CHAPTER 7

WIRELESS LOCAL AREA NETWORKS

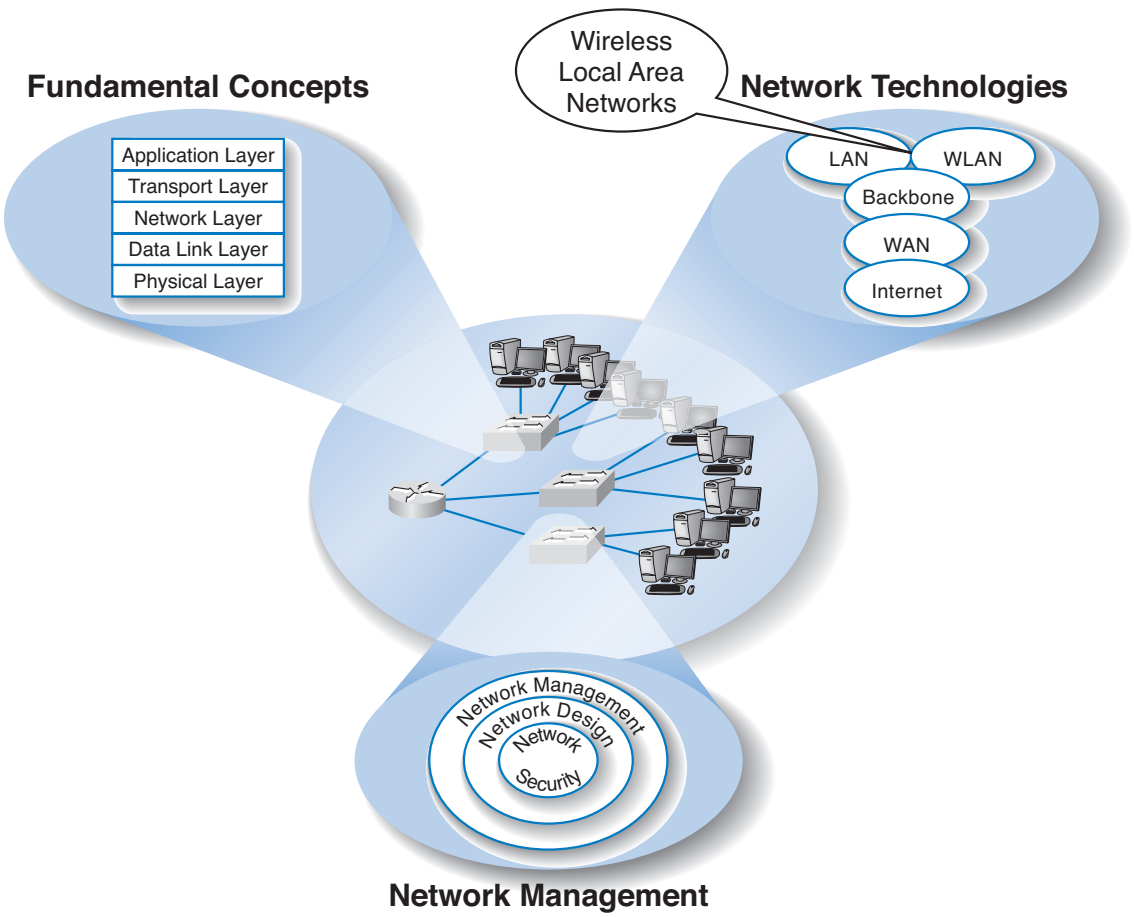

The Three Faces of Networking

ALTHOUGH TRADITIONAL wired Ethernet LANs dominate today's network environment, wireless LANs (WLANs) are common. This chapter describes the basic components of a WLAN and then examines three common wireless technologies: Wi-Fi (IEEE 802.11), WiMAX (IEEE 802.16), and Bluetooth (IEEE 802.15). The chapter ends with a discussion of best practice WLAN design, including security, and how to improve performance.

OBJECTIVES

- Understand the major components of WLANs
- Understand Wi-Fi
- Be familiar with Wi-Max
- Be familiar with Bluetooth WLANs
- Be familiar with how to improve WLAN performance
- Be familiar with WLAN security
- Understand the best practice recommendations for WLAN design

CHAPTER OUTLINE

INTRODUCTION

WLAN COMPONENTS

- Network Interface Cards
- Access Points
- Radio Frequencies

WI-FI

- Topology
- Media Access Control
- Types of Wi-Fi
- Wi-Fi as Public Internet Access

WIMAX

- Topology
- Media Access Control
- Types of WiMAX

BLUETOOTH
- Topology
- Media Access Control

THE BEST PRACTICE WLAN DESIGN
- Effective Data Rates
- Costs
- Recommendations
- Physical WLAN Design
- WLAN Security

IMPROVING WLAN PERFORMANCE
- Improving Device Performance
- Improving Circuit Capacity
- Reducing Network Demand

IMPLICATIONS FOR MANAGEMENT

SUMMARY

INTRODUCTION

The use of *Wireless LANs (WLANs)* is growing rapidly. A recent survey of network managers indicated that 90 percent of companies are using wireless LANs, usually in addition to traditional wired LANs. Wireless LANs transmit data through the air using radio transmission rather than through twisted-pair cable or fiber-optic cable. This has been one area of networking that has seen the greatest changes in a short amount of time. From a time with no widely accepted standards (2000), we have today gone to an alphabet soup of standards (e.g., 802.11a, 802.11b, 802.11g, 802.11n, 802.15, 802.16d, 802.16e).

WLANs serve the same purpose as LANs: they are used to connect a series of computers in the same small local area to each other and to a backbone network. WLANs are usually not totally wireless in that they are most commonly used to connect a set of wireless computers into a wired network. However, WLANs enable you to use the network in places where it is impractical to put a wired network (either because of cost or access). WLANs can enable staff to pull up a chair and work on the network from a lunchroom, a corridor, or an outdoor patio. WLANs also enable mobile staff to work at different locations in the office building or to move their computers easily from one location to another. WLANs are becoming popular in hospitals, for example, because they enable doctors and nurses to use laptops and tablet PCs to access patient records. WLANs are also popular in airports because they enable business travelers to connect to the Internet from any waiting area.

This chapter examines the basic components of a WLAN and then examines three commonly used WLAN technologies (Wi-Fi, WiMAX, and Bluetooth). The chapter ends

with a discussion of best practice recommendations for WLAN design and ways to improve WLAN performance.

As with Ethernet in the previous chapters, the three primary WLAN technologies (Wi-Fi, WiMAX, and Bluetooth) are layer 2 protocols that operate at the data link layer. They too must have physical hardware at layer 1 that meets their requirements and software at layers above them (e.g., TCP/IP) that enables application software to use them.

WLAN COMPONENTS

In the last chapter on LANs, we discussed the three key components of the LAN: the network interface card, the hub/switch, and the cables that connect them. WLANs use the same basic structure. There is a wireless network interface card that is built into a desktop or laptop computer (or can be added later). A wireless access point performs the same functions as a hub or switch. Finally, instead of cable, there is a set of radio frequencies that are used to transport data (see Figure 7.1).

Network Interface Cards

Each computer has a wireless *network interface card* (*NIC*) that is used to connect the computer into the WLAN. The NIC is a radio transceiver in that it sends and receives radio signals through a short range, usually only about 100 meters or 300 feet. WLAN NICs are available for laptops as PCMCIA cards and as standard cards for desktop computers, but laptop computers now come with Wi-Fi NICs built-in.

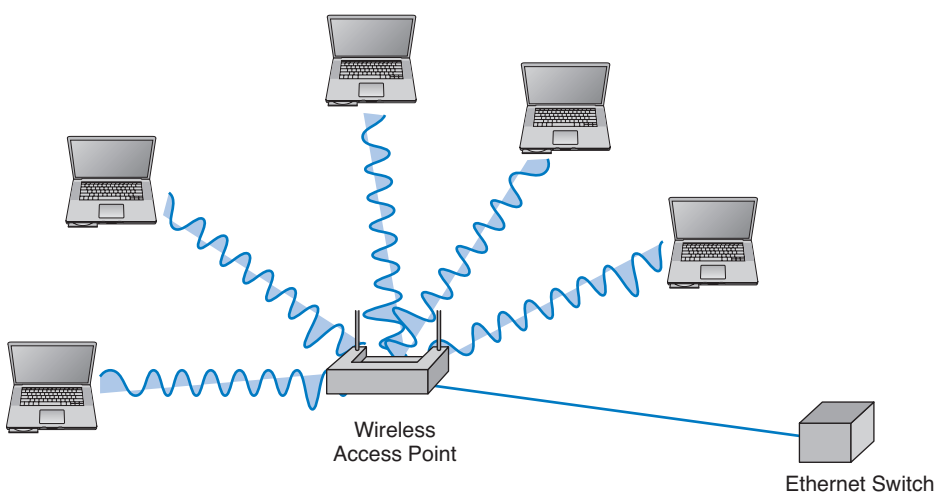

FIGURE 7.1 A wireless access point connected into an Ethernet switch.

Courtesy Alan Dennis

FIGURE 7.2 A wireless access point.

Access Points

A central wireless *access point* (*AP*) is a radio transceiver that plays the same role as a hub or switch in wired Ethernet LANs (Figure 7.2). The AP also connects the WLAN into wired LANs, typically using 100Base-T.

The AP acts as a repeater to ensure that all computers within range of the AP can hear the signals of all other computers in the WLAN. All NICs in the WLAN transmit their packets to the AP and then the AP retransmits the packet over the wireless network to its destination—or retransmits the packet over the wired network to its destination. Wireless NICs never communicate with each other directly; they always transmit through the AP. Therefore, if a message has to be transmitted from one wireless computer to another, it is transmitted twice, once from the sender to the AP and then from the AP to the destination. At first glance this may seem a bit strange because it doubles the number of transmissions in the WLAN. However, very few messages are ever sent from client computer to client computer in a WLAN. Most messages are exchanged between client computers and a server of some kind. For this reason, servers should never be placed on a WLAN. Even if they are intended to serve clients on a WLAN, they should always be placed on the wired portion of the LAN.

Most WLANs are installed using APs that have *omnidirectional antennas*, which means that the antenna transmits in all directions simultaneously. One common omnidirectional antenna is the dipole antenna shown in Figure 7.3a (nicknamed the "rubber duck" because of its flexibility). As Figure 7.3a shows, omnidirectional antennas transmit in all directions, both horizontally and vertically. The signal goes in all directions, as well as up and down, although there is often a small dead spot with no signal that is a very small area directly above the antenna.

The other type of antenna that can be used on APs is the *directional antenna* (Figure 7.3b). As the name suggests, a directional antenna projects a signal only in one direction.

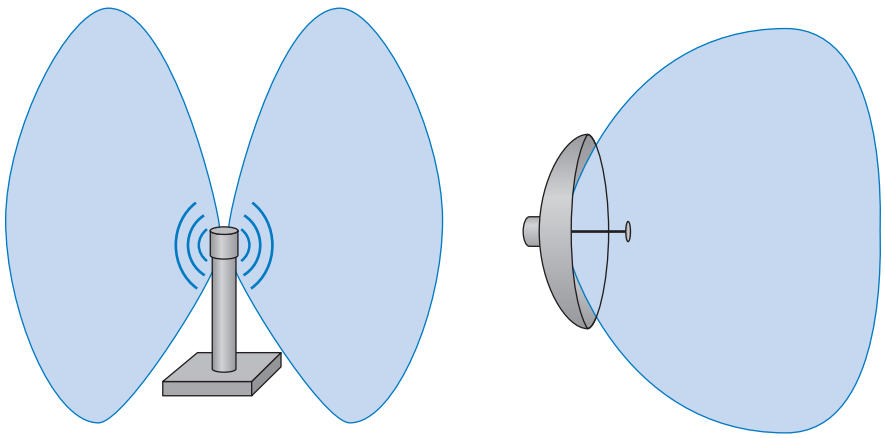

(a) Omnidirectional antenna (b) Directional antenna

FIGURE 7.3 Types of antennas.

Because the signal is concentrated in a narrower, focused area, the signal is stronger and therefore will carry farther than the signal from an AP using an omnidirectional antenna. Directional antennas are most often used on the inside of an exterior wall of a building, pointing to the inside of the building. This keeps the signal inside the building (to reduce security issues) and also has the benefit of increasing the range of the AP.

Radio Frequencies

WLANs use radio transmissions to send data between the NIC and the AP. All radio transmissions are controlled by the government so that no two radio stations attempt to transmit in the same *frequency range*. In the United States, the Federal Communications Commission (FCC) controls the airwaves. In order to transmit in a certain radio frequency band, you need to get permission.

Most countries (but not all), permit WLANs to operate in two frequency ranges that have been reserved for unlicensed transmissions: the 2.4 GHz range and the 5 GHz range[1]. Japan, for example, uses a slightly different set of frequency ranges. In this book, we will focus on the North American standards. WLANs and other unlicensed transmitters such as cordless phones and baby monitors can use these frequency ranges at will—which means that your WLAN and your cordless phone may interfere with each other. Microwave ovens also use the same frequency range and may cause interference.

The frequency range directly affects the data rates that can be transmitted. The larger the frequency range available (called the *bandwidth*), the greater the capacity of the wireless circuit and the faster data can be sent. You can think of the frequency range as the width of a pipe; larger pipes let you move more water per second, and so larger frequency ranges let you move more data per second. The 2.4 GHz range has a smaller bandwidth

[1] Some WLAN technologies operate in other frequency ranges.

than the 5 GHz range, which has nothing to do with the technology. It is just the ranges that were allocated by the FCC and chosen by standards groups. As a result, the 5 GHz range can transmit data faster than the 2.4 GHz range.

Data transmission is also affected by attenuation, which is the gradual weakening of the signal as it travels farther from the transmitter (see Chapter 3). Higher frequencies suffer attenuation more quickly than lower frequencies due to the laws of physics. As a result, transmissions in the 2.4 GHz range can travel farther and through more walls and other sources of interference than can transmissions in the 5 GHz range. As interference increases and the signal strength weakens, the effective bandwidth that can be used decreases and capacity and data rate decreases. This means that wireless technologies that use the 5 GHz can transmit over much shorter distances than technologies that use the 2.4 GHz range. The farther you move from the AP, the worse the data rates as the signal strength weakens.

When we design a WLAN it is important to ensure that the APs don't interfere with each other. If all APs transmitted on the same frequency range the transmissions of one AP would interfere with another AP. Therefore, each AP is set to transmit on a different *channel*, very much like the different channels on your TV. Each channel uses a different part of the 2.4 GHz or 5 GHz frequency range so that there is no interference among the different channels. When a computer first starts using the WLAN, its NIC searches all available channels within the appropriate frequency range and then picks the channel that has the strongest signal to use in its communications.

Figure 7.4 shows how a WLAN might be designed using 5 access points, three using omnidirectional antennas and 2 using directional antennas. This configuration uses 3 channels, with each AP configured to use a channel that does not interfere with the APs around it. The distance covered by each AP ranges from 100–500 feet, depending upon interference. Placing the APs and selecting channels to ensure that the entire area is covered and that there is no interference from APs using the same channel is an important design problem. In Figure 7.4, the two APs using channel A are at opposite ends of the building, as are the APs using channel B. The AP using channel C is placed in the middle so that its coverage overlaps but does not interfere with the others.

As the user *roams* through a building, the NIC continues to use its original channel until the signal strength starts to drop. When this happens, the NIC again listens to and

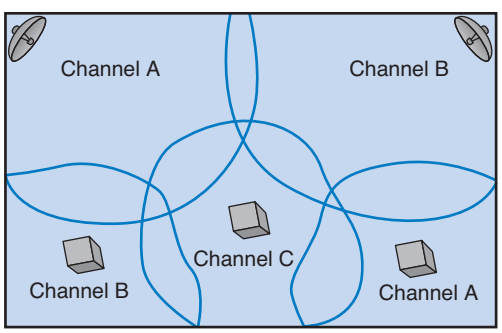

FIGURE 7.4 A WLAN using different channels.

may attempt to transmit using all of the available channels to find a new channel that has the strongest signal. Some NICs also periodically check for better channels when the channel they are using becomes busy.

One potential problem with WLANs is security. Because anyone within range of a WLAN AP can receive transmissions, eavesdropping is a serious threat. Most WLANs encrypt transmissions so that only authorized computers can decode and read the messages. Security is discussed in detail in Chapter 11 although we briefly discuss WLAN security later in this chapter.

WI-FI

Wi-Fi is the commercial name for a set of standards developed by the *IEEE 802.11* standards group. A group of vendors selling 802.11 equipment trademarked the name Wi-Fi to refer to 802.11 because they believe that consumers are more likely to buy equipment with a catchier name than 802.11. Wi-Fi is intended to evoke memories of Hi-Fi, as the original stereo music systems were called.

The 802.11 family of technologies is much like the Ethernet family. The 802.11 standards reuse many of the Ethernet 802.3 components and are designed to connect easily into Ethernet LANs. For these reasons, IEEE 802.11 is often called *wireless Ethernet*. Just as there are several different types of Ethernet (e.g., 10Base-T, 100Base-T, 1000Base-T), there are several different types of 802.11.

MANAGEMENT FOCUS

7-1 CARNIVAL CRUISE LINES GOES WIRELESS

In 2005, the cruise ship *Carnival Valor* went wireless. "Initially, we had planned to increase the number of workstations in our onboard Internet cafes and to expand Internet access to the staterooms using traditional Cat 5 cabling," says Tom McCormick, manager of network engineering for Carnival Cruise Lines. "However, using Cisco wireless technology we are able to provide wireless data access bow-to-stern and, as an added benefit, we were also able to introduce mobile VoIP (Voice over IP wireless telephones) on the same infrastructure." Passengers and crew can access the Internet using any standard laptop or PDA. The ship also provides mobile VoIP phones that can be used anywhere on board.

Designing the network was challenging due to the thick steel bulkheads throughout the ship and the heavy machinery that can often cause radio interference. The network has 217 access points and provides end-to-end voice and data coverage on all decks, including those outdoors. The access points are connected into the traditional wired Ethernet network which connects into a satellite wide area network to provide Internet access. The network also includes special purpose telephone management devices so the VoIP phones can connect into the traditional wired phone network (see Figure 7.5).

SOURCE: G. Knauer, "Voice Goes Wireless," *Packet,* Third Quarter, 2005, pp. 65–69.

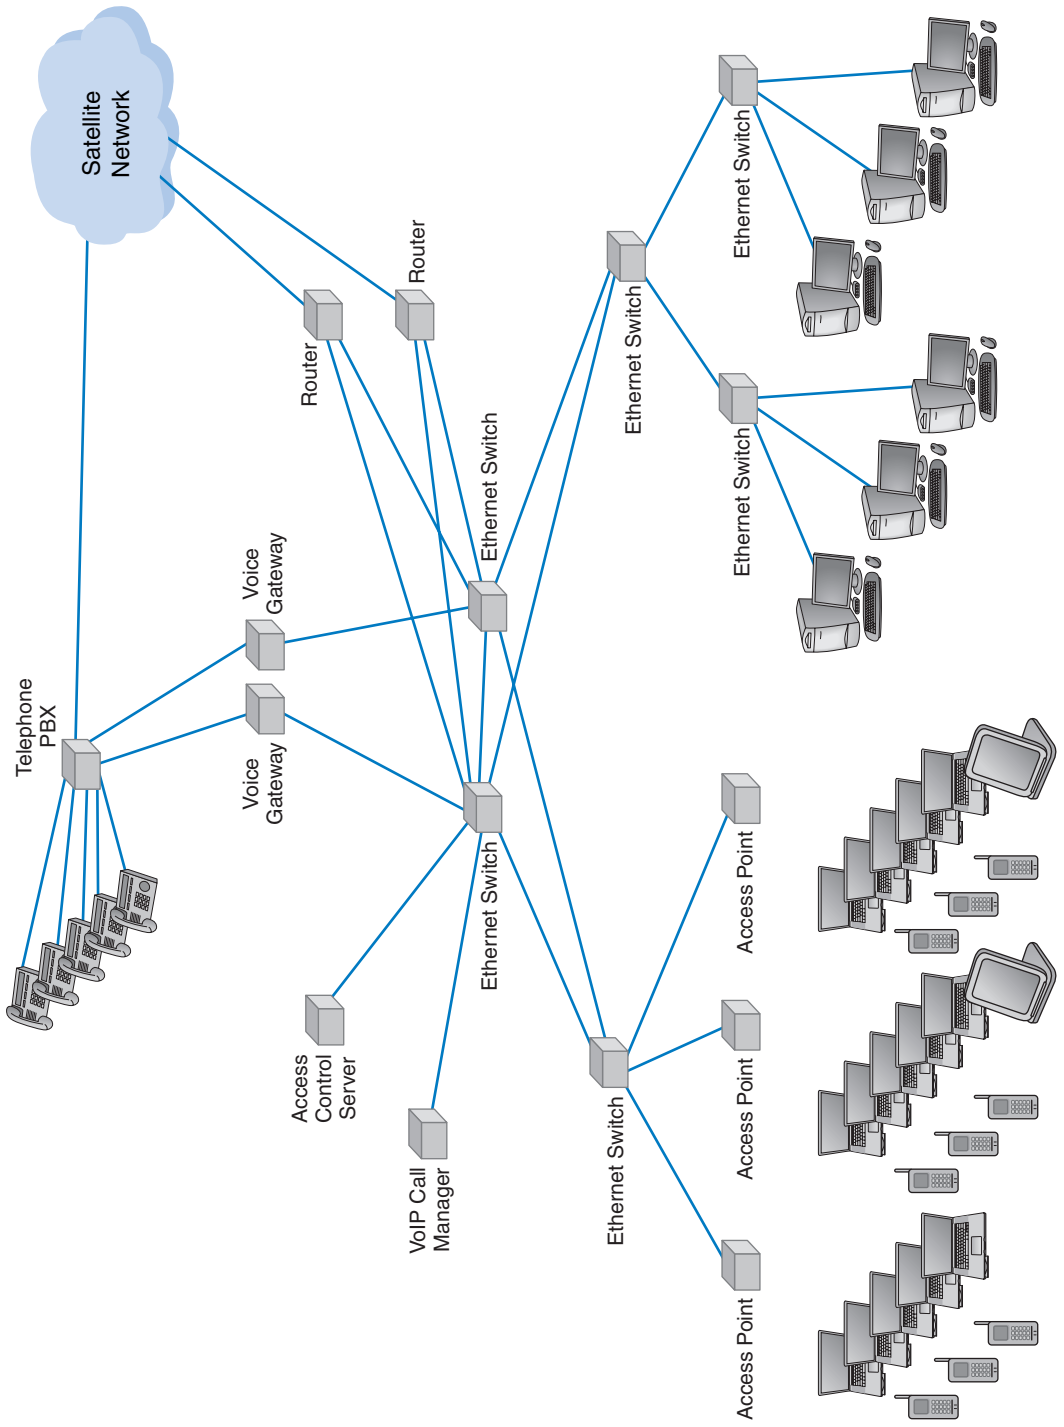

FIGURE 7.5 *Carnival Valor* cruise ship network.

Topology

The logical and physical topologies of Wi-Fi are the same as those of shared Ethernet. They are a physical star and a logical bus. There is a central AP to which all computers direct their transmissions (star), but the radio frequencies are shared (bus) so that all computers must take turns transmitting.

Media Access Control

Media access control in Wi-Fi is *Carrier Sense Multiple Access with Collision Avoidance* (*CSMA/CA*), which is similar to the contention-based CSMA/CD approach used by traditional Ethernet. With CSMA/CA, computers listen before they transmit and if no one else is transmitting, they proceed with transmission. Detecting collisions is more difficult in radio transmission than in transmission over wired networks, so Wi-Fi attempts to avoid collisions to a greater extent than traditional Ethernet. CSMA/CA simultaneously uses two media access control approaches.

Distributed Coordination Function The first media access control method is the *distributed coordination function* (*DCF*) (also called *physical carrier sense method* because it relies on the ability of computers to physically listen before they transmit). With DCF, each packet in CSMA/CA is sent using stop-and-wait ARQ. After the sender transmits one packet, it immediately stops and waits for an ACK from the receiver before attempting to send another packet. When the receiver of a packet detects the end of the packet in a transmission, it waits a fraction of a second to make sure the sender has really stopped transmitting, and then immediately transmits an ACK (or a NAK). The original sender can then send another packet, stop and wait for an ACK, and so on.

While the sender and receiver are exchanging packet and ACKs, other computers may also want to transmit. So when the sender ends its transmission, you might ask why doesn't some other computer begin transmitting before the receiver can transmit an ACK? The answer is that the physical carrier sense method is designed so that the time the receiver waits after the transmission ends before sending an ACK is significantly less time than the time a computer must listen to determine that no one else is transmitting before initiating a new transmission. Thus, the time interval between a transmission and the matching ACK is so short that no other computer has the opportunity to begin transmitting.

Point Coordination Function The second media access control technique is called the *point coordination function* (*PCF*) (also called the *virtual carrier sense method*). DCF works well in traditional Ethernet because every computer on the shared circuit receives every transmission on the shared circuit. However, in a wireless environment, this is not always true. A computer at the extreme edge of the range limit from the AP on one side may not receive transmissions from a computer on the extreme opposite edge of the AP's range limit. In Figure 7.1, all computers may be within the range of the AP, but may not be within the range of each other. In this case, if one computer transmits, the other computer on the opposite edge may not sense the other transmission and transmit at the same time causing a collision at the AP. This is called the *hidden node problem* because the computers at the opposite edges of the WLAN are hidden from each other.

MANAGEMENT FOCUS **7-2 WEST EDMONTON MALL USES WI-FI FOR TOURISM**

Being the largest shopping mall in the world (48 city blocks in size) wasn't enough for the West Edmonton Mall, located in Edmonton, Alberta, in western Canada. Every year more than 22 million visitors flock to the mall's 3 hotels, 8 amusement parks, 21 theaters, 110 restaurants, and 800 shops.

Originally planned as a network to serve the 16,500 people who work in the mall the project quickly grew to gain a competitive edge by enabling visitors to access the Internet. The first phase of the network installed 70 access points in one part of the mall and offered day passes for $11 and month passes for $35.

The response was overwhelming, from both visitors and workers. Visitors can easily surf the Web and e-mail while their children play in the water parks. Many stores have adopted mobile VoIP telephones for their employees.

The network has generated a 120 percent ROI. The next steps to expand the network to the rest of the mall are already in progress. The Mall's owner, Triple Five, is planning to install a similar network in the largest mall in the United States, the Mall of America in Bloomington, Minnesota, which it also owns.

SOURCE: Deborah Mendez-Wilson, "Untethered Utopia," *Network World*, November 21, 2005, pp. 74-76.

When the hidden node problem exists, the AP is the only device guaranteed to be able to communicate with all computers on the WLAN. Therefore, the AP must manage the shared circuit using a controlled-access technique, not the contention-based approach of traditional Ethernet. With this approach, any computer wishing to transmit first sends a *request to transmit* (*RTS*) to the AP, which may or may not be heard by all computers. The RTS requests permission to transmit and to reserve the circuit for the sole use of the requesting computer for a specified time period. If no other computer is transmitting, the AP responds with a *clear to transmit* (*CTS*), specifying the amount of time for which the circuit is reserved for the requesting computer. All computers hear the CTS and remain silent for the specified time period.

The virtual carrier sense method is optional. It can always be used, never used, or used just for packets exceeding a certain size, as set by the WLAN manager. Controlled-access methods provide poorer performance in low-traffic networks because computers must wait for permission before transmitting rather than just waiting for an unused time period. However, controlled-access techniques work better in high-traffic WLANs because without controlled access there are many collisions. Think of a large class discussion in which the instructor selects who will speak (controlled access) versus one in which any student can shout out a comment at any time.

Types of Wi-Fi

Wi-Fi is one of the fastest changing areas in networking. As we write this textbook, there are three types of Wi-Fi in current use, with a new version about to be standardized.

802.11a The IEEE 802.11a standard provides high speed wireless networking in the 5 GHz range. It provides eight channels for indoor use in the United States (plus one

channel for outdoor use). The 802.11a standard provides for more or fewer channels in other parts of the world where the radio frequency spectrum regulation is different.[2]

Each channel provides speeds of 54 Mbps under perfect conditions. Because the higher frequency 5 GHz range is used, the distance between the NIC and the AP is reduced to only 50 meters or 150 feet under perfect conditions; in practice, it is usually less. As interference increases, the data speeds decline, so the farther you are from the AP the lower the data rate you receive. Users at the extreme edge of the range or those facing interference will not be able to communicate at the 54 Mbps. Initial analyses suggest a 54-Mbps data rate is reliable and consistent only up to 50 feet from the AP. Speeds of 26 Mbps and 34 Mbps are more common, and the speed may even drop to 6 Mbps in the face of interference.

802.11b The IEEE 802.11b standard provides moderate speed wireless networking in the 2.4 GHz range. It provides three channels for indoor use in the United States. The 802.11b standard provides for more or fewer channels in other parts of the world where the radio frequency spectrum regulation is different.[3]

Each channel provides a maximum data rate of 11 Mbps. Only when there is significant interference or the signal begins to weaken because the user is moving far from the WLAN does the data rate change in an attempt to improve signal quality. Thus, for those users close to the center of the WLAN, 6–11 Mbps is the norm. The range under ideal conditions is 450 feet, although the actual range in practice is much less than this. The speed may drop to as low as 1 Mbps in the face of interference.

Thus the advantage of 802.11b over 802.11a is that in using the 2.4 GHz frequency range, 802.11b suffers less attenuation and thus the signal has greater range with less decrease in speed as distance from the AP increases. The disadvantage is that 802.11b provides lower speeds than 802.11a.

802.11g The IEEE 802.11g standard provides high speed wireless networking in the 2.4 GHz range. It provides three channels for indoor use in the United States. The 802.11g standards provides for more or fewer channels in other parts of the world whose radio frequency spectrum regulation is different. 802.11g was designed to take the best of both the 802.11a and 802.11g standards and to ultimately replace them.

Each channel provides a maximum data rate of 54 Mbps, with a range under ideal conditions of 450 feet, although the actual range in practice is much less than this. The speed may drop to as low as 6 Mbps in the face of interference.

[2] The channels are numbers 36, 40, 44, 48, 52, 56, 60, and 64, with 149 being for outdoor use. There are really many more channels as the numbers would suggest, but they overlap with these channels so only these are used in the United States. Other countries use different channels.

[3] In the United States, the channel numbers are 1, 6, and 11. As with 802.11a, there are more channels but they are not used because they overlap, and different channels are used in other countries. When 802.11b was first introduced, a four channel configuration was used (channels 1, 4, 8, and 11). With this approach, the channels overlap to some extent so if you run an AP on channel 1 and another on channel 4, there will some interference between the two APs. Field tests showed that the data rates dropped dramatically due to this interference, so although four channel configurations are possible, the best practice recommendation today is to use a three channel configuration.

802.11g is designed to be backward compatible with 802.11b, so that 802.11b devices can operate with an 802.11g access point. This will permit the many existing laptop computers that have built-in 802.11b network cards to work with the newer 802.11g access points although they will not be able to operate at the faster speeds provided by 802.11g. Newer laptops that have built-in 802.11g cards can use the same access points, so both devices can coexist in the WLAN.

However, this backward compatibility comes with a price. 802.11b devices become confused when 802.11g devices operate at high speeds near them, so when an 802.11g access point detects the presence of an 802.11b device, it prohibits 802.11g devices from operating at high speeds.

802.11n The IEEE 802.11n standard is under development as we write. Its goal is to provide very high speed wireless networking using both the 2.4 GHz and 5 GHz frequency ranges simultaneously (by using multiple sets of antennas optimized to the different frequencies) to increase the data speeds it can attain. The standard has not been finalized, but current drafts propose speeds in the 100–240 Mbps range.

As with 802.11g, 802,11n is designed to be backwards compatible with 802.11a, 802.11b, and 802.11g, so that it has the potential to co-exist with, and ultimately replace, all three prior technologies.

Wi-Fi as Public Internet Access

Wi-Fi was initially intended to provide indoor mobile wireless access to organizational LANs and backbones. Many commercial providers now offer Wi-Fi access in public places such as airports and malls, so that users can connect into the Internet and work in public locations. Several towns and commercial providers have also begun to offer outdoor Wi-Fi services as public Internet access.

There are several technical issues in providing large scale public Wi-Fi access, but none are major. The biggest obstacle is political. Some towns have offered these services at no cost or at low cost to residents, which has caused several commercial providers (e.g., Verizon) to complain that the towns are stealing business from them. Several providers have gone to court to stop towns from offering such services. Others have lobbied state governments to introduce laws to prevent towns from offering such services. Fourteen states so far have passed laws prohibiting local governments from offering free or low cost public wireless Internet services. Other states have embraced the idea of low cost public wireless Internet services and have begun encouraging local governments to act.

WIMAX

WiMAX is the commercial name for a set of standards developed by the *IEEE 802.16* standards group. The 802.16 family of technologies is much like the 802.11 family and the Ethernet family. They reuse many of the Ethernet 802.3 components and are designed to connect easily into Ethernet LANs. There are two primary types of WiMAX: fixed and mobile.

MANAGEMENT FOCUS **7-3 PUBLIC WI-FI IN TEMPE, ARIZONA**

The city of Tempe, home to Arizona State University, has become one of the early leaders in the provision of public Wi-Fi Internet access. Working with several commercial providers, Tempe installed an outdoor Wi-Fi network covering 95 percent of the city's 40 square miles. The neighboring cities of Chandler and Gilbert have also joined the project, meaning that the network eventually will cover 187 square miles. The network is built with 802.11g, meaning that all 802.11g and 802.11b devices can use it.

The network offers access to residents and visitors on an annual, monthly, or daily basis. There is a zone in the merchant district of downtown Tempe that offers free access. Access to City of Tempe and Arizona State University Web sites is also free, regardless of access location.

SOURCE: "City-wide Wi-Fi Project," Tempe City Government, www.tempe.gov/business/wifi; and www.waztempe.com.

Topology

The logical and physical topologies of wireless Ethernet are the same as those of 802.11 and shared Ethernet. They are a physical star and a logical bus. There is a central AP to which all computers direct their transmissions (star), but the radio frequencies are shared (bus) so that all computers must take turns transmitting.

Media Access Control

Unlike Ethernet, media access control for WiMAX is controlled access, using a version of the 802.11 point coordination function (PDF).

Types of WiMAX

There are two types of WiMAX.

802.16d The IEEE 802.16d standard covers fixed point wireless access, using antennas that are 12-18 inches in size. The goal is to provide wireless connections between one central access point and a set of fixed networks. The most common use of this standard is to connect a set of offices to a central office without using traditional WAN connections (which are discussed in Chapter 9). Under ideal conditions, 802.16d provides 70 Mbps data rates for up to 30 miles. Real world tests of this technology, however, suggest that the maximum effective distance, given the noisy radio frequency ranges it uses, is 5 miles, with effective data rates of 2 Mbps.

A growing use for 802.16d is to connect multiple Wi-Fi public access points to a central switch, so they can connect into the Internet. This eliminates the need to put in wires and enables a quick rollout of new technology.

802.16e The IEEE 802.16e standard is intended to provide access for mobile users in competition to outdoor Wi-Fi. It provides multiple channels, each with 28 Mbps, although

the effective data rate is about 5 Mbps. The effective range is up to 6 miles with a line of sight to the access point or 2.5 miles without a line of sight.

802.16e is a direct competitor to public access Wi-Fi and current cell phone technologies but is incompatible with both. Manufacturers will have to build in separate 802.16e chips and antennas into phones and laptops or users will need to purchase add-on 802.16e NICs.

BLUETOOTH

Bluetooth is the commercial name for the IEEE 802.15 standards, which calls it a *Wireless Personal Area Network* (*WPAN*). In case you're wondering, Bluetooth's Scandinavian inventor decided to name it after Danish King Harold Bluetooth.

Bluetooth is a strikingly different type of wireless LAN from the others discussed in this chapter. It is not intended as a general-purpose network in competition with 802.11 or 802.16 wireless LANs or 802.3 wired LANs. Its goal is to provide seamless networking of data and/or voice devices in a very small area (up to 10 meters or 30 feet, possibly to increase to about 100 meters or 300 feet with the next generation of technology). Bluetooth can be used to connect many different types of devices, such as keyboards to computers and headsets to mobile phones.

Bluetooth devices are small (about one-third of an inch square) and inexpensive. They are designed to replace short-distance cabling between devices such as keyboards, mice, and a telephone handset and base or to link your PDA to your car so that your door can unlock and automatically open as you approach. Bluetooth provides a basic data rate of 1 Mbps that can be divided into several separate voice and data channels.

Topology

A Bluetooth network is called a *piconet* and consists of no more than eight devices, but can be linked to other piconets to form a larger network. One device is considered the piconet *master*, and all other devices are *slaves*. The master controls the piconet, selecting frequencies and access control used by the master and the slaves. All messages are sent from a slave to the master and from the master to a slave. The slaves do not communicate directly. All devices share the same frequency ranges so the network behaves in the same manner as a shared bus topology.

Media Access Control

The master uses a controlled access technique similar to Wi-Fi's PCF approach. Bluetooth uses *frequency-hopping spread-spectrum (FHSS)* in which the 2.4 GHz frequency range is divided into 79 separate channels. Each channel is used in turn to transmit signals. A short burst of data is sent on one frequency and then the sender changes to another frequency channel and broadcasts another burst of data before changing to another channel. There are usually 1600 channel changes (called *hops*) per second. The master controls which channels will be used, so the master and the slave with which it is communicating are synchronized and both know which frequencies will be used at which point. This approach

also minimizes interference because if one frequency channel suffers from interference, it will soon be avoided.

Because Bluetooth operates in the same 2.4-GHz range as Wi-Fi, it has the potential to cause problems for Wi-Fi WLANs. Tests suggest that good management can prevent interference between Bluetooth and Wi-Fi. As long as no Bluetooth piconets are located within 2 meters of a Wi-Fi NIC or AP and as long as only a moderate number of Bluetooth piconets are operating in the same area as a Wi-Fi WLAN, neither the Bluetooth piconets nor the Wi-Fi WLAN appear to suffer any problems.

THE BEST PRACTICE WLAN DESIGN

As with the best practice LAN design, our recommendations for the best practice WLAN design are based primarily on the trade-off between effective data rates and costs. WiMAX and Bluetooth are not intended to be used for general networking, so we do not include them in our discussions here. Because WLANs are competitors for traditional wired LANs, we also consider the issue of LAN versus WLAN, which is perhaps the more interesting question. We also discuss the physical design of WLANs because the design can be challenging.

Effective Data Rates

As you will recall, the effective data rates of the lower network layers are the maximum practical speeds in bits that the hardware layers can be expected to provide and depend on four basic factors: nominal data rates, error rates, efficiency of the data link layer protocols used, and efficiency of the media access control protocols. Error plays a greater role in WLANs than it does in wired LANs because interference can significantly affect performance by increasing the number of retransmissions and by forcing the WLAN to drop to a slower data rate. In this analysis, we will make the *major* assumption that the APs have been well placed so that all users attempting to work on the WLAN have good signal quality and are able to operate at the maximum nominal data rate provided by the WLAN: 11 Mbps for 802.11b, 54 Mbps for 802.11a and 802.11g, and 200 Mbps for 802.11n.

Data Link Protocol Efficiency Wi-Fi uses data link layer protocols similar to those used by their wired Ethernet cousins (e.g., 100Base-T, 1000Base-T). Wi-Fi packets have a typical overhead of 51 bytes (if a short preamble is used) on 1500-byte packets, plus the ACK/NAK. However, this calculation is complicated by the fact that many of the overhead bits are transmitted at the slowest data rate, not at the maximum data rate. Assuming we have the same mix of short and full length packets and without going into all the calculations, the efficiency for 802.11b is about 85 percent and the efficiency of 802.11a, 802.11g, and 802.11n is about 75 percent.

Media Access Control Protocol Efficiency The next factor is the efficiency of the media access control protocols. Wi-Fi uses a very different media access control protocol from wired Ethernet's CSMA/CD. Chapter 6 discussed the performance charac-

teristics of CSMA/CD: gradual increases in response time delay to about 50 percent of nominal capacity, more rapid increases in delay to about 80 percent of capacity, and immense increases in delays after 80 percent that rendered the network essentially unusable.

Wi-Fi uses the PCF controlled-access technique. PCF initially imposes more fixed cost delays when traffic is low because computers must request permission before they transmit rather than just making certain there is no traffic and transmitting at will as with CSMA/CD. However, response time delays increase slowly up to about 85 to 90 percent of nominal capacity because collisions are effectively eliminated. Once this level is reached, they increase rapidly until the network is 100 percent saturated.

Wi-Fi users experience few response time delays as long as the total amount of network traffic remains under 85 to 90 percent of the nominal data rate. This means, for example, that a 802.11b WLAN with a nominal data rate of 11 Mbps can provide an effective total data rate of about 9.6 Mbps, assuming that there is no substantial interference (85 percent efficiency x 85 percent capacity x 11 Mbps = 9.6 Mbps). This capacity is shared by all computers on the WLAN, so if we had a low-traffic network with only two active computers on the one 802.11b AP, this would mean that, on average, each computer could realistically use about 5 Mbps—under perfect operating conditions. As the number of active computers increases, the average capacity drops. Under more normal operating conditions, effective data rates are also lower. Figure 7.6 shows some estimated effective data rates for Wi-Fi.

Costs

802.11g WLAN NICs and APs are modest in cost, and prices are rapidly dropping. As 802.11n products are newer, the costs are higher but should drop over the next year or two. The cost of an 802.11b AP is a bit more than a 10/100Base-T switch. Most laptops have both 802.11g and wired Ethernet NICs built-in, while more desktops only come with Ethernet NICs. The cost of an 802.11b NIC for a desktop is about $40. However, the largest cost associated with wired Ethernet LANs is not the cost of the NICs, hubs, or switches. The largest cost is the cost of installing the cables. Installing a cable can cost anywhere from $10 to $400 per cable, depending upon the condition of the building in

Technology	Operating Conditions	Low Traffic	Moderate Traffic	High Traffic
802.11a	Perfect	17 Mbps	7 Mbps	3 Mbps
	Normal	11 Mbps	5 Mbps	2 Mbps
802.11b	Perfect	5 Mbps	2 Mbps	1 Mbps
	Normal	3 Mbps	1 Mbps	500 Kbps
802.11g	Perfect	17 Mbps	7 Mbps	3 Mbps
	Normal	11 Mbps	5 Mbps	2 Mbps
802.11n	Perfect	68 Mbps	28 Mbps	12 Mbps
	Normal	44 Mbps	20 Mbps	8 Mbps

FIGURE 7.6 Effective data rate estimates for Wi-Fi.

which the cable is to be installed. It is less expensive to install cable during the construction of a new building and much more expensive to install cable after the fact in an old building.

Thus for new construction, wired LANs are less expensive than their wireless counterparts, but only by a modest amount. For installation in an existing building that lacks cabling, 802.11g WLANs may be less expensive than wired LANs.

Recommendations

There is, of course, one other major factor: mobility. Wi-Fi LANs provide the ability for computers and employees to move seamlessly throughout an indoor or outdoor area and to work in locations that wires cannot reach.

Given the trade-offs in costs and effective data rates, and the importance of mobility, there are several best practice recommendations. First, it is becoming clear that 802.11g will replace both 802.11a and 802.11b. Thus, our recommendation for WLAN design today is to adopt 802.11g. If manufacturers price the new 802.11n equipment as aggressively as initial reports suggest, then 802.11n should move very quickly into the marketplace and become the preferred technology.

Most interesting, perhaps, is the relationship between Wi-Fi and wired Ethernet. The data rates for Wi-Fi are similar to the effective data rates for wired Ethernet networks (see Chapter 6). For most networks, the wired 100Base-T recommended previously still provides the best trade-off between cost and performance. But Wi-Fi networks are a very close competitor for low-traffic environments. In cases where mobility is important or wiring is expensive, Wi-Fi may be the best practice.

Many organizations today are still installing traditional wired networks but are using WLANs as *overlay networks*. They build the usual switched Ethernet networks as the primary LAN, but also install WLANs so that employees can easily move their laptops in and out of the offices and to provide connectivity in places not normally wired such as hallways and lunch rooms.

Physical WLAN Design

We will discuss the general principles for network design in Chapter 12, but in this section we discuss some of the issues specific to the design of WLANs. Designing the physical WLAN is more challenging than designing a traditional LAN because the potential for interference means that extra care must be taken in the placement of access points. With the design of LANs there is considerable freedom in the placement of hubs and switches, subject to the maximum limits to the length of network cables. In WLANs, however, the placement of the access points needs to consider both the placement of other access points as well as the sources of interference in the building.

The physical WLAN design begins with a *site survey*. The site survey determines the feasibility of the desired coverage, the potential sources of interference, the current locations of the wired network into which the WLAN will connect, and an estimate of the number of APs required to provide coverage. While the site survey may uncover unexpected sources of interference (e.g., cordless telephones, microwave ovens, industrial equipment), the most common sources of interference are walls. WLANs work very well

when there is a clear line of sight between the AP and the wireless computer. The more walls that exist in the environment, the more the wireless signal needs to penetrate the walls and thus the weaker it becomes. The type and thickness of the wall also has an impact; traditional drywall construction provides less interference than does concrete block construction.

Although it is possible to calculate the probable range of an AP given the type of construction and the number of walls in a building, in many cases the site survey is done using a temporary AP and a computer or device that can actually measure the strength of the wireless signal. The temporary AP is installed in the area to be surveyed, and the computer or device is carried throughout the building measuring the strength of the signal. Actually measuring the strength of the signal in the environment is far more accurate than relying on estimated ranges from the vendor. The site survey will also locate the placement of power sources and the existing wired network because the AP will need power and in most cases will be connected into the wired network so that the WLAN can communicate with the rest of the network.

The design of the WLAN is simple if one AP is sufficient to cover the desired area. However, if the area is large enough to require several APs, then the design becomes more complicated. The simplest approach is to start in one corner of the coverage area and place one AP in what seems to be a good location. Then the strength of the signal is measured by walking through the area to determine the farthest point of coverage for the desired signal strength. You may have to move the AP several times until you find the placement that provides the best coverage for the corner area with little "wasted" signal outside the desired area of coverage. The exact placement of the AP depends on the environment and the type of antenna. While omnidirectional antennas are the most common, directional antennas can also be used.

This process is repeated starting in each of the different corners of the area to be covered. Once the corners have been surveyed, you begin filling in the empty coverage areas in the middle by repeating the same process.

In the above paragraphs our aim has been to design the network to provide the "desired signal strength." The signal strength determines the maximum data rate possible in the WLAN. Under ideal circumstances and if cost is not an issue, many APs will be purchased so that they can be placed close together to provide a strong signal strength that results in a data rate close to the maximum data rate provided by the AP. In general, a 15 percent overlap in coverage between APs at the desired signal strength is sufficient to provide smooth and transparent roaming from AP to AP. Each AP is set to transmit on a different wireless channel so that the APs do not interfere with each other. If cost is an issue, fewer APs will be available, and they will need to be placed farther apart to provide a lower signal strength (and slower data rates) at extreme ranges. There may even be some dead spots in less important areas.

Design becomes more difficult in a multistory building because the signals from the APs travel up and down as well as in all horizontal directions. The design must include the usual horizontal mapping but also an added vertical mapping to ensure that APs on different floors do not interfere with one another (Figure 7.7). It becomes even more difficult if your building or set of floors in a large office tower is surrounded by APs of other companies. You have to design your network not to interfere with theirs.

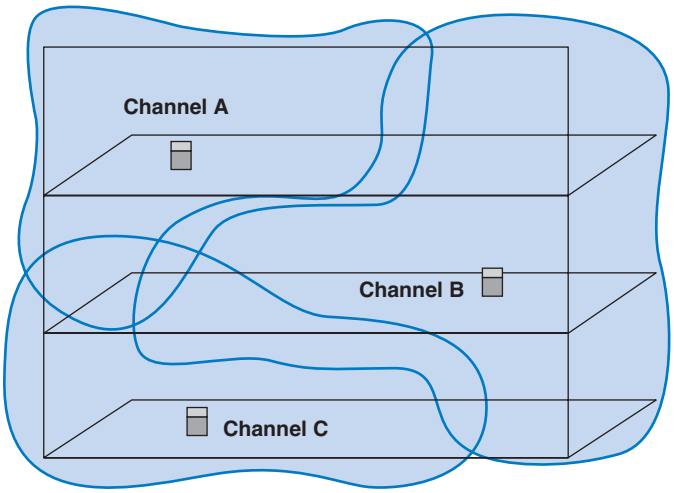

FIGURE 7.7 Multistory WLAN design.

WLAN Security

Security is important to all networks and types of technology, but it is especially important for wireless networks. In a traditional wired network such as a LAN, the only way to connect to the network is to enter the offices, find a network connection, and plug into the network. With a WLAN, anyone walking or driving within the range of an AP (even outside the offices) can begin to use the network.

Finding WLANs is quite simple. You just walk or drive around different office buildings with your WLAN-equipped client computer and see if it picks up a signal. There are also many special-purpose software tools available on the Internet that will enable you to learn more about the WLANs you discover, with the intent of helping you to break into them. This type of wireless reconnaissance is often called *wardriving* (see www.wardriving.com). *Warchalking* refers to the practice of writing symbols in chalk on sidewalks and walls to indicate the presence of an unsecured WLAN (see www.warchalking.org).

SSID The most basic security applied to WLANs is to require all client computers wanting to access an AP to include a *Service Set Identifier* (*SSID*) in all packets. Any packet with the incorrect SSID is not processed by the AP. This provides very basic security but it is easy to break. The SSID is included in all packets in plain text, so any device within range of the AP that has the right software can listen to packets and easily read the SSID they contain. Simply put, using SSID does not provide security.

WEP Another type of wireless security is *Wired Equivalent Privacy* (*WEP*). With WEP, the AP requires the user to have a *key* in order to communicate with it. All data sent to and from the AP is encrypted so that it can only be understood by computers or devices

that have the key.[4] If a computer does not have the correct WEP key, it cannot understand any messages transmitted by the access point and the access point will not accept any data that is not encrypted with the correct key. Encryption is discussed in detail in Chapter 11.

One of the problems with WEP is that the key must be manually typed into the client computer and into the AP. While this is not a major problem in a small WLAN, it does become challenging for large WLANs. Imagine the management time required when a WEP key needs to be changed in an organization with dozens of APs and hundreds of client computers (or hundreds of APs and thousands of computers).

With *Extensible Authentication Protocol* (*EAP*), the WEP keys are produced dynamically, much like the way in which a DHCP server is used to dynamically produce IP addresses. When an AP using EAP first discovers a new client computer, it requires the user to login before it will communicate with the client computer. The userid and password supplied by the user are transmitted to a login server, and if the server determines that they are valid the server generates a WEP key that will be used by the AP and client computer to communicate for this session. Once the client logs out or leaves the WLAN, the WEP is discarded and the client must login again and receive a new WEP key.

WEP has a number of serious weaknesses, and most experts agree that a determined hacker can break into a WLAN that uses only WEP security. A good way to think about WEP is that it is like locking your doors when you leave: it won't keep out a professional criminal but it will protect against a casual thief.

WPA *Wi-Fi Protected Access* (*WPA*) is a newer, more secure type of security. WPA works in ways similar to WEP and EAP: every packet is encrypted using a key, and the key can be fixed in the AP like WEP or can be assigned dynamically as users login like EAP. The difference is that the WPA key is longer than the WEP key and thus is harder to break. More importantly, the key is changed for *every packet* that is transmitted to the client. Each time a packet is transmitted, the key is changed.

802.11i *802.11i* is the newest, most secure type of WLAN security. It uses EAP to obtain a master key—in other words, the user logs in to a login server to obtain the master key. Armed with this master key, the user's computer and the AP negotiate a new key that will be used for this session until the users leaves the WLAN. 802.11i uses the Advanced Encryption Standard (AES) discussed in Chapter 11 as its encryption method.

IMPROVING WLAN PERFORMANCE

Improving the performance of WLANs is similar to improving LAN performance. You check the devices in the network (i.e., clients, and APs), the wireless circuits between the computers, and the demand placed on the network.

[4] WEP uses single-key encryption with a 40-bit or 128-bit key length. Only the data payload is encrypted (i.e., the data portion of the LLC PDU in Figure 7.4).

MANAGEMENT FOCUS

7-4 MOOCHING WI-FI

If you connect into someone else's Wi-Fi network and start using their Internet connection are you:

a. guilty of stealing from the owner because you haven't paid them
b. guilty of stealing from the ISP because you haven't paid them
c. committing an unethical but not illegal act
d. really frugal, and not unethical
e. all of the above

According to the St Petersburg, Florida police department, the answer is *a*. They arrested a man named Benjamin Smith for "willfully, knowingly, and without authorization" accessing the network of a homeowner while sitting in a car parked on the street.

According to Verizon and most ISPs, which explicitly prohibit sharing, the answer is *b*. "It's obviously not good for Verizon to have its services given away for free, just as a cable company won't want someone funneling their cable connection next door," said a Verizon spokeswoman.

According to Miss Manners, the answer is *c*. It's not nice to use other people's stuff without asking their permission.

According to Jennifer Granick, executive director of the Center for Internet and Society at Stanford Law School, the answer is *d*. "Such use [i.e., sharing] might be allowed or even encouraged [by the owner]." Unless the owner states you can't enter their network, how do you know you're not invited?

As Lee Tien, a senior staff attorney at the Electronic Frontier Foundation says "Right now, we don't have a way of saying 'Even though my wireless signal is open, I'm saying you can't use it.'" Until we do, the answer is *e*. So, tread carefully. Don't leave your WLAN unsecured or you may be legally inviting others to use it as well as your Internet connection. Likewise, don't intentionally enter someone else's WLAN and use their Internet connection or you might end up like Benjamin Smith—spending the night in jail.

SOURCE: John Cox, "Mooching Wi-Fi," *Network World*, August 8, 2005, pp. 1, 49.

Improving Device Performance

As we discussed earlier, the presence of one single computer using 802.11b to communicate with an 802.11g AP will reduce the performance of all 802.11g devices using the same WLAN because the AP slows down the 802.11g traffic so it does not confuse the 802.11b device. Therefore, if WLANs with 802.11g are widely deployed in your organization and most but not all computers use 802.11g cards, it may be possible to significantly improve performance by replacing the few remaining 802.11b cards with newer 802.11g cards.

Not all wireless cards and APs are created equal, despite the move to standardization. Some devices are better designed and thus have a stronger signal at longer ranges. Thus, sometimes performance can be improved by buying high-quality wireless cards and APs from a vendor with a reputation for quality.

Improving Circuit Capacity

The simplest way to improve circuit capacity is to upgrade from 802.11a or 802.11b to 802.11g or 802.11n. The faster speeds at greater range should enable computers to quickly see the improved performance.

Reexamining the exact placement of APs is another potential way to improve performance. APs should be placed in an area with the fewest walls between the AP and the devices on the WLAN. This means that most APs will be mounted on ceilings or high up on walls so they can transmit over the top of cubicles and other obstructions. It may be possible to significantly improve performance by placing an AP on a corridor wall rather than in a special-purpose networking closet.

If performance is significantly worse than expected, then it is important to check for sources of interference near the AP and the computers. Bluetooth devices are one source of problems for 802.11b and 802.11g devices. Cordless phones (and baby monitors) also may operate on the same frequency ranges as all three 802.11 standards (2.4 GHz and 5 GHz), so it may be necessary to remove these devices for the WLAN to operate effectively.

Another option is to try different styles of antennas for the AP. Directional antennas focus the radio energy in a smaller range of direction and therefore can produce a stronger signal (with faster throughput) at longer ranges than can omnidirectional antennas. There are also several different styles of both directional and omnidirectional antennas that may better suit different environments.

Reducing Network Demand

One of the most important design rules for improving WLAN performance is never to place a server in a WLAN. All 802.11 WLANs require that all communication is between the individual device and the AP. Therefore, if a server is placed in the WLAN all messages sent from client computers in the WLAN to the server must be sent twice: once from the client to the AP and a second time from the AP to the server. Therefore, performance in the WLAN will be improved if the server is located in the wired portion of the same LAN as the AP (ideally a switched Ethernet LAN) because this will significantly reduce the traffic on the WLAN.

TECHNICAL FOCUS 7-1 INTERFERNCE AT INDIANA UNIVERSITY

Most of the buildings at Indiana University have both wired and wireless network access. The Kelly School of Business at Indiana University has two major buildings: a modern building built in 2002 and an older building built in 1968. The new building was designed with wireless networks in mind; the old building was not. My office is in the old building.

We have one Wi-Fi access point on our floor which should provide sufficient coverage for the small office tower in which we are located. However, the walls are made of concrete which is hard for wireless signals to penetrate. Figure 7.8 shows the floor plan, the position of the AP, and the data rates that are available at different locations on the floor.

My office is located about 35 feet from the AP (less than 12 meters), which is well within the normal range for high speed access. However, because of the concrete walls, I am unable to receive a signal in most of my office.

SOURCE: "802.11g Starts Answering WLAN Range Questions," www.commsdesign.com, 2004.

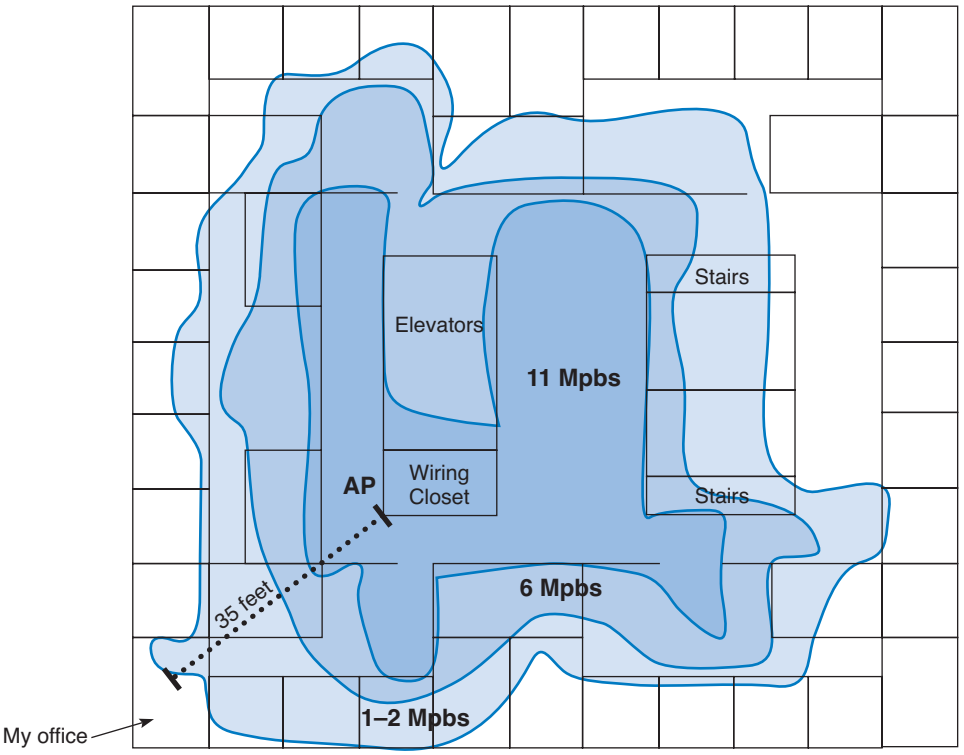

FIGURE 7.8 WLAN coverage on one floor of Indiana University's Kelley School of Business.

WLANs are most commonly used as overlay networks; they supplement existing wired LANs and are intended to be used primarily by mobile users with laptop computers. It is possible to reduce demand on the WLAN by placing wired LAN jacks in commonly used locations. For example, if there are tables or couches in a lounge that is covered by a WLAN, most mobile users will naturally sit there and use the WLAN. If response times of the WLAN become a problem, users can plug their laptops into a nearby Ethernet wall jack if offered the opportunity, thus reducing the demand on the WLAN.

IMPLICATIONS FOR MANAGEMENT

As WLANs become commonplace in organizations, accessing organizational networks or the Internet will become routine. Offices, cafeterias, break rooms, and external courtyards will be turned into wireless hotspots. Mobile workers will have access to any data, any time, and any place.

Public access wireless hotspots will become commonplace as people will come to expect the same wireless access in restaurants, malls, and courtyards as they expect in

their organizations. WLAN technology will begin to compete with traditional cell phone technology, and providers of cell phones will begin to develop and install new versions of WLAN technologies that have longer ranges.

As WLANs become widely adopted, prices will begin to drop in the same way that the costs of LAN technologies have dropped. Wireless technology will become standard in a multitude of new devices (e.g., handheld computers, shopping carts, door locks) and locations (e.g., city streets, parks, your car). We think the Internet is widespread today; in 5-10 years, it will be truly ubiquitous.

These changes will result in the development of a variety of new Internet applications designed to provide real-time data to consumers in organizations. Entirely new industry segments will be created and businesses will be created and destroyed. This also means that the amount of data flowing through organizational networks and the Internet will continue to grow at its current dizzying pace.

SUMMARY

WLAN Components The NIC is a small radio transmitter/receiver that enables a computer to transmit to and receive from the access point. The access can have a directional or omnidirectional antenna and is usually wired into a traditional wired network. Most WLANs operate in the 2.4 GHZ and 5 GHZ frequency ranges and transmit 100–500 feet.

Wi-Fi Wi-Fi is the most common type of WLAN. It uses physical star/logical bus topology with both controlled and contention-based media access control. 802.11a provides data rates up to 54 Mbps over short distances, while 802.11b provides data rates up to 11 Mbps over longer distances (up to 500 feet). 802.11g is designed to replace both of these by providing 54 Mbps over longer distances. 802.11n is designed to provide higher data rates over these same distances. Both 802.11g and 802.11n are backwards compatible, which means that they can be used with the older standards.

WiMAX WiMAX is designed to provide outdoor 70 Mbps data access over long distances, up to 30 miles, although most real world tests suggest data rates of 5 Mbps up to 6 miles is more common. 802.16d is fixed wireless WiMAX connecting multiple buildings to one center access point, while 801.16e is designed to provide access for mobile users. WiMAX is designed to replace outdoor public access Wi-Fi, but it is unclear which technology will win the battle.

Bluetooth Bluetooth is strikingly different from the other WLANs because its goal is to provide networking of data and/or voice devices in a very small area (up to 10 meters). It is designed to replace short-distance cabling between devices such as keyboards, mice, and a telephone handset. Bluetooth provides a basic data rate of 1 Mbps in the same 2.4-GHz bandwidth as Wi-Fi, but initial tests suggest that there is little interference between Bluetooth and Wi-Fi LANs provided they are not within 2 meters of each other.

Best Practice WLAN Design If mobility is important, Wi-Fi is a viable option to wired LANs. Given the trade-offs in costs and effective data rates, the best LAN for most networks is still the traditional wired LAN discussed in the previous chapter. However, as Wi-Fi becomes more mature, it will provide serious competition.

WLAN Security WLAN security is important because unlike wired LANs, anyone walking or driving by your home or office can connect unnoticed to your WLAN. Two popular approaches to WLAN security, SSID and WEP, provide some security, but neither will stop a determined hacker who knows their weaknesses. Newer security techniques, such as WPA and 802.11i, provide significantly better security.

Improving WLAN Performance WLAN performance can be improved by using name-brand equipment and by ensuring no 802.11b devices operate with an 802.11g AP because just one 802.11b device will slow down the entire WLAN. Performance can also be improved by moving to 802.11g and 802.11n, placing the APs so that fewer walls obstruct their transmission, removing interference (e.g., cordless phones), and switching to more powerful antennas. Network demand can be reduced by ensuring that no servers are placed on the WLAN and by placing additional wired LAN jacks near commonly used locations.

KEY TERMS

access point (AP)
bandwidth
Bluetooth
bus topology
channel
clear to transmit (CTS)
collision
collision avoidance (CA)
contention
CSMA/CA
data rate
directional antenna
distributed coordination function (DCF)
extensible authentication protocol (EAP)
frequency
frequency range
master
omnidirectional antenna
overlay network
physical carrier sense method
piconet
point coordination function (PCF)
request to transmit (RTS)
roaming
service set identifier (SSID)
site survey
slave
virtual carrier sense method
warchalking
wardriving
Wi-Fi
Wi-Fi protected access (WPA)
WiMAX
Wired Equivalent Privacy (WEP)
Wireless LAN (WLAN)
Wireless Personal Area Network (WPAN)
802.11a
802.11b
802.11g
802.11i
802.11n
802.15
802.16d
802.11e

QUESTIONS

1. Describe the basic components of a wireless network.
2. How do the NIC and AP work together to transmit messages in an 802.11b WLAN?
3. Compare and contrast the two types of antennas.
4. What are two ways in which an omnidirectional antenna differs from a directional antenna?
5. How does Wi-Fi perform media access control?
6. What are the types of Wi-Fi?
7. How does 802.11g differ from 802.11b and 802.11a?

8. What data rates are provided by the different types of Wi-Fi?
9. How does Wi-Fi differ from shared Ethernet in terms of topology, media access control, and error control?
10. How does roaming work?
11. Explain how CSMA/CA DCF works.
12. Explain how CSMA/CA PCF works.
13. How do the effective data rates for Wi-Fi technologies compare to their nominal data rates?
14. Explain the topology and media access control of WiMAX.
15. Compare and contrast the two types of WiMAX.
16. Is WiMAX a competitor to Wi-Fi? Explain.
17. Which type of WiMAX do you think has the greatest future prospects? Why?
18. How does a WPAN differ from a WLAN?
19. Explain the topology and media access control of Bluetooth.
20. What are the best practice recommendations for WLAN design?
21. What is a site survey and why is it important?
22. How do you decide how many APs are needed and where they should be placed for best performance?
23. How much overlap should be provided between APs? Why?
24. Why is security important for WLANs?
25. What are wardriving and warchalking?
26. Explain how SSID works.
27. Explain how WEP works.
28. Explain how EAP works.
29. Explain how 802.11i works.
30. Are today's WLANs secure? Explain.
31. What do you think WLAN security will look like in 3 years?
32. Some people believe Bluetooth is a revolution while others see it as a simple replacement for cables among devices. What do you think? Is Bluetooth a revolution?
33. Given the dramatic changes ahead in WLANs (e.g., IEEE 802.11), would you install a WLAN today? Explain.
34. If IEEE 802.11n is widely available in the next few years, what are the implications for networks of the future? Will 100Base-T still be around or will we eliminate wired offices?
35. Many of the wired and wireless LANs share the same or similar components (e.g., error control). Why?
36. What do you think are the future prospects for Wi-Fi versus WiMAX? Why?
37. What do you think the future is for public access Wi-Fi? Should towns and cities be encouraged to build or be prohibited from building such networks?

EXERCISES

7-1. Survey the WLANs used in your organization. What types of Wi-Fi and/or WiMAX are in use?

7-2. You have been hired by a small company to install a simple WLAN for their 18 Windows computers. Develop a simple WLAN and determine the total costs; that is, select AP and NICs and price them.

7-3. Investigate the current state of wireless security including ideas moving through the IEEE standards process.

7-4. If you live in a large city, explore the downtown area for warchalking. Take pictures and bring them to class.

MINI-CASES

I. General Hospital

General Hospital has five floors, each about 30,000 square feet in size, for a total of about 150,000 square feet. They want to provide a wireless overlay network in addition to their switched 100Base-T. They have a bid for 802.11g access points at a cost of $100 each and a bid for 802.11n access points at a cost of $300 each. They expect to need 200 NICs. 802.11b NICs come built into their laptops and tablets. 802.11n NICs cost about $100 each. What would you recommend? Why?

II. Central University

Central University wants to add a wireless overlay network to one 20,000 square-foot floor in its business school. They have a bid for 802.11g access points at a cost of $100 each and a bid for 802.11n access points at a cost of $300 each. Students will buy their own computers, most of which will come with 802.11g NICs. 802.11n NICs cost about $100 each (with "discount brands" selling for $85). What would you recommend? Why?

III. South West State University

South West State University installed a series of four Wi-Fi omnidirectional APs spread across the ceiling of the main floor of their library. The main floor has several large, open areas plus two dozen or so small offices spread around the outside walls. The WLAN worked well for one semester, but now more students have laptops with Wi-Fi built in, and performance has deteriorated significantly. What would you recommend that they do? Be sure to support your recommendations.

IV. Household Wireless

Your sister is building a new two-story house (which measures 50 feet long by 30 feet wide) and wants to make sure that it is capable of networking her family's three computers together. She and her husband are both consultants and work out of their home in the evenings and a few days a month (each has a separate office with a computer, plus a laptop from the office that they occasionally use). The kids also have a computer in their playroom. They have several options for networking their home:

a. Wire the two offices and playroom with Ethernet cat 5e cable and put in a 100Base-T switch for $40
b. Install one Wi-Fi access point ($85) and put Wi-Fi cards in the three computers for $70 each (their laptops already have Wi-Fi)
c. Any combination of these options.

What would you recommend? Justify your recommendation.

CASE STUDY

NEXT-DAY AIR SERVICE

See the Web site.

HANDS-ON ACTIVITY

War-Driving and War-Walking

Wireless LANS are often not secure. It is simple to bring your laptop computer into a public area and listen for wireless networks. This is called War-Driving (if you are in a car) or War-Walking (if you're walking). As long as you do not attempt to use any networks without authorization, War-Driving and War-Walking are quite legal. There are many good software tools available for War-Driving. My favorite is Net Stumbler. It is simple to use, yet powerful.

The first step is to download and install the Net Stumbler software on a laptop computer that has wireless capability. The software is available at www.netstumbler.com. Once you have installed the software, simply walk or drive to a public area and start it up. Figure 7.9 shows an example of the seven networks I discovered in my home town of Bloomington, Indiana when I walked through one building downtown. For each network, Net Stumbler displays the MAC address of the access point (or physical address if you prefer to use that term). It shows the SSID, the channel number the AP is configured to use, the speed of the network, the access point vendor (which can be disabled by the access point owner to increase security), and the type of encryption in use (if any). It also shows the signal strength both by color coding the network (green is good) and by showing the signal-to-noise ratio (SNR) and the strength of the signal and the noise.

In Figure 7.9, you can see a mix of WLANs, both 11 Mbps and 54 Mbps. The channels we usually use for 802.11b and 802.11g are channels 1, 6, and 11. In this figure, you'll see a mix of channels 1 and 6, plus one channel 8 WLAN. 802.11b and 802.11g can be configured to use four channels (1, 4, 8, and 11), although the channels overlap to some extent. So if you run an AP on channel 1 and another on channel 4, there will some interference between the two APs. The best practice recommendation that most companies follow is to use a three-channel configuration. In this building, you can see that most companies are using the three-channel configuration, but one is not; it's using the four-channel configuration.

If you click on an access point in the left panel, Net Stumbler shows you a real time graph of the signal and noise for that network. Figure 7.10 shows how the signal strength changed for one of the networks as I walked through the building. The left edge of the graph shows that the network started with a good signal (the green or light colored area at the top of the bars) was much higher than the noise (the red or dark colored area at the bottom of the bars). As I walked around, the signal became weaker; the signal was barely higher than the noise. As I walked more, the signal dropped so that it was too weak for me to detect it from the noise.

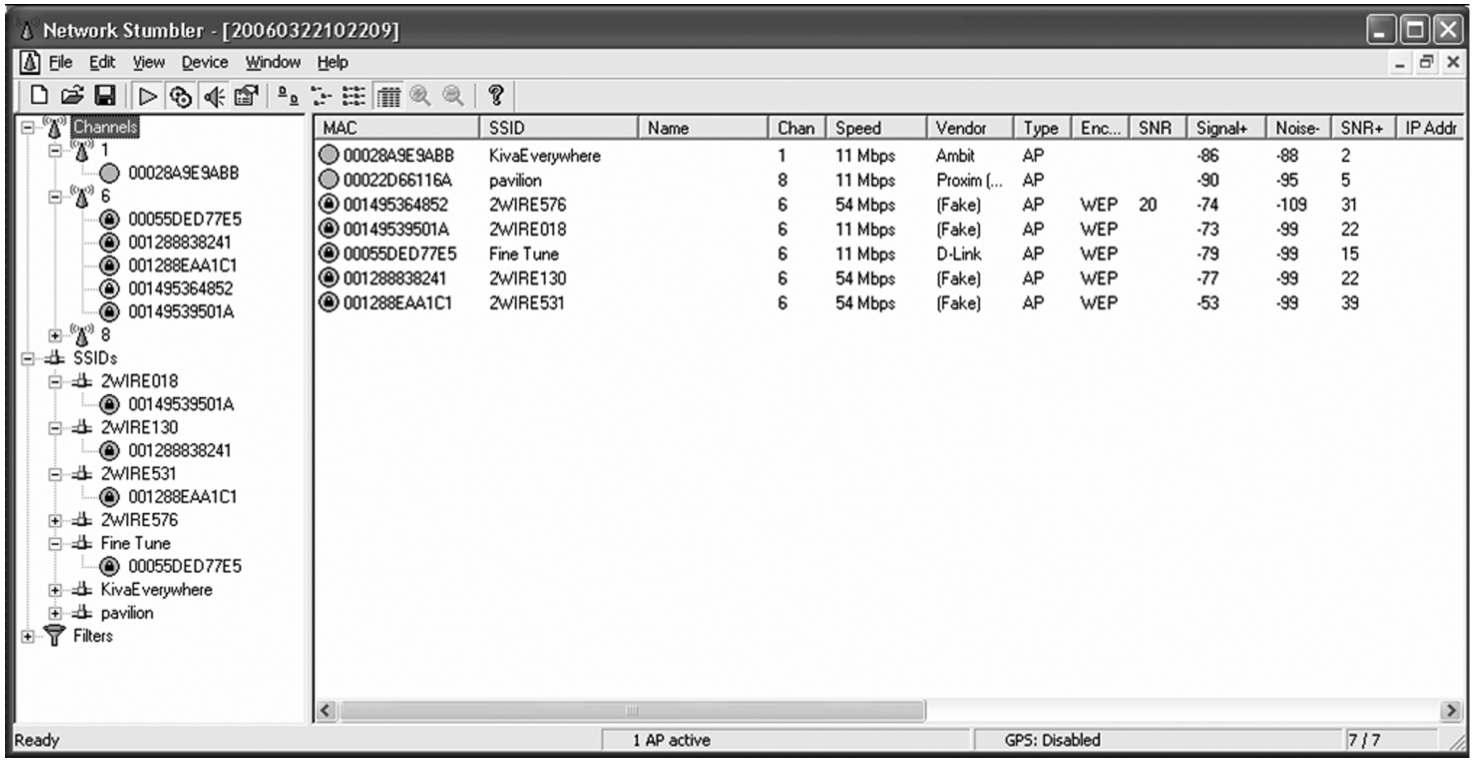

FIGURE 7.9 Networks in one downtown building.

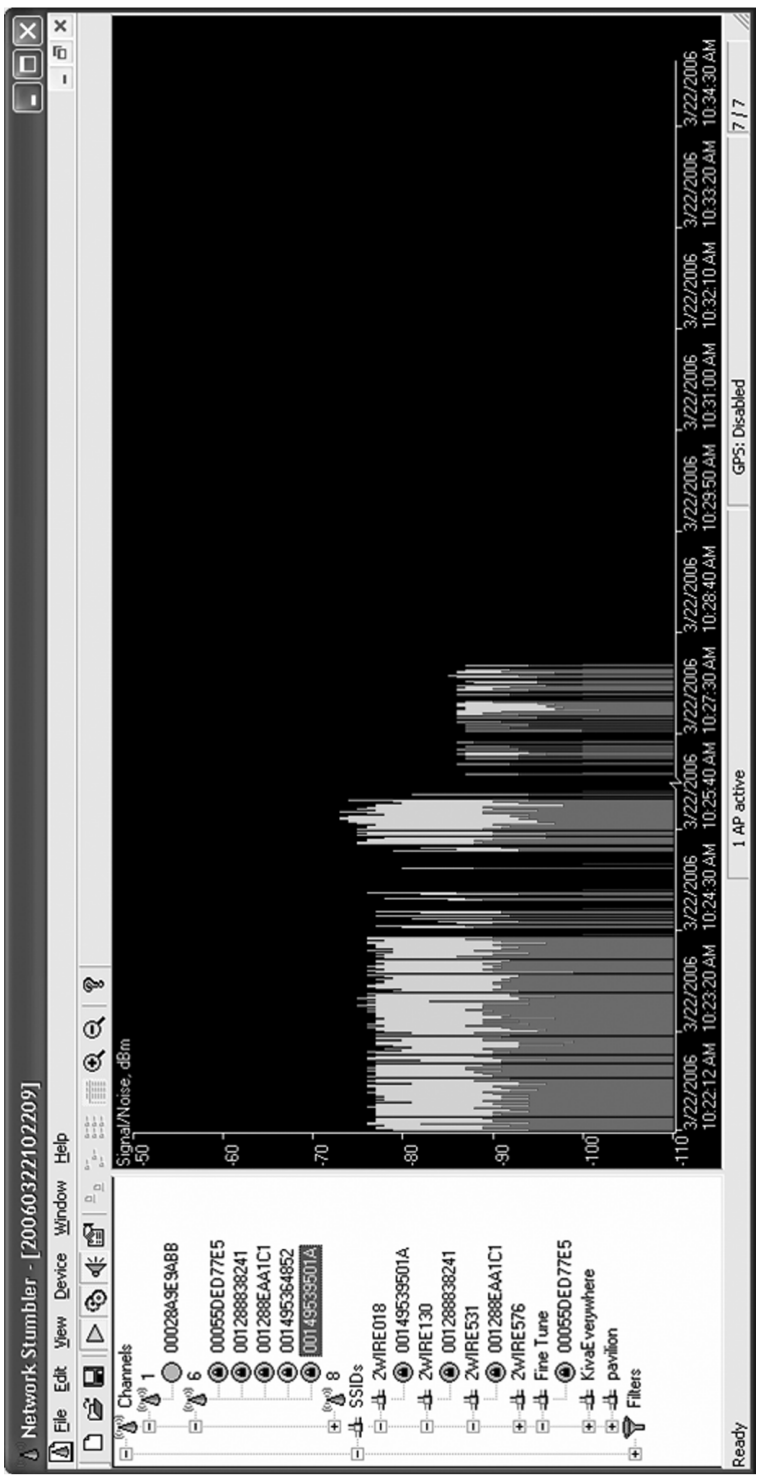

FIGURE 7.10 Changes in signal strength.