router (typically integrated together as one unit) that connects the BSS to the Internet.

Figure 6.7 ◆ IEEE 802.11 LAN architecture

Figure 6.8 shows that IEEE 802.11 stations can also **group themselves together to form an ad hoc network—a network with no central control and with no connections to the "outside world."** Here, the network is formed "on the fly," by mobile devices that have found themselves in proximity to each other, that have a need to communicate, and that find no preexisting network infrastructure in their location.

Figure 6.8 ♦ An IEEE 802.11 ad hoc network

6.3.2 Channels and Association

In 802.11, **each** *wireless station needs to associate with an AP before it can send or* receive network-layer data.

When a network administrator installs an AP, the administrator assigns **Service Set Identifier (SSID)** to the access point.

The administrator must also **assign a channel number to the AP**. To understand channel numbers, recall that 802.11 operate in the frequency range of 2.4 GHz to 2.485 GHz. Within this 85 MHz band, 802.11 **define 11 partially overlapping channels**. Any two channels are non-overlapping if and only if they are separated by four or more channels. In particular, the set of channels 1, 6, and 11 is the only set of three non-overlapping channels.

 What is a Wi-Fi Jungle?

A Wi-Fi jungle is any physical location where a wireless station receives a sufficiently strong signal from two or more APs.

Suppose there are five APs in the Wi-Fi jungle. To gain Internet access, your wireless station needs to join exactly one of the ==subnets and hence needs to **associate** with exactly one of the APs==.

Associating means the *wireless station creates a virtual wire between itself and the AP*. Specifically, only the associated AP will send data frames (that is, frames containing data, such as a datagram) to your wireless station, and your wireless station will send data frames into the Internet only through the associated AP.

But how does your wireless station associate with a particular AP? And more fundamentally, how does your wireless station know which APs, if any, are out there in the jungle?

The 802.11 standard requires that **an AP periodically** send **beacon frames**, each of which includes the AP's SSID and MAC address. **Your wireless station, knowing that APs are sending out beacon frames, scans the 11 channels**, seeking beacon frames from any APs that may be out there.

Having the available APs from the beacon frames, you (or your wireless host) select one of the APs for association.

How to Select the AP?

The 802.11 **standard does not specify an algorithm for selecting which of the available** APs to associate with; that algorithm is left up to the designers of the 802.11 firmware and software in your wireless host.

Typically, the host chooses the **AP whose beacon frame is received** with the **highest signal strength**. While high signal strength is good, signal strength is not the only AP characteristic that will determine the performance a host receives. In particular, it's possible that the selected AP may have a strong signal, but may be overloaded with other affiliated hosts (that will need to share the wireless bandwidth at that AP), while an unloaded AP is not selected due to a slightly weaker signal.

What is Passive and Active Scanning?

The process of scanning channels and listening for beacon frames is known as **passive scanning** (see Figure 6.9a). A wireless host can also perform **active scanning**, by broadcasting a probe frame that will be received by all APs within the wireless host's range, as shown in Figure 6.9b. APs respond to the probe request frame with a probe response frame. The wireless host can then choose the AP with which to associate from among the responding APs.
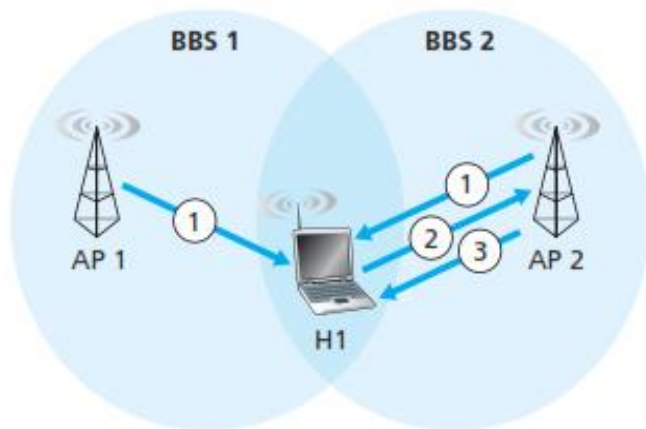
a. Passive scanning
1. Beacon frames sent from APs
2. Association Request frame sent:
   H1 to selected AP
3. Association Response frame sent:
   Selected AP to H1

a. Active scanning
1. Probe Request frame broadcast from H1
2. Probes Response frame sent from APs
3. Association Request frame sent:
   H1 to selected AP
4. Association Response frame sent:
   Selected AP to H1

Figure 6.9 ♦ Active and passive scanning for access points

The 802.11 MAC Protocol

Once a wireless station is associated with an AP, it can start sending and receiving data frames to and from the access point. **But because multiple stations may want to transmit data frames at the same time over the same channel, a multiple access protocol is needed to coordinate the transmissions**. Here, a **station** is either a wireless station or an AP.

802.11 MAC protocol is **CSMA/CA**. The "CSMA" in CSMA/CA stands for "carrier sense multiple access," meaning that each station senses the channel before transmitting, and refrains from transmitting when the channel is sensed busy, and CA stands for "Collision Avoidance".

Although both Ethernet and 802.11 use carrier-sensing random access, the two MAC protocols have important differences.

First, instead of using collision detection, 802.11 use collision-avoidance techniques.

Second, because of the relatively high bit error rates of wireless channels, 802.11 (unlike Ethernet) use a **link-layer acknowledgment/retransmission (ARQ) scheme**.

Why the 802.11 MAC protocol does *not* implement collision detection?

The 802.11 MAC protocol does *not* implement collision detection. There are two important reasons for this:

1- The ability to detect collisions requires the adapter to be able to transmit and listen at the same time. Because the strength of the received signal is typically very small compared to the strength of the transmitted signal at the 802.11 adapter, <u>it is costly to build hardware that</u> can detect a collision.

2- More importantly, even if the adapter could transmit and listen at the same time (and abort transmission when it senses a busy channel), the adapter would still not be **able to detect all** collisions**, due to the hidden terminal problem and fading.**

Because 802.11wireless LANs do not use collision detection, once a station begins to transmit a frame, *it transmits the frame in its entirety*. As one might expect, transmitting entire frames

(particularly long frames) when collisions happen can significantly degrade a multiple access protocol's performance. In order to reduce the likelihood of collisions, 802.11 **employ Collision-avoidance techniques**.

When a station in a wireless LAN sends a frame, the frame may not reach the destination station for a variety of reasons, how to deal with this?

To deal with this the 802.11 MAC protocol **uses link-layer acknowledgments**. As shown in Figure 6.10, when the destination station receives a frame, it waits a short period of time known as the **Short Inter-frame Spacing (SIFS)** and then sends back an acknowledgment frame.

If the transmitting station does not receive an acknowledgment within a given amount of time, it assumes that an error has occurred and retransmits the frame, using the CSMA/CA protocol to access the channel.

If an acknowledgment is not received after some fixed number of retransmissions, the transmitting station gives up and discards the frame.


Figure 6.10 ◆ 802.11 uses link-layer acknowledgments

How the 802.11 CSMA/CA Protocol Reduce Collisions? (That is Collision A voidance)

Suppose that a station (wireless station or an AP) has a frame to transmit.

1. If initially the station senses the channel idle, it transmits its frame after a short period of time known as the **Distributed Inter-frame Space (DIFS);** see Figure 6.10.

2. Otherwise, the station chooses a random **back off value and counts down this value** when the channel is sensed idle. While the channel is sensed busy, the counter value remains frozen.

3. When the counter reaches zero (note that this can only occur while the channel is sensed idle), the station transmits the entire frame and then waits for an acknowledgment.

4. If an acknowledgment is received, the transmitting station knows that its frame has been correctly received at the destination station. If the station has another frame to send, it begins the CSMA/CA protocol at step 2.

If the acknowledgment isn't received, the transmitting station reenters the back off phase in step 2, with the random value chosen larger.

Dealing with Hidden Terminals

The 802.11 MAC protocol also includes a scheme that helps avoid collisions even in the presence of hidden terminals.

From figure (6.11), let's now consider why hidden terminals can be problematic.

Suppose Station H1 is transmitting a frame and halfway through H1's transmission, Station H2 wants to send a frame to the AP. H2, not hearing the transmission from H1, will first wait a DIFS interval and then transmit the frame, resulting in a collision. The channel will therefore be wasted during the entire period of H1's transmission as well as during H2's transmission.

Figure 6.11 ♦ Hidden terminal example: H1 is hidden from H2, and vice versa

In order to avoid this problem, the IEEE 802.11 protocol allows a station to use a short **Request to Send (RTS)** control frame and a short **Clear to Send (CTS)** control frame to *reserve* access to the channel.

1- When a sender wants to send a DATA frame, it can first send an RTS frame to the AP, <u>indicating the total time</u> **required to transmit the DATA frame and the acknowledgment (ACK) frame**.

2- When the AP receives the RTS frame, it responds by broadcasting a CTS frame. This <u>**CTS frame serves two purposes**</u>: It gives the

Thus, in Figure 6.12, before transmitting a DATA frame, H1 first broadcasts an RTS frame, which is heard by all stations in its circle, including the AP. The AP then responds with a CTS frame, which is heard by all stations within its range, including H1 and H2. Station H2, having heard the CTS, will not **transmitting for the time specified in the CTS frame**. The RTS, CTS, DATA, and ACK frames are shown in Figure 6.12.


Figure 6.12 ◆ Collision avoidance using the RTS and CTS frames