based data traffic (if enabled) occurs alternately between contention periods.

An access point polls stations according to a polling list, then switches to a contention period when stations use DCF. This process enables support for both synchronous and asynchronous modes of operation. No known wireless NICs or access points on the market today, however, implement PCF.

An issue with PCF is that not many vendors support it in their products. As a result, it's generally not an option available to users. Future products, however, might support PCF in order to offer quality of service (QoS) mechanisms.

The following sections summarize primary 802.11 MAC functions.

Scanning

The 802.11 standard defines both passive and active scanning, whereby a radio NIC searches for access points. Passive scanning is mandatory where each NIC scans individual channels to find the best access-point signal. Periodically, access points broadcast a beacon, and the radio NIC receives these beacons while scanning and takes note of the corresponding signal strengths. The beacons contain information about the access point, including SSID and supported data rates. The radio NIC can use this information along with the signal strength to compare access points and decide on which one to use.

Optional active scanning is similar, except the radio NIC initiates the process by broadcasting a probe frame, and all access points within range respond with a probe response. Active scanning enables a radio NIC to receive immediate response from access points, without waiting for a beacon transmission. The issue, however, is that active scanning imposes additional overhead on the network because of the transmission of probe and corresponding response frames.

Exhibit 1010
Page 165 of 292

Stations set to ad hoc mode form are what the 802.11 standard refers to as an independent basic service set. In this mode, one of the stations always sends a beacon, which alerts new stations of the network presence. The responsibility of transmitting the beacon is based upon each station waiting for a beacon interval to expire and an additional random time. A station transmits a beacon if after the beacon interval and random time the station does not receive a beacon from another station. This distributes the responsibility for sending beacons among all stations.

Authentication

Authentication is the process of proving identity, and the 802.11 standard specifies two forms: open system authentication and shared key authentication. Open system authentication is mandatory, and it's a two-step process. A radio NIC initiates the process by sending an authentication request frame to the access point. The access point replies with an authentication response frame containing approval or disapproval of authentication indicated in the status code field in the frame body.

Shared key authentication is an optional four-step process that bases authentication on whether the authenticating device has the correct WEP key. The radio NIC starts by sending an authentication request frame to the access point. The access point then places challenge text into the frame body of a response frame and sends it to the radio NIC. The radio NIC uses its WEP key to encrypt the challenge text and then sends it back to the access point in another authentication frame. The access point decrypts the challenge text and compares it to the initial text. If the text is equivalent, the access point assumes that the radio NIC has the correct key. The access point finishes the sequence by sending an authentication frame to the radio NIC with the approval or disapproval. Many hackers know how to break through shared key authentication, however, so it's not a good idea to depend on it for a high level of security.

Exhibit 1010
Page 166 of 292

Association

When authenticated, the radio NIC must associate with the access point before sending data frames. **Association** is necessary to synchronize the radio NIC and access point with important information, such as supported data rates. The radio NIC initiates the association by sending an association request frame containing elements such as SSID and supported data rates. The access point responds by sending an association response frame containing an association ID along with other information regarding the access point. Once the radio NIC and access point complete the association process, they can send data frames to each other.

WEP

With the optional WEP enabled, the wireless NIC encrypts the body (not the header) of each frame before transmission using a common key; and the receiving station decrypts the frame upon receipt using the common key. The 802.11 standard does not specify a key distribution method, which makes 802.11 wireless LANs vulnerable to eavesdroppers. The 802.11i version of the standard, however, is improving 802.11 security by incorporating 802.1x and stronger encryption into the standard.

RTS/CTS

The optional request-to-send and clear-to-send (RTS/CTS) function allows the access point to control use of the medium for stations activating RTS/CTS. With most radio NICs, users can set a maximum frame-length threshold for when the radio NIC activates RTS/CTS. For example, a frame length of 1,000 bytes triggers RTS/CTS for all frames larger than 1,000 bytes. The use of RTS/CTS alleviates hidden node problems (where two or more radio NICs can't hear each other and they are associated with the same access point).

Exhibit 1010
Page 167 of 292

If the radio NIC activates RTS/CTS, it first sends an RTS frame to an access point before sending a data frame. The access point then responds with a CTS frame, indicating that the radio NIC can send the data frame. With the CTS frame, the access point provides a value in the duration field of the frame header that holds off other stations from transmitting until after the radio NIC initiating the RTS can send its data frame. This avoids collisions between hidden nodes. The RTS/CTS handshake continues for each frame, as long as the frame size exceeds the threshold set in the corresponding radio NIC.

Power Save Mode

The optional power save mode that a user can turn on enables the radio NIC to conserve battery power when there is no need to send data. With power save mode on, the radio NIC indicates its desire to enter a sleep state to the access point through a status bit located in the header of each frame. The access point takes note of each radio NIC wanting to enter power save mode and buffers packets corresponding to the sleeping station.

In order to still receive data frames, the sleeping NIC must wake up periodically (at the right time) to receive regular beacon transmissions coming from the access point. These beacons identify whether sleeping stations have frames buffered at the access point and are waiting for delivery to their respective destinations. The radio NICs having awaiting frames will request them from the access point. After receiving the frames, the radio NIC can go back to sleep.

Fragmentation

The optional fragmentation function enables an 802.11 station to divide data packets into smaller frames. This is to avoid needing to retransmit large frames in the presence of RF interference. The bit errors resulting from RF interference are likely to affect a single frame, and it requires less overhead to retransmit a smaller frame rather than a larger one. As with RTS/CTS, users can set a maximum frame-

length threshold for when the radio NIC activates fragmentation. If the frame size is larger than the threshold, the radio NIC breaks the packet into multiple frames, with each frame no larger than the threshold value.

802.11 Physical Layers

Several 802.11 Physical layers satisfy a variety of application requirements. The following sections provide a summary of each of the 802.11 Physical layers.

Initial 802.11

The initial 802.11 standard ratified in 1997 includes frequency hopping spread spectrum (FHSS) and direct sequence spread spectrum (DSSS) physical layers operating in the 2.4-GHz band with data rates of up to 2 Mbps. FHSS transmits a wideband signal that spans the entire 2.4-GHz band. It's possible to tune FHSS access points to as many as 15 different hopping patterns that don't interfere with each other, which enables up to 15 FHSS access points to effectively operate in the same area.

Because the current version of 802.11 FHSS has only a maximum data rate of 2 Mbps, not many companies sell FHSS solutions for an indoor wireless LAN. Much faster 802.11a, 802.11b, and 802.11g wireless LANs are now available. Also, FHSS doesn't interoperate with any of the other 802.11 physical layers. FHSS does, however, provide a very solution for outdoor, point-to-multipoint systems. This is because FHSS is more resilient to the RF interference that might be present in outdoor environments.

802.11 DSSS also operates only up to 2 Mbps, but it interoperates with the newer 802.11b physical layer. As a result, a user having an 802.11 DSSS radio NIC in her laptop can interface with an 802.11b access point. This situation is unlikely, however, because 802.11 DSSS radio NICs are not sold anymore.

Exhibit 1010
Page 169 of 292

802.11a

Toward the end of 1999, the IEEE released 802.11a, which defines operation in the 5-GHz band using Orthogonal Frequency Division Multiplexing (OFDM) with data rates up to 54 Mbps. Products, however, didn't become available until 2000, primarily because of the difficulties in developing circuitry in the 5-GHz band.

802.11a operates up to 54 Mbps in the 5-GHz band using OFDM with a range up to 100 feet depending on the actual data rate. 802.11a access points and radio NICs just became available in late 2001; therefore, the installed base of 802.11a wireless LANs today is relatively small as compared to 802.11b. Because of this, carefully consider interoperability issues that might result if you choose to deploy 802.11a networks.

A strong advantage of 802.11a is that it offers the highest capacity because of 12 separate, non-overlapping channels. This makes it a good choice for supporting a high concentration of users and higher-performance applications, such as video streaming. In addition to outperforming 802.11b systems, 802.11a has greater capacity than 802.11g.

Another advantage of 802.11a is that the 5-GHz band is uncrowded, which enables users to achieve higher levels of performance. Most interfering devices, such as microwave ovens and cordless phones, operate in the 2.4-GHz band. With less potential for RF interference, the deployment of a wireless LAN is less risky.

A potential issue of 802.11a is limited range, mainly because of operation in the higher frequency 5-GHz band. At 54 Mbps, you'll have a range less than 100 feet in most facilities. This requires a greater number of access points to fully cover a facility, as compared to an 802.11b system.

If you compare the operation of 802.11b and 802.11a, however, an 802.11a user has a higher data rate at the same range as an 802.11b user until the 802.11a user loses connectivity. The 802.11b user, how-

ever, can continue operating at lower data rates—1 or 2 Mbps, for example—at longer ranges than 802.11a.

A definite problem is that 802.11a and 802.11b/g are not compatible. For example, a user equipped with an 802.11b radio card can't associate with an 802.11a access point. The opposite scenario also applies. Vendors are solving this problem by introducing multimode radio cards that implement both 802.11a and 802.11b.

An 802.11a modulator converts the binary signal into an analog waveform through the use of different modulation types, depending on which data rate is chosen. With 6-Mbps operation, for example, the PMD uses binary phase shift keying (BPSK), which shifts the phase of the transmit center frequency to represent different data bit patterns. The higher data rates, such as 54 Mbps, employ quadrature amplitude modulation (QAM) to represent data bits by varying the transmit center frequency with different amplitude levels in addition to phase shifts.

802.11b

Along with 802.11a, IEEE ratified 802.11b, which is a higher-rate extension to the initial direct sequence standard in the 2.4-GHz band—with data rates up to 11 Mbps. 802.11b access points and radio NICs have been available since 1999; therefore, most wireless LANs installed today are 802.11b compliant.

A significant advantage of 802.11b is its relatively long-range properties. With 802.11b, you can achieve a range of 300 feet in most indoor facilities. The superior range allows the deployment of wireless LANs with fewer access points to cover a facility as compared to 802.11a.

A disadvantage of 802.11b is that you're limited to three non-overlapping channels in the 2.4 GHz band. The 802.11 standard specifies 14 channels (only channels 1 through 11 are available in the U.S.) for configuring access points; but each channel occupies roughly one third of

Exhibit 1010
Page 171 of 292

the overall 2.4-GHz band while transmitting a signal. Most companies utilize only channels 1, 6, and 11 to ensure access points don't interfere with each other. This limits overall capacity of 802.11b, which makes it most suitable for supporting medium performance applications, such as e-mail and web surfing.

Another disadvantage of 802.11b is the potential for RF interference from other radio devices. For example, a 2.4-GHz cordless phone severely interferes with an 802.11b wireless LAN, which significantly reduces the performance for users. Microwave ovens and other devices operating in the 2.4-GHz band can also cause interference.

802.11b uses DSSS to disperse the data frame signal over a 22-MHz portion of the 2.4-GHz frequency band. This results in greater immunity to RF interference as compared to narrowband signaling, which is why the FCC deems the operation of spread spectrum systems license free.

The 802.11b modulator converts the spread binary signal into an analog waveform through the use of different modulation types, depending on which data rate is chosen. For example with 1-Mbps operation, the PMD uses differential binary phase shift keying (DBPSK). This isn't really as complex as it sounds. The modulator merely shifts the phase of the center transmit frequency to distinguish a binary 1 from a binary 0 within the data stream.

For 2-Mbps transmission, the PMD uses differential quadrature phase shift keying (DQPSK), which is similar to DBPSK except four possible phase shifts that represents every two data bits. This is a clever process that enables the data stream to be sent at 2 Mbps while using the same amount of bandwidth as the one sent at 1 Mbps. The modulator uses similar methods for the higher 5.5-Mbps and 11-Mbps data rates.

802.11g

IEEE ratified the 802.11g standard in 2003, which is compatible with 802.11b and increases performance up to 54 Mbps in the 2.4-GHz band using OFDM.

A strong advantage of 802.11g is that it's backward compatible with 802.11b. Companies with existing 802.11b networks can generally upgrade their access points to become 802.11g compliant through simple firmware upgrades. This provides an effective migration path for wireless LANs. An issue, however, is that the presence of 802.11b client devices within an 802.11g environment requires protection mechanisms that limit the performance of the overall wireless LAN. The problem is that 802.11b devices can't understand when 802.11g devices are transmitting because of a difference in modulation types. As a result, both types of devices must announce their impending use of the medium using a commonly understood modulation type.

The disadvantages of 802.11b, such as potential for RF interference and limit of three non-overlapping channels, still apply to 802.11g because of operation in the 2.4-GHz band. As a result, 802.11g networks have capacity constraints as compared to 802.11a.

2.4 GHz or 5 GHz?

When deploying a wireless LAN, companies must make a decision on whether to use NICs and access points designed to operate in the 2.4-GHz or 5-GHz band —or both. Not too long ago, the choice of frequency band was easy, when only 2.4-GHz (802.11b) products were available. Now, 802.11b and 802.11g products are both available that operate in the 2.4-GHz band, while 802.11a uses the 5-GHz band. This can cause confusion when designing a wireless LAN, so take a look at what is necessary to consider when making this critical decision.

Exhibit 1010
Page 173 of 292

When assessing the pros and cons of 2.4-GHz and 5-GHz systems, be sure to first define requirements. This provides a solid basis for defining all design elements. Without firm requirements, you'll be making the choice on flimsy ground.

The following are requirements for consideration when deciding between 2.4-GHz or 5-GHz solutions:

- Geographical Location— Consider the geographical location of where the wireless LAN will operate. A 2.4-GHz wireless LAN has regulatory acceptance throughout most of the world; however, the use of 5 GHz for wireless LANs is somewhat limited. For example, the U.S. allows operation of 5-GHz wireless LANs, but other countries do not. Your location might require you to use the 2.4-GHz band regardless of other requirements.
- Performance— The 5-GHz bands have much greater spectrum available. Each of the 12 non-overlapping channels in this band has 20 MHz of bandwidth. This means significantly better performance as compared to the 2.4-GHz band. The entire 2.4-GHz band is 80 MHz wide, which allows only three non-overlapping channels. If high performance is an important requirement, lean toward the 5-GHz band.
- Facility Size— As frequency increases, range generally decreases. As a result, 5-GHz systems generally have less range than ones operating in the 2.4-GHz band. The selection of a 5-GHz wireless LAN could require a greater number of access points, which can result in higher costs. As a result, you might benefit by deploying 2.4-GHz systems in larger facilities unless high performance is critical. Keep in mind, however, that 5-GHz systems might have equal or even better range in some situations.

Exhibit 1010
Page 174 of 292

- RF Interference— 2.4GHz wireless LANs can experience interference from cordless phones, microwaves, and other wireless LANs. The interfering signals degrade the performance of an 802.11b wireless LAN by periodically blocking users and access points from accessing the shared air medium. If it's not possible to reduce potential interference to an acceptable level, consider deploying a 5-GHz system, which is relatively free from interfering sources. There are some 5-GHz phones now on the market, but it's much more possible to avoid this interference because of the many non-overlapping channels that 802.11a offers.
- Interoperability— 2.4-GHz and 5-GHz systems are not directly compatible, and few users and access points operate in the 5-GHz band. Consequently, it might be best to deploy a 2.4-GHz solution if you have little control over the NICs that users have in their PDAs and laptops. This applies mostly to universities and public wireless LAN hotspots. Your application might require you to implement 2.4 GHz to support the more common 802.11b-equipped users.

  Vendors, however, offer dual-band radio NICs and access points, which reduces interoperability problems. Someone equipped with a dual-band radio NIC can associate with either a 2.4-GHz (802.11b/g) or 5-GHz (802.11a) access point. As a greater number of users begin equipping their devices with the dual-band radio NICs, the interoperability issue will diminish.
- Security— Security of the wireless LAN is of great concern to most companies. By minimizing the propagation of radio waves outside the controlled area of a facility, a wireless network is more secure because of the reduction of the potential for eavesdropping and denial of service (DoS) attacks. As a result, 5-GHz systems can pro-

Exhibit 1010
Page 175 of 292