Creating and Administering Wireless Networks

802.11® Wireless Networks

The Definitive Guide

O'REILLY®

Matthew S. Gast

INTEL-1031, Part 1
10,291,449

802.11® Wireless Networks

The Definitive Guide

802.11® Wireless Networks

The Definitive Guide

Matthew S. Gast

O'REILLY®

Beijing • Cambridge • Farnham • Köln • Paris • Sebastopol • Taipei • Tokyo

802.11® Wireless Networks: The Definitive Guide
by Matthew S. Gast

Copyright © 2002 O'Reilly Media, Inc. All rights reserved.
Printed in the United States of America.

Published by O'Reilly Media, Inc., 1005 Gravenstein Highway North,
Sebastopol, CA 95472.

O'Reilly Media, Inc. books may be purchased for educational, business, or sales promotional use. Online editions are also available for most titles (*safari.oreilly.com*). For more information contact our corporate/institutional sales department: (800) 998-9938 or *corporate@oreilly.com*.

Editor: Mike Loukides

Production Editor: Matt Hutchinson

Cover Designer: Ellie Volckhausen

Printing History:

April 2002: First Edition.

Nutshell Handbook, the Nutshell Handbook logo, and the O'Reilly logo are registered trademarks of O'Reilly Media, Inc. *802.11® Wireless Networks: The Definitive Guide*, the image of a horseshoe bat, and related trade dress are trademarks of O'Reilly Media, Inc. Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and O'Reilly Media, Inc. was aware of a trademark claim, the designations have been printed in caps or initial caps.

802.11® and all 802.11-based trademarks and logos are trademarks or registered trademarks of IEEE, Inc. in the United States and other countries. O'Reilly Media, Inc. is independent of IEEE.

While every precaution has been taken in the preparation of this book, the author and publisher assume no responsibility for errors or omissions, or for damages resulting from the use of the information contained herein.

This book uses RepKover™, a durable and flexible lay-flat binding.

ISBN: 0-596-00183-5
[M]

[11/04]

CHAPTER 1

Introduction to Wireless Networks

Over the past five years, the world has become increasingly mobile. As a result, traditional ways of networking the world have proven inadequate to meet the challenges posed by our new collective lifestyle. If users must be connected to a network by physical cables, their movement is dramatically reduced. Wireless connectivity, however, poses no such restriction and allows a great deal more free movement on the part of the network user. As a result, wireless technologies are encroaching on the traditional realm of "fixed" or "wired" networks. This change is obvious to anybody who drives on a regular basis. One of the "life and death" challenges to those of us who drive on a regular basis is the daily gauntlet of erratically driven cars containing mobile phone users in the driver's seat.

We are on the cusp of an equally profound change in computer networking. Wireless telephony has been successful because it enables people to connect with each other regardless of location. New technologies targeted at computer networks promise to do the same for Internet connectivity. The most successful wireless networking technology this far has been 802.11.

Why Wireless?

To dive into a specific technology at this point is getting a bit ahead of the story, though. Wireless networks share several important advantages, no matter how the protocols are designed, or even what type of data they carry.

The most obvious advantage of wireless networking is *mobility*. Wireless network users can connect to existing networks and are then allowed to roam freely. A mobile telephone user can drive miles in the course of a single conversation because the phone connects the user through cell towers. Initially, mobile telephony was expensive. Costs restricted its use to highly mobile professionals such as sales managers and important executive decision makers who might need to be reached at a moment's notice regardless of their location. Mobile telephony has proven to be a

useful service, however, and now it is relatively common in the United States and extremely common among Europeans.*

Likewise, wireless data networks free software developers from the tethers of an Ethernet cable at a desk. Developers can work in the library, in a conference room, in the parking lot, or even in the coffee house across the street. As long as the wireless users remain within the range of the base station, they can take advantage of the network. Commonly available equipment can easily cover a corporate campus; with some work, more exotic equipment, and favorable terrain, you can extend the range of an 802.11 network up to a few miles.

Wireless networks typically have a great deal of *flexibility*, which can translate into rapid deployment. Wireless networks use a number of base stations to connect users to an existing network. The infrastructure side of a wireless network, however, is qualitatively the same whether you are connecting one user or a million users. To offer service in a given area, you need base stations and antennas in place. Once that infrastructure is built, however, adding a user to a wireless network is mostly a matter of authorization. With the infrastructure built, it must be configured to recognize and offer services to the new users, but authorization does not require more infrastructure. Adding a user to a wireless network is a matter of configuring the infrastructure, but it does not involve running cables, punching down terminals, and patching in a new jack.†

Flexibility is an important attribute for service providers. One of the markets that many 802.11 equipment vendors have been chasing is the so-called "hot spot" connectivity market. Airports and train stations are likely to have itinerant business travelers interested in network access during connection delays. Coffeehouses and other public gathering spots are social venues in which network access is desirable. Many cafes already offer Internet access; offering Internet access over a wireless network is a natural extension of the existing Internet connectivity. While it is possible to serve a fluid group of users with Ethernet jacks, supplying access over a wired network is problematic for several reasons. Running cables is time-consuming and expensive and may also require construction. Properly guessing the correct number of cable drops is more an art than a science. With a wireless network, though, there is no need to suffer through construction or make educated (or wild) guesses about demand. A simple wired infrastructure connects to the Internet, and then the wireless network can

* While most of my colleagues, acquaintances, and family in the U.S. have mobile telephones, it is still possible to be a holdout. In Europe, it seems as if everybody has a mobile phone—one cab driver in Finland I spoke with while writing this book took great pride in the fact that his family of four had six mobile telephones!

† This simple example ignores the challenges of scale. Naturally, if the new users will overload the existing infrastructure, the infrastructure itself will need to be beefed up. Infrastructure expansion can be expensive and time-consuming, especially if it involves legal and regulatory approval. However, my basic point holds: adding a user to a wireless network can often be reduced to a matter of configuration (moving or changing bits) while adding a user to a fixed network requires making physical connections (moving atoms), and moving bits is easier than moving atoms.

accommodate as many users as needed. Although wireless LANs have somewhat limited bandwidth, the limiting factor in networking a small hot spot is likely to be the cost of WAN bandwidth to the supporting infrastructure.

Flexibility may be particularly important in older buildings because it reduces the need for constructions. Once a building is declared historical, remodeling can be particularly difficult. In addition to meeting owner requirements, historical preservation agencies must be satisfied that new construction is not desecrating the past. Wireless networks can be deployed extremely rapidly in such environments because there is only a small wired network to install.

Flexibility has also led to the development of grassroots community networks. With the rapid price erosion of 802.11 equipment, bands of volunteers are setting up shared wireless networks open to visitors. Community networks are also extending the range of Internet access past the limitations for DSL into communities where high-speed Internet access has been only a dream. Community networks have been particularly successful in out-of-the way places that are too rugged for traditional wireline approaches.

Like all networks, wireless networks transmit data over a network medium. The medium is a form of electromagnetic radiation.* To be well-suited for use on mobile networks, the medium must be able to cover a wide area so clients can move throughout a coverage area. The two media that have seen the widest use in local-area applications are infrared light and radio waves. Most portable PCs sold now have infrared ports that can make quick connections to printers and other peripherals. However, infrared light has limitations; it is easily blocked by walls, partitions, and other office construction. Radio waves can penetrate most office obstructions and offer a wider coverage range. It is no surprise that most, if not all, 802.11 products on the market use the radio wave physical layer.

Radio Spectrum: The Key Resource

Wireless devices are constrained to operate in a certain frequency band. Each band has an associated *bandwidth*, which is simply the amount of frequency space in the band. Bandwidth has acquired a connotation of being a measure of the data capacity of a link. A great deal of mathematics, information theory, and signal processing can be used to show that higher-bandwidth slices can be used to transmit more information. As an example, an analog mobile telephony channel requires a 20-kHz bandwidth. TV signals are vastly more complex and have a correspondingly larger bandwidth of 6 MHz.

* Laser light is also used by some wireless networking applications, but the extreme focus of a laser beam makes it suited only for applications in which the ends are stationary. "Fixed wireless" applications, in which lasers replace other access technology such as leased telephone circuits, are a common application.

The use of a radio spectrum is rigorously controlled by regulatory authorities through *licensing* processes. In the U.S., regulation is done by the Federal Communications Commission (FCC). Many FCC rules are adopted by other countries throughout the Americas. European allocation is performed by CEPT's European Radiocommunications Office (ERO). Other allocation work is done by the International Telecommunications Union (ITU). To prevent overlapping uses of the radio waves, frequency is allocated in bands, which are simply ranges of frequencies available to specified applications. Table 1-1 lists some common frequency bands used in the U.S.

Table 1-1. Common U.S. frequency bands

Band	Frequency range
UHF ISM	902–928 MHz
S-Band	2–4 GHz
S-Band ISM	2.4–2.5 GHz
C-Band	4–8 GHz
C-Band satellite downlink	3.7–4.2 GHz
C-Band Radar (weather)	5.25–5.925 GHz
C-Band ISM	5.725–5.875 GHz
C-Band satellite uplink	5.925–6.425 GHz
X-Band	8–12 GHz
X-Band Radar (police/weather)	8.5–10.55 GHz
Ku-Band	12–18 GHz
Ku-Band Radar (police)	13.4–14 GHz
	15.7–17.7 GHz

The ISM bands

In Table 1-1, there are three bands labeled ISM, which is an abbreviation for industrial, scientific, and medical. ISM bands are set aside for equipment that, broadly speaking, is related to industrial or scientific processes or is used by medical equipment. Perhaps the most familiar ISM-band device is the microwave oven, which operates in the 2.4-GHz ISM band because electromagnetic radiation at that frequency is particularly effective for heating water.

I pay special attention to the ISM bands because that's where 802.11 devices operate. The more common 802.11b devices operate in S-band ISM. The ISM bands are generally license-free, provided that devices are low-power. How much sense does it make to require a license for microwave ovens, after all? Likewise, you don't need a license to set up and operate a wireless network.

The Limits of Wireless Networking

Wireless networks do not replace fixed networks. The main advantage of mobility is that the network user is moving. Servers and other data center equipment must access data, but the physical location of the server is irrelevant. As long as the servers do not move, they may as well be connected to wires that do not move.

The speed of wireless networks is constrained by the available bandwidth. Information theory can be used to deduce the upper limit on the speed of a network. Unless the regulatory authorities are willing to make the unlicensed spectrum bands bigger, there is an upper limit on the speed of wireless networks. Wireless-network hardware tends to be slower than wired hardware. Unlike the 10-GB Ethernet standard, wireless-network standards must carefully validate received frames to guard against loss due to the unreliability of the wireless medium.

Using radio waves as the network medium poses several challenges. Specifications for wired networks are designed so that a network will work as long as it respects the specifications. Radio waves can suffer from a number of propagation problems that may interrupt the radio link, such as multipath interference and shadows.

Security on any network is a prime concern. On wireless networks, it is often a critical concern because the network transmissions are available to anyone within range of the transmitter with the appropriate antenna. On a wired network, the signals stay in the wires and can be protected by strong physical-access control (locks on the doors of wiring closets, and so on). On a wireless network, sniffing is much easier because the radio transmissions are designed to be processed by any receiver within range. Furthermore, wireless networks tend to have fuzzy boundaries. A corporate wireless network may extend outside the building. It is quite possible that a parked car across the street could be receiving the signals from your network. As an experiment on one of my trips to San Francisco, I turned on my laptop to count the number of wireless networks near a major highway outside the city. I found eight without expending any significant effort. A significantly more motivated investigator would undoubtedly have discovered many more networks by using a much more sensitive antenna mounted outside the steel shell of the car.

A Network by Any Other Name...

Wireless networking is a hot industry segment. Several wireless technologies have been targeted primarily for data transmission. Bluetooth is a standard used to build small networks between peripherals: a form of "wireless wires," if you will. Most people in the industry are familiar with the hype surrounding Bluetooth. I haven't met many people who have used devices based on the Bluetooth specification.

Third-generation (3G) mobile telephony networks are also a familiar source of hype. They promise data rates of megabits per cell, as well as the "always on" connections

that have proven to be quite valuable to DSL and cable modem customers. In spite of the hype and press from 3G equipment vendors, the rollout of commercial 3G services has been continually pushed back.

In contrast to Bluetooth and 3G, equipment based on the IEEE 802.11 standard has been an astounding success. While Bluetooth and 3G may be successful in the future, 802.11 is a success *now*. Apple initiated the pricing moves that caused the market for 802.11 equipment to explode in 1999. Price erosion made the equipment affordable and started the growth that continues today.

This is a book about 802.11 networks. 802.11 goes by a variety of names, depending on who is talking about it. Some people call 802.11 *wireless Ethernet*, to emphasize its shared lineage with the traditional wired Ethernet (802.3). More recently, the Wireless Ethernet Compatibility Alliance (WECA) has been pushing its *Wi-Fi* ("wireless fidelity") certification program.* Any 802.11 vendor can have its products tested for interoperability. Equipment that passes the test suite can use the Wi-Fi mark. For newer products based on the 802.11a standard, WECA will allow use of the *Wi-Fi5* mark. The "5" reflects the fact that 802.11a products use a different frequency band of around 5 GHz.

Table 1-2 is a basic comparison of the different 802.11 standards. Products based on 802.11 were initially released in 1997. 802.11 included an infrared (IR) layer that was never widely deployed, as well as two spread-spectrum radio layers: frequency hopping (FH) and direct sequence (DS). (The differences between these two radio layers is described in Chapter 10.) Initial 802.11 products were limited to 2 Mbps, which is quite slow by modern network standards. The IEEE 802.11 working group quickly began working on faster radio layers and standardized both 802.11a and 802.11b in 1999. Products based on 802.11b were released in 1999 and can operate at speeds of up to 11 Mbps. 802.11a uses a third radio technique called orthogonal frequency division multiplexing (OFDM). 802.11a operates in a different frequency band entirely and currently has regulatory approval only in the United States. As you can see from the table, 802.11 already provides speeds faster than 10BASE-T Ethernet and is reasonably competitive with Fast Ethernet.

Table 1-2. Comparison of 802.11 standards

IEEE standard	Speed	Frequency band	Notes
802.11	1 Mbps 2 Mbps	2.4 GHz	First standard (1997). Featured both frequency-hopping and direct-sequence modulation techniques.
802.11a	up to 54 Mbps	5 GHz	Second standard (1999), but products not released until late 2000.
802.11b	5.5 Mbps 11 Mbps	2.4 GHz	Third standard, but second wave of products. The most common 802.11 equipment as this book was written.
802.11g	up to 54 Mbps	2.4 GHz	Not yet standardized.

* More details on WECA and the Wi-Fi certification can be found at *http://www.wi-fi.org/*.

CHAPTER 2

Overview of 802.11 Networks

Before studying the details of anything, it often helps to get a general "lay of the land." A basic introduction is often necessary when studying networking topics because the number of acronyms can be overwhelming. Unfortunately, 802.11 takes acronyms to new heights, which makes the introduction that much more important. To understand 802.11 on anything more than a superficial basis, you must get comfortable with some esoteric terminology and a herd of three-letter acronyms. This chapter is the glue that binds the entire book together. Read it for a basic understanding of 802.11, the concepts that will likely be important to users, and how the protocol is designed to provide an experience as much like Ethernet as possible. After that, move on to the low-level protocol details or deployment, depending on your interests and needs.

Part of the reason why this introduction is important is because it introduces the acronyms used throughout the book. With 802.11, the introduction serves another important purpose. 802.11 is superficially similar to Ethernet. Understanding the background of Ethernet helps slightly with 802.11, but there is a host of additional background needed to appreciate how 802.11 adapts traditional Ethernet technology to a wireless world. To account for the differences between wired networks and the wireless media used by 802.11, a number of additional management features were added. At the heart of 802.11 is a white lie about the meaning of media access control (MAC). Wireless network interface cards are assigned 48-bit MAC addresses, and, for all practical purposes, they look like Ethernet network interface cards. In fact, the MAC address assignment is done from the same address pool so that 802.11 cards have unique addresses even when deployed into a network with wired Ethernet stations.

To outside network devices, these MAC addresses appear to be fixed, just as in other IEEE 802 networks; 802.11 MAC addresses go into ARP tables alongside Ethernet addresses, use the same set of vendor prefixes, and are otherwise indistinguishable from Ethernet addresses. The devices that comprise an 802.11 network (access points and other 802.11 devices) know better. There are many differences between an 802.11 device and an Ethernet device, but the most obvious is that 802.11 devices

are mobile; they can easily move from one part of the network to another. The 802.11 devices on your network understand this and deliver frames to the current location of the mobile station.

IEEE 802 Network Technology Family Tree

802.11 is a member of the IEEE 802 family, which is a series of specifications for local area network (LAN) technologies. Figure 2-1 shows the relationship between the various components of the 802 family and their place in the OSI model.

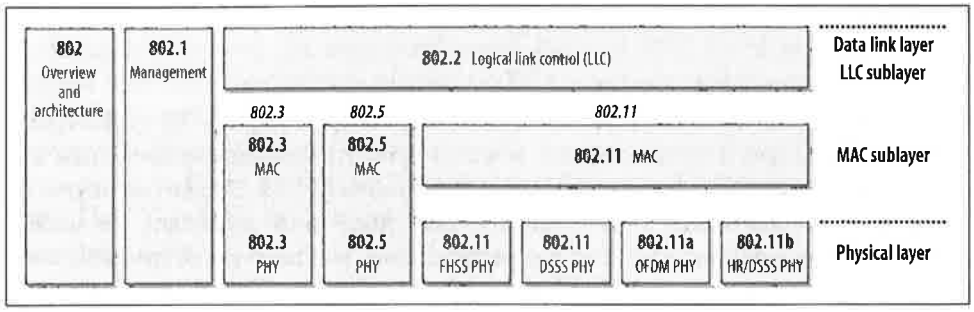

Figure 2-1. The IEEE 802 family and its relation to the OSI model

IEEE 802 specifications are focused on the two lowest layers of the OSI model because they incorporate both physical and data link components. All 802 networks have both a MAC and a Physical (PHY) component. The MAC is a set of rules to determine how to access the medium and send data, but the details of transmission and reception are left to the PHY.

Individual specifications in the 802 series are identified by a second number. For example, 802.3 is the specification for a Carrier Sense Multiple Access network with Collision Detection (CSMA/CD), which is related to (and often mistakenly called) Ethernet, and 802.5 is the Token Ring specification. Other specifications describe other parts of the 802 protocol stack. 802.2 specifies a common link layer, the Logical Link Control (LLC), which can be used by any lower-layer LAN technology. Management features for 802 networks are specified in 802.1. Among 802.1's many provisions are bridging (802.1d) and virtual LANs, or VLANs (802.1q).

802.11 is just another link layer that can use the 802.2/LLC encapsulation. The base 802.11 specification includes the 802.11 MAC and two physical layers: a frequency-hopping spread-spectrum (FHSS) physical layer and a direct-sequence spread-spectrum (DSSS) physical layer. Later revisions to 802.11 added additional physical layers. 802.11b specifies a high-rate direct-sequence layer (HR/DSSS); products based on 802.11b hit the marketplace in 1999 and make up the bulk of the installed base. 802.11a describes a physical layer based on orthogonal frequency division multiplexing (OFDM); products based on 802.11a were released as this book was completed.

To say that 802.11 is "just another link layer for 802.2" is to omit the details in the rest of this book, but 802.11 is exciting precisely because of these details. 802.11 allows for mobile network access; in accomplishing this goal, a number of additional features were incorporated into the MAC. As a result, the 802.11 MAC may seem baroquely complex compared to other IEEE 802 MAC specifications.

The use of radio waves as a physical layer requires a relatively complex PHY, as well. 802.11 splits the PHY into two generic components: the Physical Layer Convergence Procedure (PLCP), to map the MAC frames onto the medium, and a Physical Medium Dependent (PMD) system to transmit those frames. The PLCP straddles the boundary of the MAC and physical layers, as shown in Figure 2-2. In 802.11, the PLCP adds a number of fields to the frame as it is transmitted "in the air."

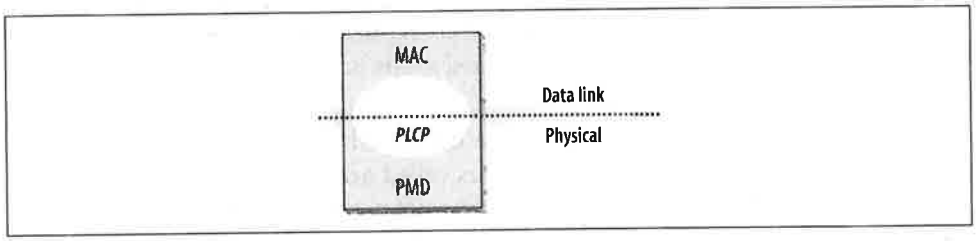

Figure 2-2. PHY components

All this complexity begs the question of how much you actually need to know. As with any technology, the more you know, the better off you will be. The 802.11 protocols have many knobs and dials that you can tweak, but most 802.11 implementations hide this complexity. Many of the features of the standard come into their own only when the network is congested, either with a lot of traffic or with a large number of wireless stations. Today's networks tend not to push the limits in either respect. At any rate, I can't blame you for wanting to skip the chapters about the protocols and jump ahead to the chapters about planning and installing an 802.11 network. After you've read this chapter, you can skip ahead to Chapters 12–17 and return to the chapters on the protocol's inner workings when you need (or want) to know more.

802.11 Nomenclature and Design

802.11 networks consist of four major physical components, which are summarized in Figure 2-3. The components are:

Distribution system
: When several access points are connected to form a large coverage area, they must communicate with each other to track the movements of mobile stations. The distribution system is the logical component of 802.11 used to forward frames to their destination. 802.11 does not specify any particular technology for

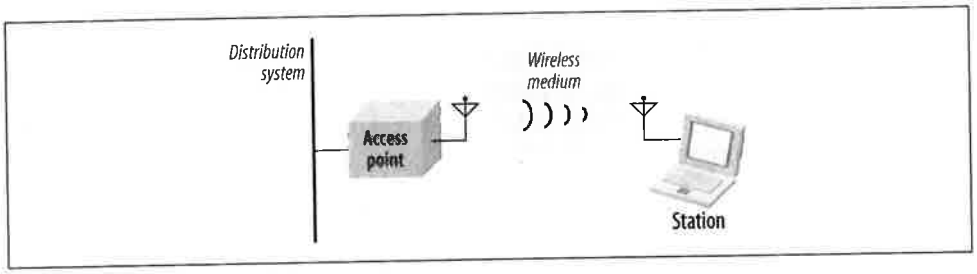

Figure 2-3. Components of 802.11 LANs

the distribution system. In most commercial products, the distribution system is implemented as a combination of a bridging engine and a distribution system medium, which is the backbone network used to relay frames between access points; it is often called simply the backbone network. In nearly all commercially successful products, Ethernet is used as the backbone network technology.

Access points
: Frames on an 802.11 network must be converted to another type of frame for delivery to the rest of the world. Devices called access points perform the wireless-to-wired bridging function. (Access points perform a number of other functions, but bridging is by far the most important.)

Wireless medium
: To move frames from station to station, the standard uses a wireless medium. Several different physical layers are defined; the architecture allows multiple physical layers to be developed to support the 802.11 MAC. Initially, two radio frequency (RF) physical layers and one infrared physical layer were standardized, though the RF layers have proven far more popular.

Stations
: Networks are built to transfer data between stations. Stations are computing devices with wireless network interfaces. Typically, stations are battery-operated laptop or handheld computers. There is no reason why stations must be portable computing devices, though. In some environments, wireless networking is used to avoid pulling new cable, and desktops are connected by wireless LANs.

Types of Networks

The basic building block of an 802.11 network is the *basic service set* (BSS), which is simply a group of stations that communicate with each other. Communications take place within a somewhat fuzzy area, called the *basic service area*, defined by the propagation characteristics of the wireless medium.* When a station is in the basic

* All of the wireless media used will propagate in three dimensions. From that perspective, the service area should perhaps be called the service *volume*. However, the term area is widely used and accepted.

service area, it can communicate with the other members of the BSS. BSSs come in two flavors, both of which are illustrated in Figure 2-4.

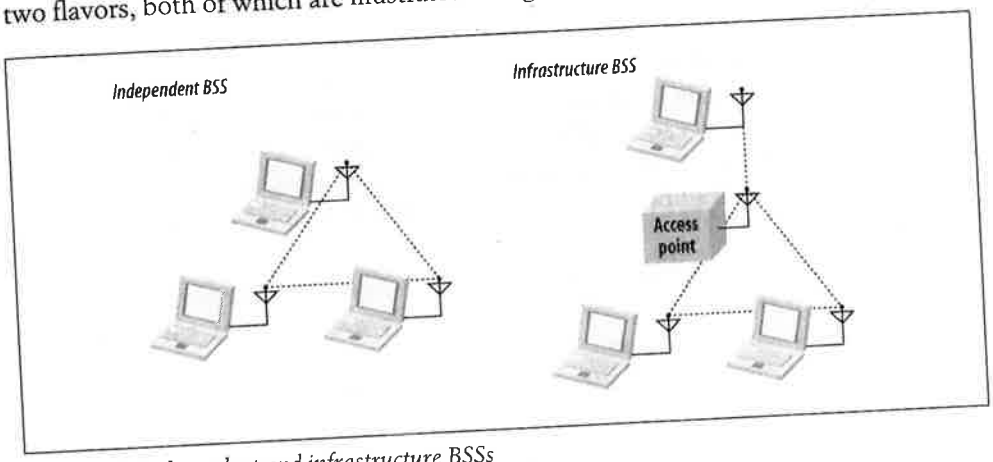

Figure 2-4. Independent and infrastructure BSSs

Independent networks

On the left is an *independent BSS* (IBSS). Stations in an IBSS communicate directly with each other and thus must be within direct communication range. The smallest possible 802.11 network is an IBSS with two stations. Typically, IBSSs are composed of a small number of stations set up for a specific purpose and for a short period of time. One common use is to create a short-lived network to support a single meeting in a conference room. As the meeting begins, the participants create an IBSS to share data. When the meeting ends, the IBSS is dissolved.* Due to their short duration, small size, and focused purpose, IBSSs are sometimes referred to as *ad hoc BSSs* or *ad hoc networks*.

Infrastructure networks

On the right side of Figure 2-4 is an *infrastructure BSS* (never called an IBSS). Infrastructure networks are distinguished by the use of an access point. Access points are used for all communications in infrastructure networks, including communication between mobile nodes in the same service area. If one mobile station in an infrastructure BSS needs to communicate with a second mobile station, the communication must take two hops. First, the originating mobile station transfers the frame to the access point. Second, the access point transfers the frame to the destination station. With all communications relayed through an access point, the basic service area corresponding to an infrastructure BSS is defined by the points in which transmissions from the access point can be received. Although the multihop transmission takes

* IBSSs have found a similar use at LAN parties throughout the world.

more transmission capacity than a directed frame from the sender to the receiver, it has two major advantages:

- An infrastructure BSS is defined by the distance from the access point. All mobile stations are required to be within reach of the access point, but no restriction is placed on the distance between mobile stations themselves. Allowing direct communication between mobile stations would save transmission capacity but at the cost of increased physical layer complexity because mobile stations would need to maintain neighbor relationships with all other mobile stations within the service area.
- Access points in infrastructure networks are in a position to assist with stations attempting to save power. Access points can note when a station enters a power-saving mode and buffer frames for it. Battery-operated stations can turn the wireless transceiver off and power it up only to transmit and retrieve buffered frames from the access point.

In an infrastructure network, stations must *associate* with an access point to obtain network services. Association is the process by which mobile station joins an 802.11 network; it is logically equivalent to plugging in the network cable on an Ethernet. It is not a symmetric process. Mobile stations always initiate the association process, and access points may choose to grant or deny access based on the contents of an association request. Associations are also exclusive on the part of the mobile station: a mobile station can be associated with only one access point.* The 802.11 standard places no limit on the number of mobile stations that an access point may serve. Implementation considerations may, of course, limit the number of mobile stations an access point may serve. In practice, however, the relatively low throughput of wireless networks is far more likely to limit the number of stations placed on a wireless network.

Extended service areas

BSSs can create coverage in small offices and homes, but they cannot provide network coverage to larger areas. 802.11 allows wireless networks of arbitrarily large size to be created by linking BSSs into an *extended service set* (ESS). An ESS is created by chaining BSSs together with a backbone network. 802.11 does not specify a particular backbone technology; it requires only that the backbone provide a specified set of services. In Figure 2-5, the ESS is the union of the four BSSs (provided that all the access points are configured to be part of the same ESS). In real-world deployments, the degree of overlap between the BSSs would probably be much greater than

* One reviewer noted that a similar restriction was present in traditional Ethernet networks until the development of VLANs and specifically asked how long this restriction was likely to last. I am not intimately involved with the standardization work, so I cannot speak to the issue directly. I do, however, agree that it is an interesting question.

the overlap in Figure 2-5. In real life, you would want to offer continuous coverage within the extended service area; you wouldn't want to require that users walk through the area covered by BSS3 when en route from BSS1 to BSS2.

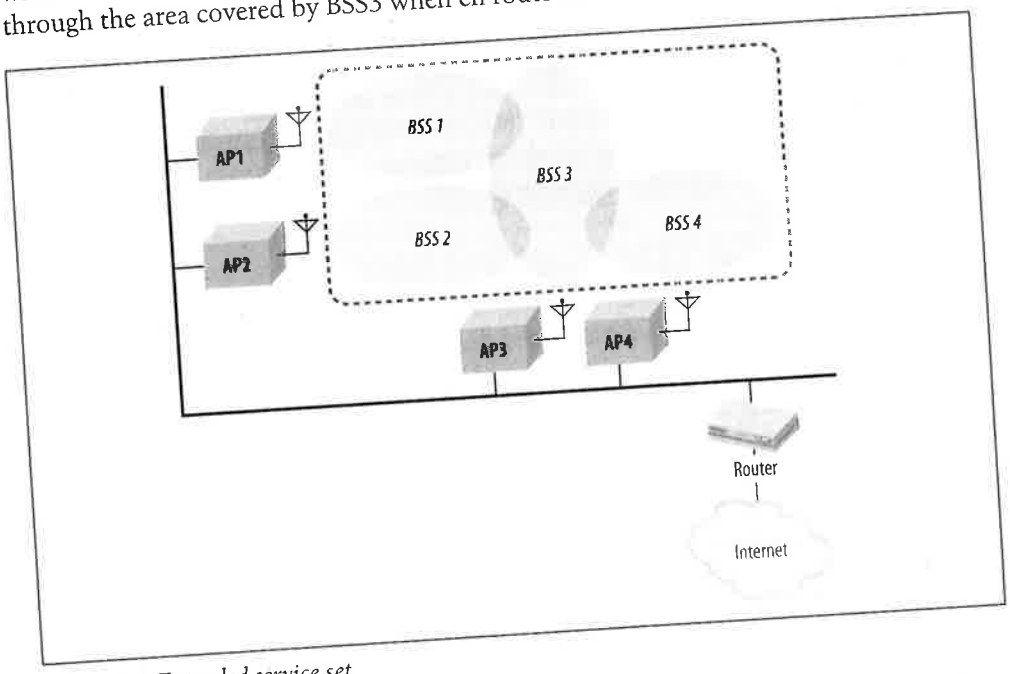

Figure 2-5. Extended service set

Stations within the same ESS may communicate with each other, even though these stations may be in different basic service areas and may even be moving between basic service areas. For stations in an ESS to communicate with each other, the wireless medium must act like a single layer 2 connection. Access points act as bridges, so direct communication between stations in an ESS requires that the backbone network also be a layer 2 connection. Any link-layer connection will suffice. Several access points in a single area may be connected to a single hub or switch, or they can use virtual LANs if the link-layer connection must span a large area.

802.11 supplies link-layer mobility within an ESS but only if the backbone network is a single link-layer domain, such as a shared Ethernet or a VLAN. This important constraint on mobility is often a major factor in 802.11 network design.

Extended service areas are the highest-level abstraction supported by 802.11 networks. Access points in an ESS operate in concert to allow the outside world to use a single MAC address to talk to a station somewhere within the ESS. In Figure 2-5, the router uses a single MAC address to deliver frames to a mobile station; the access point with which that mobile station is associated delivers the frame. The router

remains ignorant of the location of the mobile station and relies on the access points to deliver the frame.

The Distribution System, Revisited

With an understanding of how an extended service set is built, I'd like to return to the concept of the distribution system. 802.11 describes the distribution system in terms of the services it provides to wireless stations. While these services will be described in more detail later in this chapter, it is worth describing their operation at a high level.

The distribution system provides mobility by connecting access points. When a frame is given to the distribution system, it is delivered to the right access point and relayed by that access point to the intended destination.

The distribution system is responsible for tracking where a station is physically located and delivering frames appropriately. When a frame is sent to a mobile station, the distribution system is charged with the task of delivering it to the access point serving the mobile station. As an example, consider the router in Figure 2-5. The router simply uses the MAC address of a mobile station as its destination. The distribution system of the ESS pictured in Figure 2-5 must deliver the frame to the right access point. Obviously, part of the delivery mechanism is the backbone Ethernet, but the backbone network cannot be the entire distribution system because it has no way of choosing between access points. In the language of 802.11, the backbone Ethernet is the *distribution system medium*, but it is not the entire distribution system.

To find the rest of the distribution system, we need to look to the access points themselves. Most access points currently on the market operate as bridges. They have at least one wireless network interface and at least one Ethernet network interface. The Ethernet side can be connected to an existing network, and the wireless side becomes an extension of that network. Relaying frames between the two network media is controlled by a bridging engine.

Figure 2-6 illustrates the relationship between the access point, backbone network, and the distribution system. The access point has two interfaces connected by a bridging engine. Arrows indicate the potential paths to and from the bridging engine. Frames may be sent by the bridge to the wireless network; any frames sent by the bridge's wireless port are transmitted to all associated stations. Each associated station can transmit frames to the access point. Finally, the backbone port on the bridge can interact directly with the backbone network. The distribution system in Figure 2-6 is composed of the bridging engine plus the wired backbone network..

Every frame sent by a mobile station in an infrastructure network must use the distribution system. It is easy to understand why interaction with hosts on the backbone network must use the distribution system. After all, they are connected to the distribution

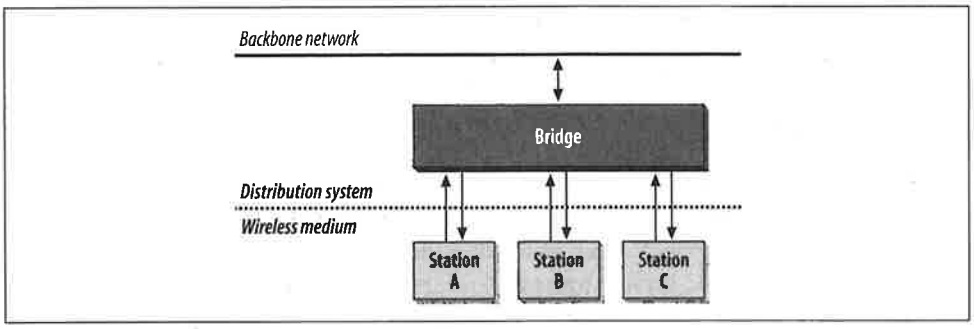

Figure 2-6. Distribution system in common 802.11 access point implementations

system medium. Wireless stations in an infrastructure network depend on the distribution system to communicate with each other because they are not directly connected to each other. The only way for station A to send a frame to station B is by relaying the frame through the bridging engine in the access point. However, the bridge is a component of the distribution system. While what exactly makes up the distribution system may seem like a narrow technical concern, there are some features of the 802.11 MAC that are closely tied to its interaction with the distribution system.

Inter-access point communication as part of the distribution system

Included with this distribution system is a method to manage associations. A wireless station is associated with only one access point at a time. If a station is associated with one access point, all the other access points in the ESS need to learn about that station. In Figure 2-5, AP4 must know about all the stations associated with AP1. If a wireless station associated with AP4 sends a frame to a station associated with AP1, the bridging engine inside AP4 must send the frame over the backbone Ethernet to AP1 so it can be delivered to its ultimate destination. To fully implement the distribution system, access points must inform other access points of associated stations. Naturally, many access points on the market use an *inter-access point protocol* (IAPP) over the backbone medium. There is, however, no standardized method for communicating association information to other members of an ESS. Proprietary technology is giving way to standardization, however. One of the major projects in the IEEE 802.11 working group is the standardization of the IAPP.

Wireless bridges and the distribution system

Up to this point, I have tacitly assumed that the distribution system was an existing fixed network. While this will often be the case, the 802.11 specification explicitly supports using the wireless medium itself as the distribution system. The wireless distribution system configuration is often called a "wireless bridge" configuration because it allows network engineers to connect two LANs at the link layer. Wireless bridges can be used to quickly connect distinct physical locations and are well-suited for use by access providers. Most 802.11 access points on the market now support

the wireless bridge configuration, though it may be necessary to upgrade the firmware on older units.

Network Boundaries

Because of the nature of the wireless medium, 802.11 networks have fuzzy boundaries. In fact, some degree of fuzziness is desirable. As with mobile telephone networks, allowing basic service areas to overlap increases the probability of successful transitions between basic service areas and offers the highest level of network coverage. The basic service areas on the right of Figure 2-7 overlap significantly. This means that a station moving from BSS2 to BSS4 is not likely to lose coverage; it also means that AP3 (or, for that matter, AP4) can fail without compromising the network too badly. On the other hand, if AP2 fails, the network is cut into two disjoint parts, and stations in BSS1 lose connectivity when moving out of BSS1 and into BSS3 or BSS4.

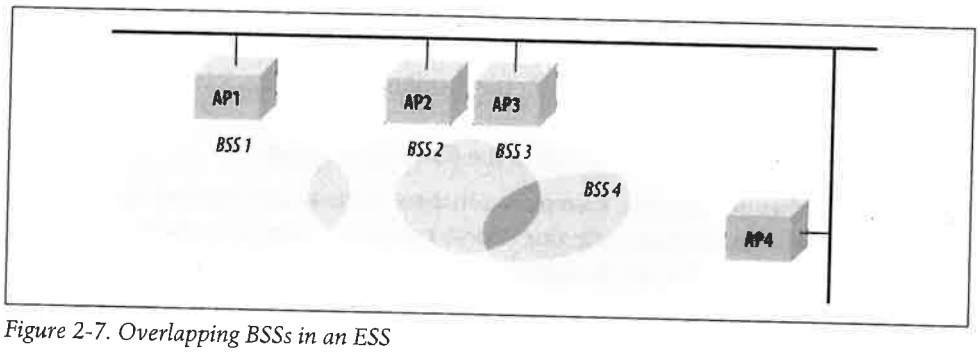

Figure 2-7. Overlapping BSSs in an ESS

Different types of 802.11 networks may also overlap. Independent BSSs may be created within the basic service area of an access point. Figure 2-8 illustrates spatial overlap. An access point appears at the top of the figure; its basic service area is shaded. Two stations are operating in infrastructure mode and communicate only with the access point. Three stations have been set up as an independent BSS and communicate with each other. Although the five stations are assigned to two different BSSs, they may share the same wireless medium. Stations may obtain access to the medium only by using the rules specified in the 802.11 MAC; these rules were carefully designed to enable multiple 802.11 networks to coexist in the same spatial area. Both BSSs must share the capacity of a single radio channel, so there may be adverse performance implications from co-located BSSs.

802.11 Network Operations

From the outset, 802.11 was designed to be just another link layer to higher-layer protocols. Network administrators familiar with Ethernet will be immediately comfortable

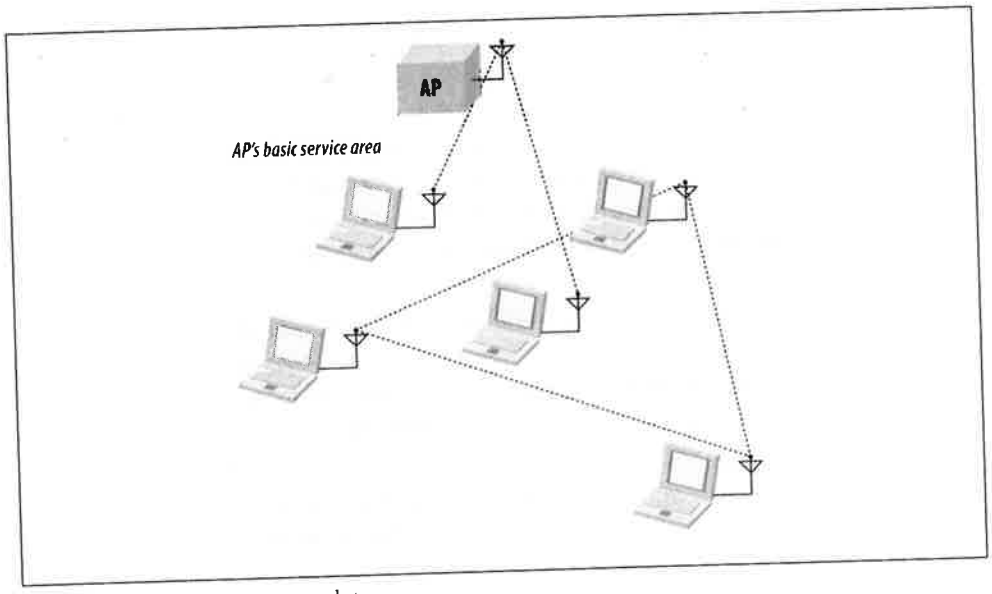

Figure 2-8. Overlapping network types

with 802.11. The shared heritage is deep enough that 802.11 is sometimes referred to as "wireless Ethernet."

The core elements present in Ethernet are present in 802.11. Stations are identified by 48-bit IEEE 802 MAC addresses. Conceptually, frames are delivered based on the MAC address. Frame delivery is unreliable, though 802.11 incorporates some basic reliability mechanisms to overcome the inherently poor qualities of the radio channels it uses.*

From a user's perspective, 802.11 might just as well be Ethernet. Network administrators, however, need to be conversant with 802.11 at a much deeper level. Providing MAC-layer mobility while following the path blazed by previous 802 standards requires a number of additional services and more complex framing.

Network Services

One way to define a network technology is to define the services it offers and allow equipment vendors to implement those services in whatever way they see fit. 802.11 provides nine services. Only three of the services are used for moving data; the remaining six are management operations that allow the network to keep track of the mobile nodes and deliver frames accordingly.

* I don't mean "poor" in an absolute sense. But the reliability of wireless transmission is really not comparable to the reliability of a wired network.

The services are described in the following list and summarized in Table 2-1:

Distribution

This service is used by mobile stations in an infrastructure network every time they send data. Once a frame has been accepted by an access point, it uses the distribution service to deliver the frame to its destination. Any communication that uses an access point travels through the distribution service, including communications between two mobile stations associated with the same access point.

Integration

Integration is a service provided by the distribution system; it allows the connection of the distribution system to a non-IEEE 802.11 network. The integration function is specific to the distribution system used and therefore is not specified by 802.11, except in terms of the services it must offer.

Association

Delivery of frames to mobile stations is made possible because mobile stations register, or associate, with access points. The distribution system can then use the registration information to determine which access point to use for any mobile station. Unassociated stations are not "on the network," much like workstations with unplugged Ethernet cables. 802.11 specifies the function that must be provided by the distribution system using the association data, but it does not mandate any particular implementation.

Reassociation

When a mobile station moves between basic service areas within a single extended service area, it must evaluate signal strength and perhaps switch the access point with which it is associated. Reassociations are initiated by mobile stations when signal conditions indicate that a different association would be beneficial; they are never initiated by the access point. After the reassociation is complete, the distribution system updates its location records to reflect the reachability of the mobile station through a different access point.

Disassociation

To terminate an existing association, stations may use the disassociation service. When stations invoke the disassociation service, any mobility data stored in the distribution system is removed. Once disassociation is complete, it is as if the station is no longer attached to the network. Disassociation is a polite task to do during the station shutdown process. The MAC is, however, designed to accommodate stations that leave the network without formally disassociating.

Authentication

Physical security is a major component of a wired LAN security solution. Network attachment points are limited, often to areas in offices behind perimeter access control devices. Network equipment can be secured in locked wiring closets, and data jacks in offices and cubicles can be connected to the network only when needed. Wireless networks cannot offer the same level of physical security,

however, and therefore must depend on additional authentication routines to ensure that users accessing the network are authorized to do so. Authentication is a necessary prerequisite to association because only authenticated users are authorized to use the network. (In practice, though, many access points are configured for "open-system" authentication and will authenticate any station.)

Deauthentication

Deauthentication terminates an authenticated relationship. Because authentication is needed before network use is authorized, a side effect of deauthentication is termination of any current association.

Privacy

Strong physical controls can prevent a great number of attacks on the privacy of data in a wired LAN. Attackers must obtain physical access to the network medium before attempting to eavesdrop on traffic. On a wired network, physical access to the network cabling is a subset of physical access to other computing resources. By design, physical access to wireless networks is a comparatively simpler matter of using the correct antenna and modulation methods. To offer a similar level of privacy, 802.11 provides an optional privacy service called Wired Equivalent Privacy (WEP). WEP is not ironclad security—in fact, it has been proven recently that breaking WEP is easily within the capabilities of any laptop (for more information, see Chapter 5). Its purpose is to provide roughly equivalent privacy to a wired network by encrypting frames as they travel across the 802.11 air interface. Depending on your level of cynicism, you may or may not think that WEP achieves its goal; after all, it's not that hard to access the Ethernet cabling in a traditional network. In any case, do not assume that WEP provides more than minimal security. It prevents other users from casually appearing on your network, but that's about all.*

MSDU delivery

Networks are not much use without the ability to get the data to the recipient. Stations provide the MAC Service Data Unit (MSDU) delivery service, which is responsible for getting the data to the actual endpoint.

Table 2-1. Network services

Service	Station or distribution service?	Description
Distribution	Distribution	Service used in frame delivery to determine destination address in infrastructure networks
Integration	Distribution	Frame delivery to an IEEE 802 LAN outside the wireless network

* One of O'Reilly's offices had a strange situation in which apparent "interlopers" appeared on the network. They eventually discovered that their ESS overlapped a company in a neighboring office building, and "foreign" laptops were simply associating with the access point that had the strongest signal. WEP solves problems like this but will not withstand a deliberate attack on your network.

Table 2-1. Network services (continued)

Service	Station or distribution service?	Description
Association	Distribution	Used to establish the AP which serves as the gateway to a particular mobile station
Reassociation	Distribution	Used to change the AP which serves as the gateway to a particular mobile station
Disassociation	Distribution	Removes the wireless station from the network
Authentication	Station	Establishes identity prior to establishing association
Deauthentication	Station	Used to terminate authentication, and by extension, association
Privacy	Station	Provides protection against eavesdropping
MSDU delivery	Station	Delivers data to the recipient

Station services

Station services are part of every 802.11-compliant station and must be incorporated by any product claiming 802.11 compliance. Station services are provided by both mobile stations and the wireless interface on access points. Stations provide frame delivery services to allow message delivery, and, in support of this task, they may need to use the authentication services to establish associations. Stations may also wish to take advantage of privacy functions to protect messages as they traverse the vulnerable wireless link.

Distribution system services

Distribution system services connect access points to the distribution system. The major role of access points is to extend the services on the wired network to the wireless network; this is done by providing the distribution and integration services to the wireless side. Managing mobile station associations is the other major role of the distribution system. To maintain association data and station location information, the distribution system provides the association, reassociation, and disassociation services.

Mobility Support

Mobility is the major motivation for deploying an 802.11 network. Stations can move while connected to the network and transmit frames while in motion. Mobility can cause one of three types of transition:

No transition
: When stations do not move out of their current access point's service area, no transition is necessary. This state occurs because the station is not moving or it is

moving within the basic service area of its current access point.* (Arguably, this isn't a transition so much as the absence of a transition, but it is defined in the specification.)

BSS transition

Stations continuously monitor the signal strength and quality from all access points administratively assigned to cover an extended service area. Within an extended service area, 802.11 provides MAC layer mobility. Stations attached to the distribution system can send out frames addressed to the MAC address of a mobile station and let the access points handle the final hop to the mobile station. Distribution system stations do not need to be aware of a mobile station's location as long as it is within the same extended service area.

Figure 2-9 illustrates a BSS transition. The three access points in the picture are all assigned to the same ESS. At the outset, denoted by *t=1*, the laptop with an 802.11 network card is sitting within AP1's basic service area and is associated with AP1. When the laptop moves out of AP1's basic service area and into AP2's at *t=2*, a BSS transition occurs. The mobile station uses the reassociation service to associate with AP2, which then starts sending frames to the mobile station.

BSS transitions require the cooperation of access points. In this scenario, AP2 needs to inform AP1 that the mobile station is now associated with AP2. 802.11 does not specify the details of the communications between access points during BSS transitions. A standardized IAPP is a likely result of future work within the 802.11 working group.

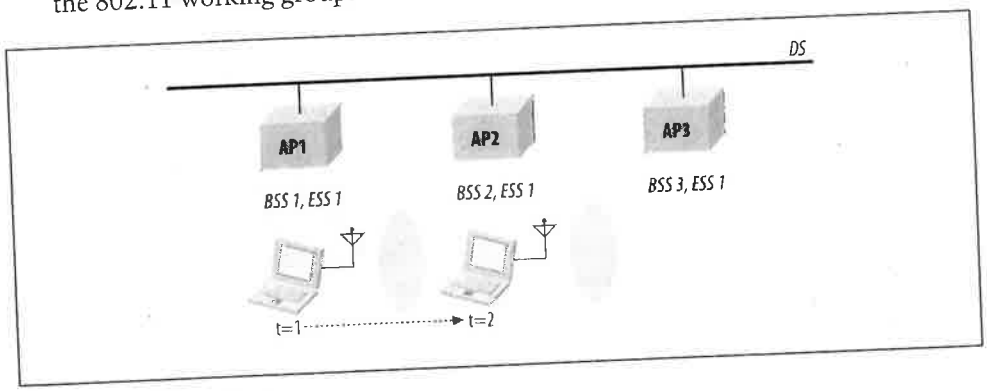

Figure 2-9. BSS transition

* Although my explanation makes it sound as if the "no motion" and "local motion" substates are easily distinguishable, they are not. The underlying physics of RF propagation can make it impossible to tell whether a station is moving because the signal strength can vary with the placement of objects in the room, which, of course, includes the people who may be walking around.

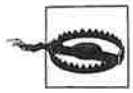

Because inter-access point communications are not standardized, mobility between access points supplied by different vendors is not guaranteed.

ESS transition

An ESS transition refers to the movement from one ESS to a second distinct ESS. 802.11 does not support this type of transition, except to allow the station to associate with an access point in the second ESS once it leaves the first. Higher-layer connections are almost guaranteed to be interrupted. It would be fair to say that 802.11 supports ESS transitions only to the extent that it is relatively easy to attempt associating with an access point in the new extended service area. Maintaining higher-level connections requires support from the protocol suites in question. In the case of TCP/IP, Mobile IP is required to seamlessly support an ESS transition.

Figure 2-10 illustrates an ESS transition. Four basic service areas are organized into two extended service areas. Seamless transitions from the lefthand ESS to the righthand ESS are not supported. ESS transitions are supported only because the mobile station will quickly associate with an access point in the second ESS. Any active network connections are likely to be dropped when the mobile station leaves the first ESS.

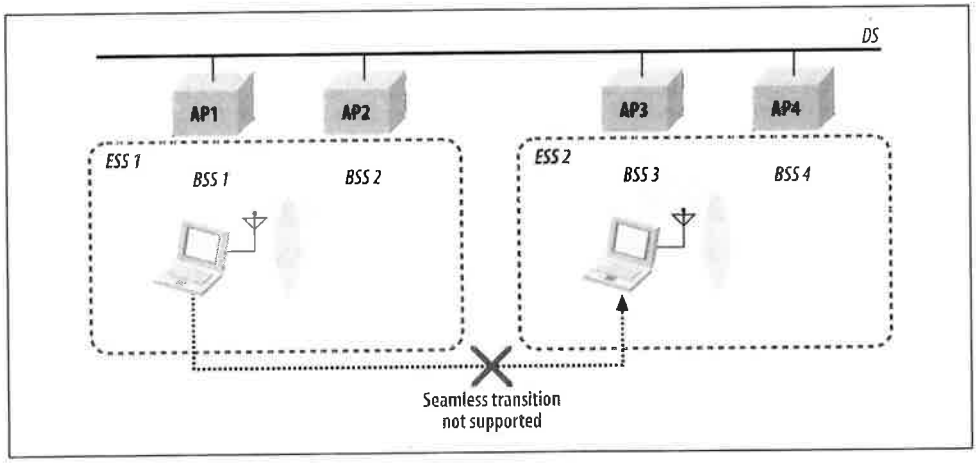

Figure 2-10. ESS transition

CHAPTER 3
The 802.11 MAC

This chapter begins our exploration of the 802.11 standard in depth. Chapter 2 provided a high-level overview of the standard and discussed some of its fundamental attributes. You are now at a fork in the book. Straight ahead lies a great deal of information on the 802.11 specifications. It is possible, however, to build a wired network without a thorough and detailed understanding of the protocols, and the same is true for wireless networks. However, there are a number of situations in which you may need a deeper knowledge of the machinery under the hood:

- Although 802.11 has been widely and rapidly adopted, security issues have continued to grab headlines. Network managers will undoubtedly be asked to comment on security issues, especially in any wireless LAN proposals. To understand and participate in these discussions, read Chapter 5. As I write this, WEP has been fully broken and the IEEE is forging a successor to it based on 802.1x.* Though the final form of the new and improved security framework has not yet become apparent, it will almost surely be based on 802.1x, which is described in Chapter 6.
- Troubleshooting wireless networks is similar to troubleshooting wired networks but can be much more complex. As always, a trusty packet sniffer can be an invaluable aid. To take full advantage of a packet sniffer, though, you need to understand what the packets mean to interpret your network's behavior.
- Tuning a wireless network is tied intimately to a number of parameters in the specification. To understand the behavior of your network and what effect the optimizations will have requires a knowledge of what those parameters really do.
- Device drivers may expose low-level knobs and dials for you to play with. Most drivers provide good defaults for all of the parameters, but some give you freedom to experiment. Open source software users have the source code and are free to experiment with any and all settings.

* And as we go to press, 802.1x has reportedly been broken.

- A number of interesting features of the standard have not been implemented by the current products, but they may be implemented later. As these features are rolled out, you may need to know what they are and how to use them.

As with many other things in life, the more you know, the better off you are. Ethernet is usually trouble-free, but serious network administrators have long known that when you do run into trouble, there is no substitute for thorough knowledge of how the network is working. To some extent, 802.11 networks have had a "free ride" the past few years. Because they were cool, users were forgiving when they failed; wireless connectivity was a privilege, not a right. And since there were relatively few networks and relatively few users on those networks, the networks were rarely subjected to severe stresses. An Ethernet that has only a half dozen nodes is not likely to be a source of problems; problems occur when you add a few high-capacity servers, a few hundred users, and the associated bridges and routers to glue everything together. There is no reason to believe that wireless will be any different. A couple of access points serving a half dozen users will not reveal any problems. But when the user community grows to a few dozen, and you have several overlapping wireless networks, each with its own set of access points, you can expect to see the effects of stress.

That is why you should read this chapter. Now on to the details.

The key to the 802.11 specification is the MAC. It rides on every physical layer and controls the transmission of user data into the air. It provides the core framing operations and the interaction with a wired network backbone. Different physical layers may provide different transmission speeds, all of which are supposed to interoperate.

802.11 does not depart from the previous IEEE 802 standards in any radical way. The standard successfully adapts Ethernet-style networking to radio links. Like Ethernet, 802.11 uses a carrier sense multiple access (CSMA) scheme to control access to the transmission medium. However, collisions waste valuable transmission capacity, so rather than the collision detection (CSMA/CD) employed by Ethernet, 802.11 uses collision avoidance (CSMA/CA). Also like Ethernet, 802.11 uses a distributed access scheme with no centralized controller. Each 802.11 station uses the same method to gain access to the medium. The major differences between 802.11 and Ethernet stem from the differences in the underlying medium.

This chapter provides some insight into the motivations of the MAC designers by describing some challenges they needed to overcome and describes the rules used for access to the medium, as well as the basic frame structure. If you simply want to understand the basic frame sequences that you will see on an 802.11 network, skip ahead to the end of this chapter. For further information on the MAC, consult its formal specification in Clause 9 of the 802.11 standard; detailed MAC state diagrams are in Annex C.

Challenges for the MAC

Differences between the wireless network environment and the traditional wired environment create challenges for network protocol designers. This section examines a number of the hurdles that the 802.11 designers faced.

RF Link Quality

On a wired Ethernet, it is reasonable to transmit a frame and assume that the destination receives it correctly. Radio links are different, especially when the frequencies used are unlicensed ISM bands. Even narrowband transmissions are subject to noise and interference, but unlicensed devices must assume that interference will exist and work around it. The designers of 802.11 considered ways to work around the radiation from microwave ovens and other RF sources. In addition to the noise, multipath fading may also lead to situations in which frames cannot be transmitted because a node moves into a dead spot.

Unlike many other link layer protocols, 802.11 incorporates positive acknowledgments. All transmitted frames must be acknowledged, as shown in Figure 3-1. If any part of the transfer fails, the frame is considered lost.

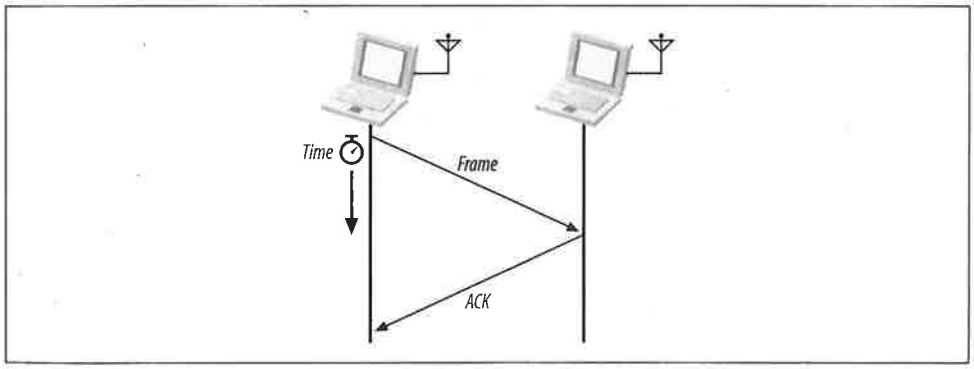

Figure 3-1. Positive acknowledgment of data transmissions

The sequence in Figure 3-1 is an *atomic* operation. 802.11 allows stations to lock out contention during atomic operations so that atomic sequences are not interrupted by other stations attempting to use the transmission medium.

The Hidden Node Problem

In Ethernet networks, stations depend on the reception of transmissions to perform the carrier sensing functions of CSMA/CD. Wires in the physical medium contain the signals and distribute them to network nodes. Wireless networks have fuzzier

boundaries, sometimes to the point where each node may not be able to communicate with every other node in the wireless network, as in Figure 3-2.

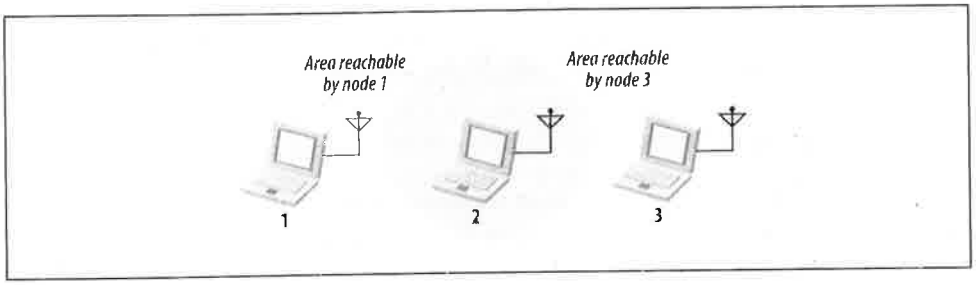

Figure 3-2. Nodes 1 and 3 are "hidden"

In the figure, node 2 can communicate with both nodes 1 and 3, but something prevents nodes 1 and 3 from communicating directly. (The obstacle itself is not relevant; it could be as simple as nodes 1 and 3 being as far away from 2 as possible, so the radio waves cannot reach the full distance from 1 to 3.) From the perspective of node 1, node 3 is a "hidden" node. If a simple transmit-and-pray protocol was used, it would be easy for node 1 and node 3 to transmit simultaneously, thus rendering node 2 unable to make sense of anything. Furthermore, nodes 1 and 3 would not have any indication of the error because the collision was local to node 2.

Collisions resulting from hidden nodes may be hard to detect in wireless networks because wireless transceivers are generally half-duplex; they don't transmit and receive at the same time. To prevent collisions, 802.11 allows stations to use Request to Send (RTS) and Clear to Send (CTS) signals to clear out an area. Figure 3-3 illustrates the procedure.

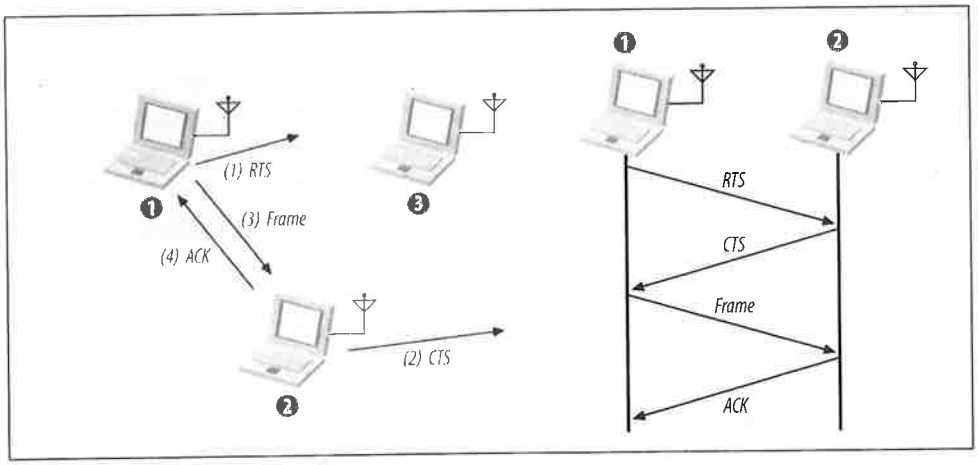

Figure 3-3. RTS/CTS clearing

In Figure 3-3, node 1 has a frame to send; it initiates the process by sending an RTS frame. The RTS frame serves several purposes: in addition to reserving the radio link for transmission, it silences any stations that hear it. If the target station receives an RTS, it responds with a CTS. Like the RTS frame, the CTS frame silences stations in the immediate vicinity. Once the RTS/CTS exchange is complete, node 1 can transmit its frames without worry of interference from any hidden nodes. Hidden nodes beyond the range of the sending station are silenced by the CTS from the receiver. When the RTS/CTS clearing procedure is used, any frames must be positively acknowledged.

The multiframe RTS/CTS transmission procedure consumes a fair amount of capacity, especially because of the additional latency incurred before transmission can commence. As a result, it is used only in high-capacity environments and environments with significant contention on transmission. For lower-capacity environments, it is not necessary.

You can control the RTS/CTS procedure by setting the *RTS threshold* if the device driver for your 802.11 card allows you to adjust it. The RTS/CTS exchange is performed for frames larger than the threshold. Frames shorter than the threshold are simply sent.

MAC Access Modes and Timing

Access to the wireless medium is controlled by coordination functions. Ethernet-like CSMA/CA access is provided by the distributed coordination function (DCF). If contention-free service is required, it can be provided by the point coordination function (PCF), which is built on top of the DCF. Contention-free services are provided only in infrastructure networks. The coordination functions are described in the following list and illustrated in Figure 3-4:

DCF
: The DCF is the basis of the standard CSMA/CA access mechanism. Like Ethernet, it first checks to see that the radio link is clear before transmitting. To avoid collisions, stations use a random backoff after each frame, with the first transmitter seizing the channel. In some circumstances, the DCF may use the CTS/RTS clearing technique to further reduce the possibility of collisions.

PCF
: Point coordination provides contention-free services. Special stations called *point coordinators* are used to ensure that the medium is provided without contention. Point coordinators reside in access points, so the PCF is restricted to infrastructure networks. To gain priority over standard contention-based services, the PCF allows stations to transmit frames after a shorter interval. The PCF is not widely implemented and is described in Chapter 8.

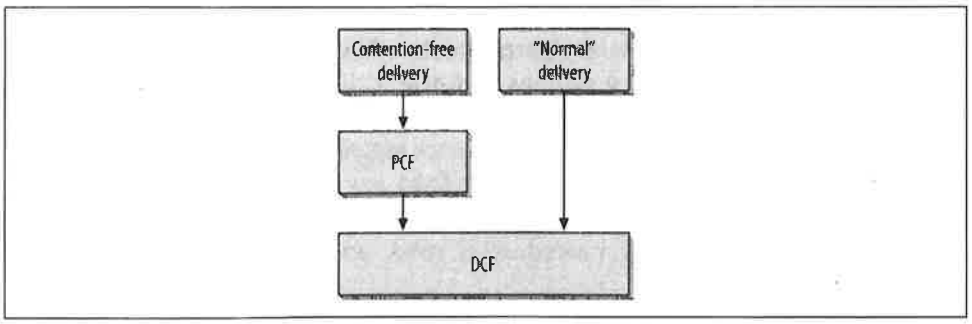

Figure 3-4. MAC coordination functions

Carrier-Sensing Functions and the Network Allocation Vector

Carrier sensing is used to determine if the medium is available. Two types of carrier-sensing functions in 802.11 manage this process: the physical carrier-sensing and virtual carrier-sensing functions. If either carrier-sensing function indicates that the medium is busy, the MAC reports this to higher layers.

Physical carrier-sensing functions are provided by the physical layer in question and depend on the medium and modulation used. It is difficult (or, more to the point, expensive) to build physical carrier-sensing hardware for RF-based media, because transceivers can transmit and receive simultaneously only if they incorporate expensive electronics. Furthermore, with hidden nodes potentially lurking everywhere, physical carrier-sensing cannot provide all the necessary information.

Virtual carrier-sensing is provided by the Network Allocation Vector (NAV). Most 802.11 frames carry a duration field, which can be used to reserve the medium for a fixed time period. The NAV is a timer that indicates the amount of time the medium will be reserved. Stations set the NAV to the time for which they expect to use the medium, including any frames necessary to complete the current operation. Other stations count down from the NAV to 0. When the NAV is nonzero, the virtual carrier-sensing function indicates that the medium is busy; when the NAV reaches 0, the virtual carrier-sensing function indicates that the medium is idle.

By using the NAV, stations can ensure that atomic operations are not interrupted. For example, the RTS/CTS sequence in Figure 3-3 is atomic. Figure 3-5 shows how the NAV protects the sequence from interruption. (This is a standard format for a number of diagrams in this book that illustrate the interaction of multiple stations with the corresponding timers.) Activity on the medium by stations is represented by the shaded bars, and each bar is labeled with the frame type. Interframe spacing is depicted by the lack of any activity. Finally, the NAV timer is represented by the bars on the NAV line at the bottom of the figure. The NAV is carried in the frame headers on the RTS and CTS frames; it is depicted on its own line to show how the NAV

relates to actual transmissions in the air. When a NAV bar is present on the NAV line, stations should defer access to the medium because the virtual carrier-sensing mechanism will indicate a busy medium.

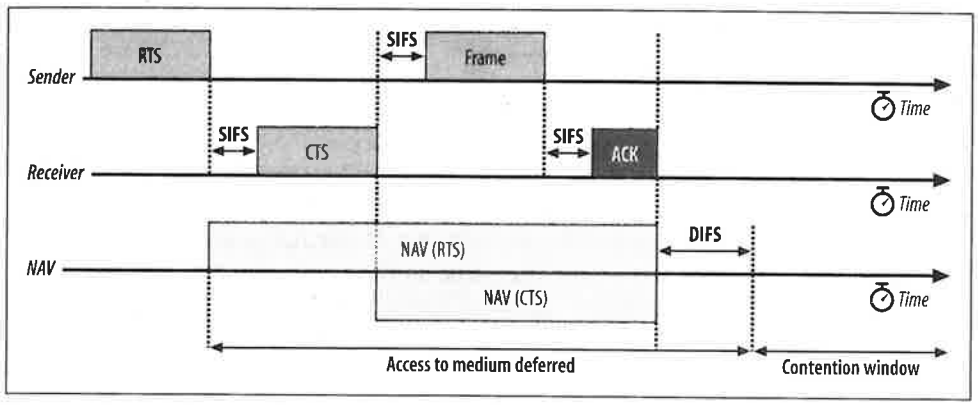

Figure 3-5. Using the NAV for virtual carrier sensing

To ensure that the sequence is not interrupted, node 1 sets the NAV in its RTS to block access to the medium while the RTS is being transmitted. All stations that hear the RTS defer access to the medium until the NAV elapses.

RTS frames are not necessarily heard by every station in the network. Therefore, the recipient of the intended transmission responds with a CTS that includes a shorter NAV. This NAV prevents other stations from accessing the medium until the transmission completes. After the sequence completes, the medium can be used by any station after distributed interframe space (DIFS), which is depicted by the contention window beginning at the right side of the figure.

RTS/CTS exchanges may be useful in crowded areas with multiple overlapping networks. Every station on the same physical channel receives the NAV and defers access appropriately, even if the stations are configured to be on different networks.

Interframe Spacing

As with traditional Ethernet, the interframe spacing plays a large role in coordinating access to the transmission medium. 802.11 uses four different interframe spaces. Three are used to determine medium access; the relationship between them is shown in Figure 3-6.

We've already seen that as part of the collision avoidance built into the 802.11 MAC, stations delay transmission until the medium becomes idle. Varying interframe spacings create different priority levels for different types of traffic. The logic behind this is simple: high-priority traffic doesn't have to wait as long after the medium has become idle. Therefore, if there is any high-priority traffic waiting, it grabs the network before low-priority frames have a chance to try. To assist with interoperability

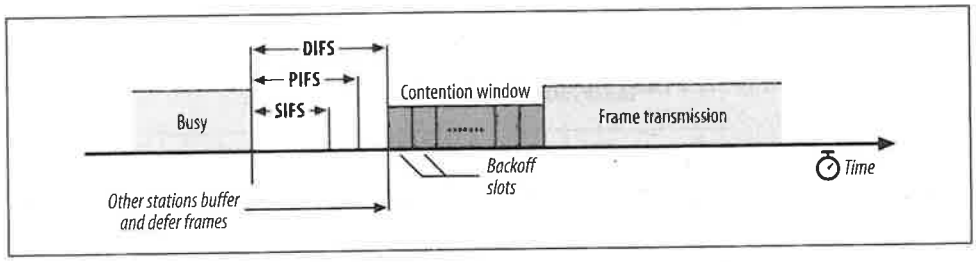

Figure 3-6. Interframe spacing relationships

between different data rates, the interframe space is a fixed amount of time, independent of the transmission speed. (This is only one of the many problems caused by having different physical layers use the same radio resources, which are different modulation techniques.) Different physical layers, however, can specify different interframe space times.

Short interframe space (SIFS)
: The SIFS is used for the highest-priority transmissions, such as RTS/CTS frames and positive acknowledgments. High-priority transmissions can begin once the SIFS has elapsed. Once these high-priority transmissions begin, the medium becomes busy, so frames transmitted after the SIFS has elapsed have priority over frames that can be transmitted only after longer intervals.

PCF interframe space (PIFS)
: The PIFS, sometimes erroneously called the priority interframe space, is used by the PCF during contention-free operation. Stations with data to transmit in the contention-free period can transmit after the PIFS has elapsed and preempt any contention-based traffic.

DCF interframe space (DIFS)
: The DIFS is the minimum medium idle time for contention-based services. Stations may have immediate access to the medium if it has been free for a period longer than the DIFS.

Extended interframe space (EIFS)
: The EIFS is not illustrated in Figure 3-6 because it is not a fixed interval. It is used only when there is an error in frame transmission.

Interframe spacing and priority

Atomic operations start like regular transmissions: they must wait for the DIFS before they can begin. However, the second and any subsequent steps in an atomic operation take place using the SIFS, rather than during the DIFS. This means that the second (and subsequent) parts of an atomic operation will grab the medium before another type of frame can be transmitted. By using the SIFS and the NAV, stations can seize the medium for as long as necessary.

In Figure 3-5, for example, the short interframe space is used between the different units of the atomic exchange. After the sender gains access to the medium, the receiver replies with a CTS after the SIFS. Any stations that might attempt to access the medium at the conclusion of the RTS would wait for one DIFS interval. Partway through the DIFS interval, though, the SIFS interval elapses, and the CTS is transmitted.

Contention-Based Access Using the DCF

Most traffic uses the DCF, which provides a standard Ethernet-like contention-based service. The DCF allows multiple independent stations to interact without central control, and thus may be used in either IBSS networks or in infrastructure networks.

Before attempting to transmit, each station checks whether the medium is idle. If the medium is not idle, stations defer to each other and employ an orderly exponential backoff algorithm to avoid collisions.

In distilling the 802.11 MAC rules, there is a basic set of rules that are always used, and additional rules may be applied depending on the circumstances. Two basic rules apply to all transmissions using the DCF:

1. If the medium has been idle for longer than the DIFS, transmission can begin immediately. Carrier sensing is performed using both a physical medium-dependent method and the virtual (NAV) method.
    a. If the previous frame was received without errors, the medium must be free for at least the DIFS.
    b. If the previous transmission contained errors, the medium must be free for the amount of the EIFS.
2. If the medium is busy, the station must wait for the channel to become idle. 802.11 refers to the wait as *access deferral*. If access is deferred, the station waits for the medium to become idle for the DIFS and prepares for the exponential backoff procedure.

Additional rules may apply in certain situations. Many of these rules depend on the particular situation "on the wire" and are specific to the results of previous transmissions.

1. Error recovery is the responsibility of the station sending a frame. Senders expect acknowledgments for each transmitted frame and are responsible for retrying the transmission until it is successful.
    a. Positive acknowledgments are the only indication of success. Atomic exchanges must complete in their entirety to be successful. If an acknowledgment is expected and does not arrive, the sender considers the transmission lost and must retry.
    b. All unicast data must be acknowledged.

c. Any failure increments a retry counter, and the transmission is retried. A failure can be due to a failure to gain access to the medium or a lack of an acknowledgment. However, there is a longer congestion window when transmissions are retried (see next section).

2. Multiframe sequences may update the NAV with each step in the transmission procedure. When a station receives a medium reservation that is longer than the current NAV, it updates the NAV. Setting the NAV is done on a frame-by-frame basis and is discussed in much more detail in the next chapter.
3. The following types of frames can be transmitted after the SIFS and thus receive maximum priority: acknowledgments, the CTS in an RTS/CTS exchange sequence, and fragments in fragment sequences.
   a. Once a station has transmitted the first frame in a sequence, it has gained control of the channel. Any additional frames and their acknowledgments can be sent using the short interframe space, which locks out any other stations.
   b. Additional frames in the sequence update the NAV for the expected additional time the medium will be used.
4. Extended frame sequences are required for higher-level packets that are larger than configured thresholds.
   a. Packets larger than the RTS threshold must have RTS/CTS exchange.
   b. Packets larger than the fragmentation threshold must be fragmented.

Error Recovery with the DCF

Error detection and correction is up to the station that begins an atomic frame exchange. When an error is detected, the station with data must resend the frame. Errors must be detected by the sending station. In some cases, the sender can infer frame loss by the lack of a positive acknowledgment from the receiver. Retry counters are incremented when frames are retransmitted.

Each frame or fragment has a single retry counter associated with it. Stations have two retry counters: the *short retry count* and the *long retry count*. Frames that are shorter than the RTS threshold are considered to be short; frames longer than the threshold are long. Depending on the length of the frame, it is associated with either a short or long retry counter. Frame retry counts begin at 0 and are incremented when a frame transmission fails.

The short retry count is reset to 0 when:

- A CTS frame is received in response to a transmitted RTS
- A MAC-layer acknowledgment is received after a nonfragmented transmission
- A broadcast or multicast frame is received

The long retry count is reset to 0 when:

- A MAC-layer acknowledgment is received for a frame longer than the RTS threshold
- A broadcast or multicast frame is received

In addition to the associated retry count, fragments are given a maximum "lifetime" by the MAC. When the first fragment is transmitted, the lifetime counter is started. When the lifetime limit is reached, the frame is discarded and no attempt is made to transmit any remaining fragments.

Using the retry counters

Like most other network protocols, 802.11 provides reliability through retransmission. Data transmission happens within the confines of an atomic sequence, and the entire sequence must complete for a transmission to be successful. When a station transmits a frame, it must receive an acknowledgment from the receiver or it will consider the transmission to have failed. Failed transmissions increment the retry counter associated with the frame (or fragment). If the retry limit is reached, the frame is discarded, and its loss is reported to higher-layer protocols.

One of the reasons for having short frames and long frames is to allow network administrators to customize the robustness of the network for different frame lengths. Large frames require more buffer space, so one potential application of having two separate retry limits is to decrease the long retry limit to decrease the amount of buffer space required.

Backoff with the DCF

After frame transmission has completed and the DIFS has elapsed, stations may attempt to transmit congestion-based data. A period called the *contention window* or *backoff window* follows the DIFS. This window is divided into slots. Slot length is medium-dependent; higher-speed physical layers use shorter slot times. Stations pick a random slot and wait for that slot before attempting to access the medium; all slots are equally likely selections. When several stations are attempting to transmit, the station that picks the first slot (the station with the lowest random number) wins.

As in Ethernet, the backoff time is selected from a larger range each time a transmission fails. Figure 3-7 illustrates the growth of the contention window as the number of transmissions increases, using the numbers from the direct-sequence spread-spectrum (DSSS) physical layer. Other physical layers use different sizes, but the principle is identical. Contention window sizes are always 1 less than a power of 2 (e.g., 31, 63, 127, 255). Each time the retry counter increases, the contention window moves to the next greatest power of two. The size of the contention window is limited by the

physical layer. For example, the DS physical layer limits the contention window to 1023 transmission slots.

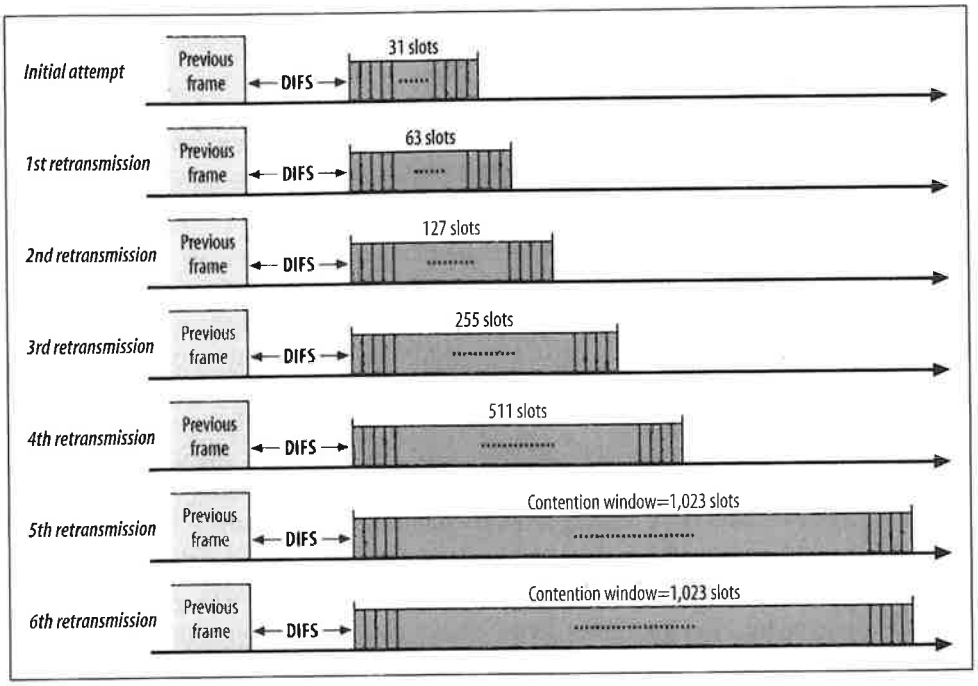

Figure 3-7. DSSS contention window size

When the contention window reaches its maximum size, it remains there until it can be reset. Allowing long contention windows when several competing stations are attempting to gain access to the medium keeps the MAC algorithms stable even under maximum load. The contention window is reset to its minimum size when frames are transmitted successfully, or the associated retry counter is reached, and the frame is discarded.

Fragmentation and Reassembly

Higher-level packets and some large management frames may need to be broken into smaller pieces to fit through the wireless channel. Fragmentation may also help improve reliability in the presence of interference. The primary sources of interference with 802.11 LANs are microwave ovens, with which they share the 2.4-GHz ISM band. Electromagnetic radiation is generated by the magnetron tube during its ramp-up and ramp-down, so microwaves emit interference half the time.*

* In the US, appliances are powered by 60-Hz alternating current, so microwaves interfere for about 8 milliseconds (ms) out of every 16-ms cycle. Much of the rest of the world uses 50-Hz current, and interference takes place for 10 ms out of the 20-ms cycle.

Wireless LAN stations may attempt to fragment transmissions so that interference affects only small fragments, not large frames. By immediately reducing the amount of data that can be corrupted by interference, fragmentation may result in a higher effective throughput.

Fragmentation takes place when a higher-level packet's length exceeds the fragmentation threshold configured by the network administrator. Fragments all have the same frame sequence number but have ascending fragment numbers to aid in reassembly. Frame control information also indicates whether more fragments are coming. All of the fragments that comprise a frame are normally sent in a *fragmentation burst*, which is shown in Figure 3-8. This figure also incorporates an RTS/CTS exchange, because it is common for the fragmentation and RTS/CTS thresholds to be set to the same value. The figure also shows how the NAV and SIFS are used in combination to control access to the medium.

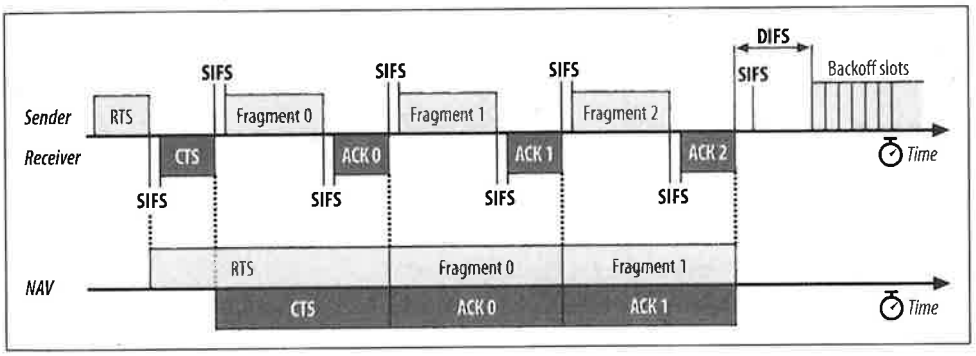

Figure 3-8. Fragmentation burst

Fragments and their acknowledgments are separated by the SIFS, so a station retains control of the channel during a fragmentation burst. The NAV is also used to ensure that other stations don't use the channel during the fragmentation burst. As with any RTS/CTS exchange, the RTS and CTS both set the NAV from the expected time to the end of the first fragments in the air. Subsequent fragments then form a chain. Each fragment sets the NAV to hold the medium until the end of the acknowledgment for the next frame. Fragment 0 sets the NAV to hold the medium until ACK 1, fragment 1 sets the NAV to hold the medium until ACK 2, and so on. After the last fragment and its acknowledgment have been sent, the NAV is set to 0, indicating that the medium will be released after the fragmentation burst completes.

Frame Format

To meet the challenges posed by a wireless data link, the MAC was forced to adopt several unique features, not the least of which was the use of four address fields. Not all frames use all the address fields, and the values assigned to the address fields may

change depending on the type of MAC frame being transmitted. Details on the use of address fields in different frame types are presented in Chapter 4.

Figure 3-9 shows the generic 802.11 MAC frame. All diagrams in this section follow the IEEE conventions in 802.11. Fields are transmitted from left to right, and the most significant bits appear last.

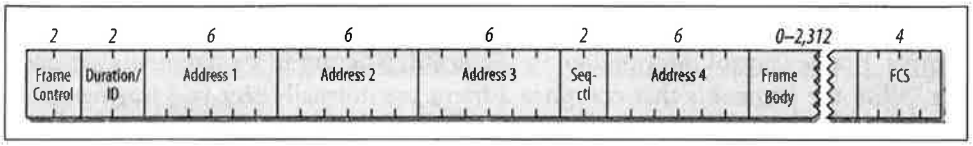

Figure 3-9. Generic 802.11 MAC frame

802.11 MAC frames do not include some of the classic Ethernet frame features, most notably the type/length field and the preamble. The preamble is part of the physical layer, and encapsulation details such as type and length are present in the header on the data carried in the 802.11 frame.

Frame Control

Each frame starts with a two-byte Frame Control subfield, shown in Figure 3-10. The components of the Frame Control subfield are:

Protocol version

Two bits indicate which version of the 802.11 MAC is contained in the rest of the frame. At present, only one version of the 802.11 MAC has been developed; it is assigned the protocol number 0. Other values will appear when the IEEE standardizes changes to the MAC that render it incompatible with the initial specification.

Type and subtype fields

Type and subtype fields identify the type of frame used. To cope with noise and unreliability, a number of management functions are incorporated into the 802.11 MAC. Some, such as the RTS/CTS operations and the acknowledgments, have already been discussed. Table 3-1 shows how the type and subtype identifiers are used to create the different classes of frames.

In Table 3-1, bit strings are written most-significant bit first, which is the reverse of the order used in Figure 3-10. Therefore, the frame type is the third bit in the frame control field followed by the second bit (b3 b2), and the subtype is the seventh bit, followed by the sixth, fifth, and fourth bits (b7 b6 b5 b4).

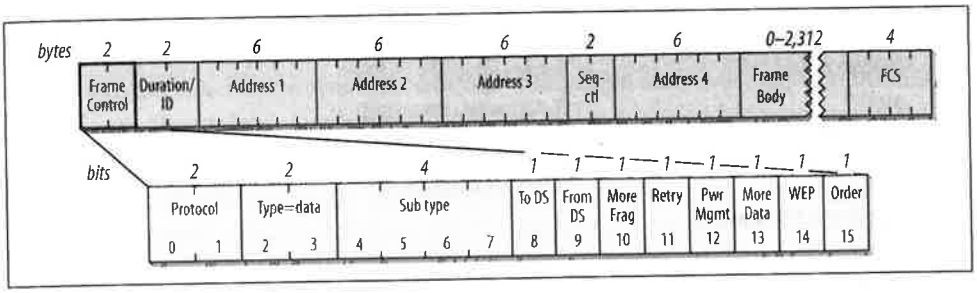

Figure 3-10. Frame control field

Table 3-1. Type and subtype identifiers

Subtype value	Subtype name
Management frames (type=00)[a]	
0000	Association request
0001	Association response
0010	Reassociation request
0011	Reassociation response
0100	Probe request
0101	Probe response
1000	Beacon
1001	Announcement traffic indication message (ATIM)
1010	Disassociation
1011	Authentication
1100	Deauthentication
Control frames (type=01)[b]	
1010	Power Save (PS)-Poll
1011	RTS
1100	CTS
1101	Acknowledgment (ACK)
1110	Contention-Free (CF)-End
1111	CF-End+CF-Ack
Data frames (type=10)[c]	
0000	Data
0001	Data+CF-Ack
0010	Data+CF-Poll
0011	Data+CF-Ack+CF-Poll
0100	Null data (no data transmitted)

Table 3-1. Type and subtype identifiers (continued)

Subtype value	Subtype name
0101	CF-Ack (no data transmitted)
0110	CF-Poll (no data transmitted)
0111	Data+CF-Ack+CF-Poll
(Frame type 11 is reserved)	

[a]Management subtypes 0110–0111 and 1101–1111 are reserved and not currently used.
[b]Control subtypes 0000–1001 are reserved and not currently used.
[c]Data subtypes 1000–1111 are reserved and not currently used.

ToDS and FromDS bits

These bits indicate whether a frame is destined for the distribution system. All frames on infrastructure networks will have one of the distribution system's bits set. Table 3-2 shows how these bits are interpreted. As Chapter 4 will explain, the interpretation of the address fields depends on the setting of these bits.

Table 3-2. Interpreting the ToDS and FromDS bits

	To DS=0	To DS=1
From DS=0	All management and control frames Data frames within an IBSS (never infrastructure data frames)	Data frames transmitted from a wireless station in an infrastructure network
From DS=1	Data frames received for a wireless station in an infrastructure network	Data frames on a "wireless bridge"

More fragments bit

This bit functions much like the "more fragments" bit in IP. When a higher-level packet has been fragmented by the MAC, the initial fragment and any following nonfinal fragments set this bit to 1. Some management frames may be large enough to require fragmentation; all other frames set this bit to 0.

Retry bit

From time to time, frames may be retransmitted. Any retransmitted frames set this bit to 1 to aid the receiving station in eliminating duplicate frames.

Power management bit

Network adapters built on 802.11 are often built to the PC Card form factor and used in battery-powered laptop or handheld computers. To conserve battery life, many small devices have the ability to power down parts of the network interface. This bit indicates whether the sender will be in a power-saving mode after the completion of the current atomic frame exchange. One indicates that the station will be in power-save mode, and 0 indicates that the station will be active. Access points perform a number of important management functions and are not allowed to save power, so this bit is always 0 in frames transmitted by an access point.

More data bit

To accommodate stations in a power-saving mode, access points may buffer frames received from the distribution system. An access point sets this bit to indicate that at least one frame is available and is addressed to a dozing station.

WEP bit

Wireless transmissions are inherently easier to intercept than transmissions on a fixed network. 802.11 defines a set of encryption routines called Wired Equivalent Privacy (WEP) to protect and authenticate data. When a frame has been processed by WEP, this bit is set to 1, and the frame changes slightly. WEP is described in more detail in Chapter 5.

Order bit

Frames and fragments can be transmitted in order at the cost of additional processing by both the sending and receiving MACs. When the "strict ordering" delivery is employed, this bit is set to 1.

Duration/ID Field

The Duration/ID field follows the frame control field. This field has several uses and takes one of the three forms shown in Figure 3-11.

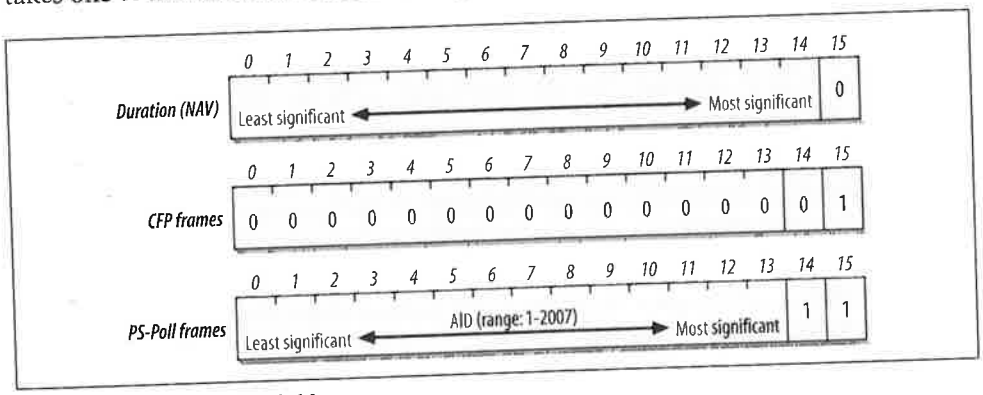

Figure 3-11. Duration/ID field

Duration: setting the NAV

When bit 15 is 0, the duration/ID field is used to set the NAV. The value represents the number of microseconds that the medium is expected to remain busy for the transmission currently in progress. All stations must monitor the headers of all frames they receive and update the NAV accordingly. Any value that extends the amount of time the medium is busy updates the NAV and blocks access to the medium for additional time.

Frames transmitted during contention-free periods

During the contention-free periods, bit 14 is 0 and bit 15 is 1. All other bits are 0, so the duration/ID field takes a value of 32,768. This value is interpreted as a NAV. It allows any stations that did not receive the Beacon* announcing the contention-free period to update the NAV with a suitably large value to avoid interfering with contention-free transmissions.

PS-Poll frames

Bits 14 and 15 are both set to 0 in PS-Poll frames. Mobile stations may elect to save battery power by turning off antennas. Dozing stations must wake up periodically. To ensure that no frames are lost, stations awaking from their slumber transmit a PS-Poll frame to retrieve any buffered frames from the access point. Along with this request, waking stations incorporate the association ID (AID) that indicates which BSS they belong to. The AID is included in the PS-Poll frame and may range from 1–2,007. Values from 2,008–16,383 are reserved and not used.

Address Fields

An 802.11 frame may contain up to four address fields. The address fields are numbered because different fields are used for different purposes depending on the frame type (details are found in Chapter 4). The general rule of thumb is that Address 1 is used for the receiver, Address 2 for the transmitter, with the Address 3 field used for filtering by the receiver.

Addressing in 802.11 follows the conventions used for the other IEEE 802 networks, including Ethernet. Addresses are 48 bits long. If the first bit sent to the physical medium is a 0, the address represents a single station (unicast). When the first bit is a 1, the address represents a group of physical stations and is called a *multicast* address. If all bits are 1s, then the frame is a *broadcast* and is delivered to all stations connected to the wireless medium.

48-bit addresses are used for a variety of purposes:

Destination address
: As in Ethernet, the destination address is the 48-bit IEEE MAC identifier that corresponds to the final recipient: the station that will hand the frame to higher protocol layers for processing.

Source address
: This is the 48-bit IEEE MAC identifier that identifies the source of the transmission. Only one station can be the source of a frame, so the Individual/Group bit is always 0 to indicate an individual station.

* Beacon frames are a subtype of management frames, which is why "Beacon" is capitalized.

Receiver address

This is a 48-bit IEEE MAC identifier that indicates which wireless station should process the frame. If it is a wireless station, the receiver address is the destination address. For frames destined to a node on an Ethernet connected to an access point, the receiver is the wireless interface in the access point, and the destination address may be a router attached to the Ethernet.

Transmitter address

This is a 48-bit IEEE MAC address to identify the wireless interface that transmitted the frame onto the wireless medium. The transmitter address is used only in wireless bridging.

Basic Service Set ID (BSSID)

To identify different wireless LANs in the same area, stations may be assigned to a BSS. In infrastructure networks, the BSSID is the MAC address used by the wireless interface in the access point. Ad hoc networks generate a random BSSID with the Universal/Local bit set to 1 to prevent conflicts with officially assigned MAC addresses.

The number of address fields used depends on the type of frame. Most data frames use three fields for source, destination, and BSSID. The number and arrangement of address fields in a data frame depends on how the frame is traveling relative to the distribution system. Most transmissions use three addresses, which is why only three of the four addresses are contiguous in the frame format.

Sequence Control Field

This 16-bit field is used for both defragmentation and discarding duplicate frames. It is composed of a 4-bit fragment number field and a 12-bit sequence number field, as shown in Figure 3-12.

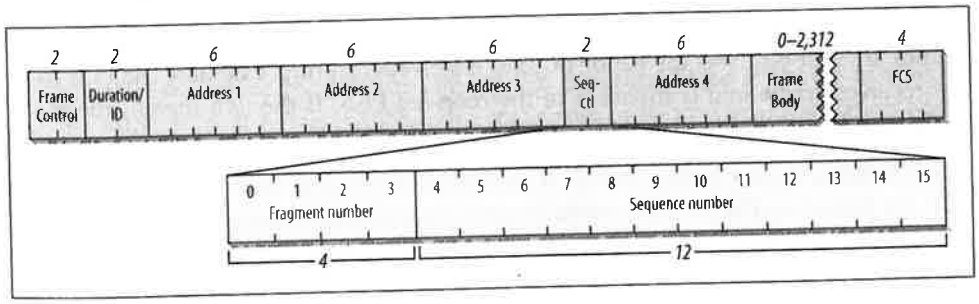

Figure 3-12. Sequence Control field

Higher-level frames are each given a sequence number as they are passed to the MAC for transmission. The sequence number subfield operates as a modulo-4096 counter

of the frames transmitted. It begins at 0 and increments by 1 for each higher-level packet handled by the MAC. If higher-level packets are fragmented, all fragments will have the same sequence number. When frames are retransmitted, the sequence number is not changed.

What differs between fragments is the fragment number. The first fragment is given a fragment number of 0. Each successive fragment increments the fragment number by one. Retransmitted fragments keep their original sequence numbers to assist in reassembly.

Frame Body

The frame body, also called the Data field, moves the higher-layer payload from station to station. 802.11 can transmit frames with a maximum payload of 2,304 bytes of higher-level data. (Implementations must support frame bodies of 2,312 bytes to accommodate WEP overhead.) 802.2 LLC headers use 8 bytes for a maximum network protocol payload of 2,296 bytes. Preventing fragmentation must be done at the protocol layer. On IP networks, Path MTU Discovery (RFC 1191) will prevent the transmission of frames with Data fields larger than 1,500 bytes.

Frame Check Sequence

As with Ethernet, the 802.11 frame closes with a frame check sequence (FCS). The FCS is often referred to as the cyclic redundancy check (CRC) because of the underlying mathematical operations. The FCS allows stations to check the integrity of received frames. All fields in the MAC header and the body of the frame are included in the FCS. Although 802.3 and 802.11 use the same method to calculate the FCS, the MAC header used in 802.11 is different from the header used in 802.3, so the FCS must be recalculated by access points.

When frames are sent to the wireless interface, the FCS is calculated before those frames are sent out over the RF or IR link. Receivers can then calculate the FCS from the received frame and compare it to the received FCS. If the two match, there is a high probability that the frame was not damaged in transit.

On Ethernets, frames with a bad FCS are simply discarded, and frames with a good FCS are passed up the protocol stack. On 802.11 networks, frames that pass the integrity check may also require the receiver to send an acknowledgment. For example, data frames that are received correctly must be positively acknowledged, or they are retransmitted. 802.11 does not have a negative acknowledgment for frames that fail the FCS; stations must wait for the acknowledgment timeout before retransmitting.

Encapsulation of Higher-Layer Protocols Within 802.11

Like all other 802 link layers, 802.11 can transport any network-layer protocol. Unlike Ethernet, 802.11 relies on 802.2 logical-link control (LLC) encapsulation to carry higher-level protocols. Figure 3-13 shows how 802.2 LLC encapsulation is used to carry an IP packet. In the figure, the "MAC headers" for 802.1h and RFC 1042 might be the 12 bytes of source and destination MAC address information on Ethernet or the long 802.11 MAC header from the previous section.

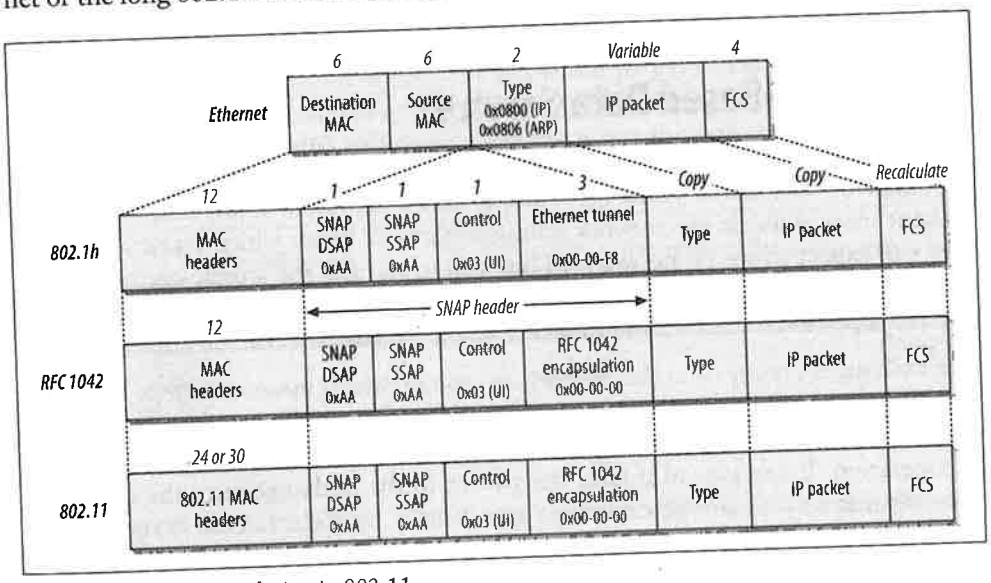

Figure 3-13. IP encapsulation in 802.11

Two different methods can be used to encapsulate LLC data for transmission. One is described in RFC 1042, and the other in 802.1h. As you can see in Figure 3-13, though, the two methods are quite similar. An Ethernet frame is shown in the top line of Figure 3-13. It has a MAC header composed of source and destination MAC addresses, a type code, the embedded packet, and a frame check field. In the IP world, the Type code is either 0x0800 (2048 decimal) for IP itself, or 0x0806 (2054 decimal) for the Address Resolution Protocol (ARP).

Both RFC 1042 and 802.1h are derivatives of 802.2's *sub-network access protocol* (SNAP). The MAC addresses are copied into the beginning of the encapsulation frame, and then a SNAP header is inserted. SNAP headers begin with a *destination service access point* (DSAP) and a *source service access point* (SSAP). After the addresses, SNAP includes a Control header. Like high-level data link control (HDLC)

and its progeny, the Control field is set to 0x03 to denote unnumbered information (UI), a category that maps well to the best-effort delivery of IP datagrams. The last field inserted by SNAP is an organizationally unique identifier (OUI). Initially, the IEEE hoped that the 1-byte service access points would be adequate to handle the number of network protocols, but this proved to be an overly optimistic assessment of the state of the world. As a result, SNAP copies the type code from the original Ethernet frame.

Products usually have a software option to toggle between the two encapsulation types. Of course, products on the same network must use the same type of encapsulation.

Contention-Based Data Service

The additional features incorporated into 802.11 to add reliability lead to a confusing tangle of rules about which types of frames are permitted at any point. They also make it more difficult for network administrators to know which frame exchanges they can expect to see on networks. This section clarifies the atomic exchanges that move data on an 802.11 LAN. (Most management frames are announcements to interested parties in the area and transfer information in only one direction.)

The exchanges presented in this section are atomic, which means that they should be viewed as a single unit. As an example, unicast data is always acknowledged to ensure delivery. Although the exchange spans two frames, the exchange itself is a single operation. If any part of it fails, the parties to the exchange retry the operation. Two distinct sets of atomic exchanges are defined by 802.11. One is used by the DCF for contention-based service; those exchanges are described in this chapter. A second set of exchanges is specified for use with the PCF for contention-free services. Frame exchanges used with contention-free services are intricate and harder to understand. Since very few (if any) commercial products implement contention-free service, these exchanges are not described.

Frame exchanges under the DCF dominate the 802.11 MAC. According to the rules of the DCF, all products are required to provide best-effort delivery. To implement the contention-based MAC, stations process MAC headers for every frame while they are active. Exchanges begin with a station seizing an idle medium after the DIFS.

Broadcast and Multicast Data or Management Frames

Broadcast and multicast frames have the simplest frame exchanges because there is no acknowledgment. Framing and addressing are somewhat more complex in 802.11, so the types of frames that match this rule are the following:

- Broadcast data frames with a broadcast address in the Address1 field
- Multicast data frames with a multicast address in the Address1 field
- Broadcast management frames with a broadcast address in the Address1 field (Beacon, Probe Request, and IBSS ATIM frames)

Frames destined for group addresses cannot be fragmented and are not acknowledged. The entire atomic sequence is a single frame, sent according to the rules of the contention-based access control. After the previous transmission concludes, all stations wait for the DIFS and begin counting down the random delay intervals in the contention window.

Because the frame exchange is a single-frame sequence, the NAV is set to 0. With no further frames to follow, there is no need to use the virtual carrier-sense mechanism to lock other stations out of using the medium. After the frame is transmitted, all stations wait through the DIFS and begin counting down through the contention window for any deferred frames. See Figure 3-14.

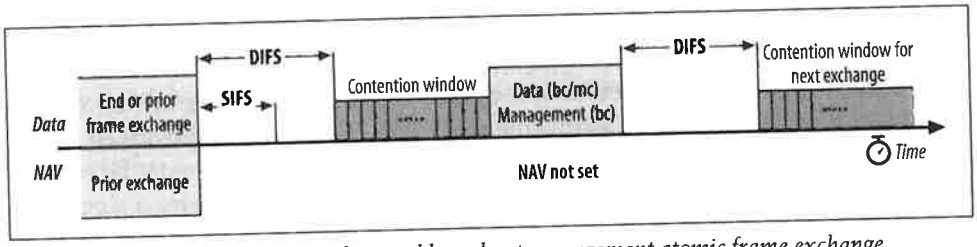

Figure 3-14. Broadcast/multicast data and broadcast management atomic frame exchange

Depending on the environment, frames sent to group addresses may have lower service quality because the frames are not acknowledged. Some stations may therefore miss broadcast or multicast traffic, but there is no facility built into the MAC for retransmitting broadcast or multicast frames.

Unicast Frames

Frames that are destined for a single station are called *directed* data by the 802.11 standard. This book uses the more common term *unicast*. Unicast frames must be acknowledged to ensure reliability, which means that a variety of mechanisms can be used to improve efficiency. All the sequences in this section apply to any unicast frame and thus can apply to management frames and data frames. In practice, these operations are usually observed only with data frames.

Basic positive acknowledgment (final fragment)

Reliable transmission between two stations is based on simple positive acknowledgments. Unicast data frames must be acknowledged, or the frame is assumed to be lost. The most basic case is a single frame and its accompanying acknowledgment, as shown in Figure 3-15.

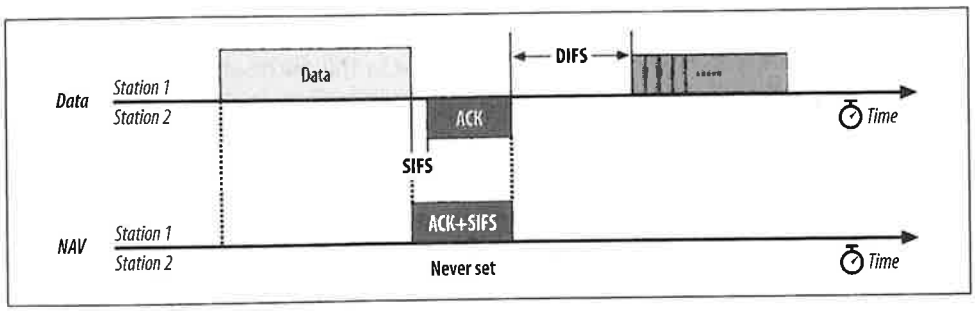

Figure 3-15. Basic positive acknowledgment of data

The frame uses the NAV to reserve the medium for the frame, its acknowledgment, and the intervening SIFS. By setting a long NAV, the sender locks the virtual carrier for the entire sequence, guaranteeing that the recipient of the frame can send the acknowledgment. Because the sequence concludes with the ACK, no further virtual carrier locking is necessary, and the NAV in the ACK is set to 0.

Fragmentation

Many higher-layer network protocols, including IP, incorporate fragmentation. The disadvantage of network-layer fragmentation is that reassembly is performed by the final destination; if any of the fragments are lost, the entire packet must be retransmitted. Link layers may incorporate fragmentation to boost speed over a single hop with a small MTU.* 802.11 can also use fragmentation to help avoid interference. Radio interference is often in the form of short, high-energy bursts and is frequently synchronized with the AC power line. Breaking a large frame into small frames allows a larger percentage of the frames to arrive undamaged. The basic fragmentation scheme is shown in Figure 3-16.

The last two frames exchanged are the same as in the previous sequence, and the NAV is set identically. However, all previous frames use the NAV to lock the medium for the next frame. The first data frame sets the NAV for a long enough period to accommodate its ACK, the next fragment, and the acknowledgment following the next fragment. To indicate that it is a fragment, the MAC sets the More Fragments bit in the frame control field to 1. All nonfinal ACKs continue to extend the lock for the next data fragment and its ACK. Subsequent data frames then continue to lengthen the NAV to include successive acknowledgments until the final data frame, which sets the More Fragments bit to 0, and the final ACK, which sets the NAV to 0. No limit is placed on the number of fragments, but the total frame length must be shorter than any constraint placed on the exchange by the PHY.

Fragmentation is controlled by the fragmentation threshold parameter in the MAC. Most network card drivers allow you to configure this parameter. Any frames larger

* This is the approach used by Multi-link PPP (RFC 1990).

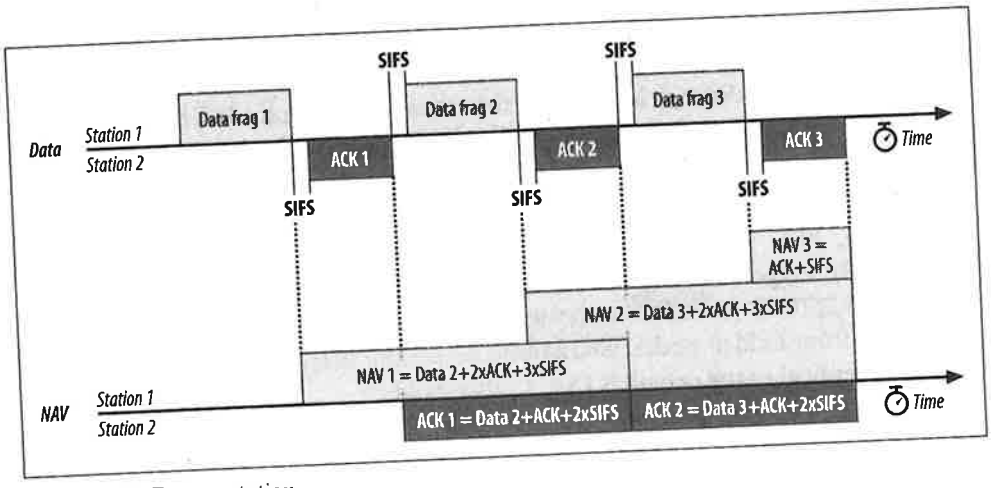

Figure 3-16. Fragmentation

than the fragmentation threshold are fragmented in an implementation-dependent way. Network administrators can change the fragmentation threshold to tune network behavior. Higher fragmentation thresholds mean that frames are delivered with less overhead, but the cost to a lost or damaged frame is much higher because more data must be discarded and retransmitted. Low fragmentation thresholds have much higher overhead, but they offer increased robustness in the face of hostile conditions.

RTS/CTS

To guarantee reservation of the medium and uninterrupted data transmission, a station can use the RTS/CTS exchange. Figure 3-17 shows this process. The RTS/CTS exchange acts exactly like the initial exchange in the fragmentation case, except that the RTS frame does not carry data. The NAV in the RTS allows the CTS to complete, and the CTS is used to reserve access for the data frame.

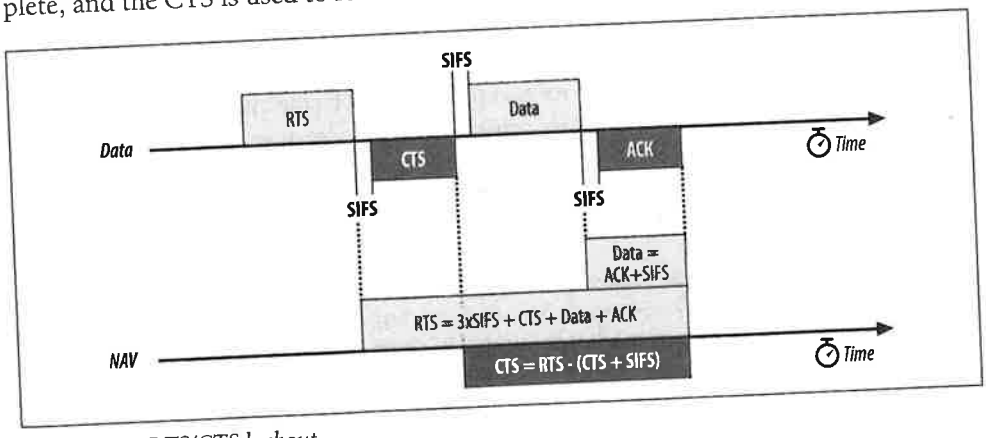

Figure 3-17. RTS/CTS lockout

RTS/CTS can be used for all frame exchanges, none of them, or something in between. Like fragmentation, RTS/CTS behavior is controlled by a threshold set in the driver software. Frames larger than the threshold are preceded by an RTS/CTS exchange to clear the medium, while smaller frames are simply transmitted.

RTS/CTS with fragmentation

In practice, the RTS/CTS exchange is often combined with fragmentation (Figure 3-18). Fragmented frames are usually quite long and thus benefit from the use of the RTS/CTS procedure to ensure exclusive access to the medium, free from contention from hidden nodes. Some vendors set the default fragmentation threshold to be identical to the default RTS/CTS threshold.

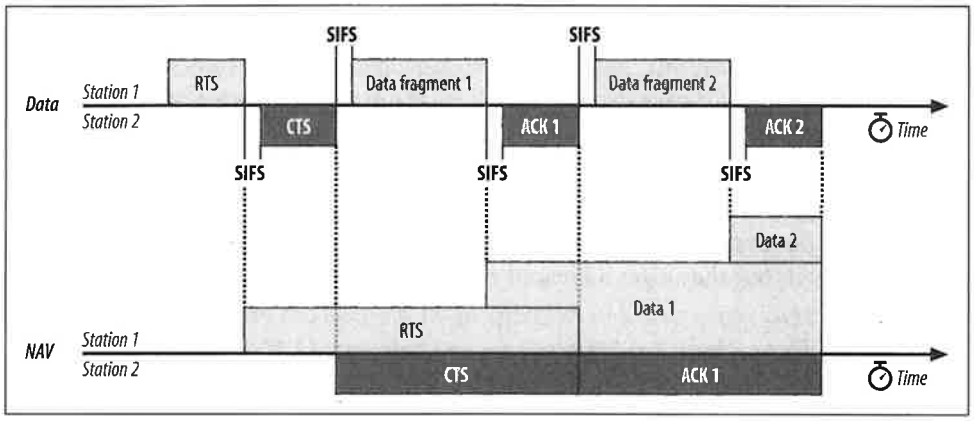

Figure 3-18. RTS/CTS with fragmentation

Power-Saving Sequences

The most power-hungry components in RF systems are the amplifiers used to boost a signal immediately prior to transmission and to boost the received signal to an intelligible level immediately after its reception. 802.11 stations can maximize battery life by shutting down the radio transceiver and sleeping periodically. During sleeping periods, access points buffer any unicast frames for sleeping stations. These frames are announced by subsequent Beacon frames. To retrieve buffered frames, newly awakened stations use PS-Poll frames.

Immediate response

Access points can respond immediately to the PS-Poll. After a short interframe space, an access point may transmit the frame. Figure 3-19 shows an implied NAV as a result of the PS-Poll frame. The PS-Poll frame contains an Association ID in the Duration/ID field so that the access point can determine which frames were buffered for the mobile station. However, the MAC specification requires all stations receiving a

PS-Poll to update the NAV with an implied value equal to a short interframe space and one ACK. Although the NAV is too short for the data frame, the access point acquires that the medium and all stations defer access for the entire data frame. At the conclusion of the data frame, the NAV is updated to reflect the value in the header of the data frame.

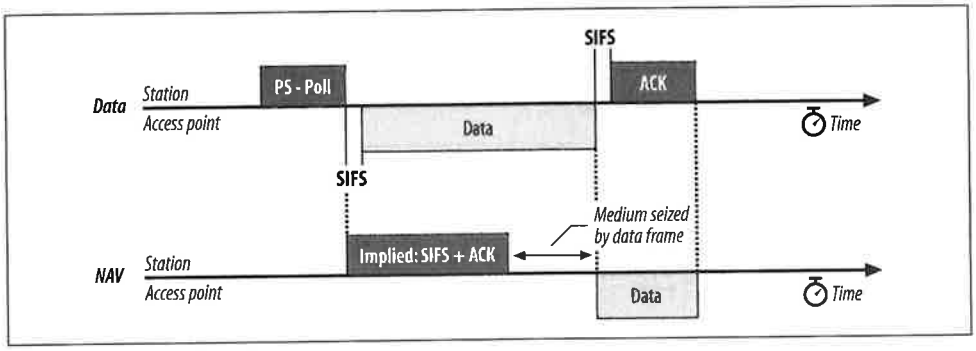

Figure 3-19. Immediate PS-Poll response

If the buffered frame is large, it may require fragmentation. Figure 3-20 illustrates an immediate PS-Poll response requiring fragmentation. Like all other stations, access points typically have a configurable fragmentation threshold.

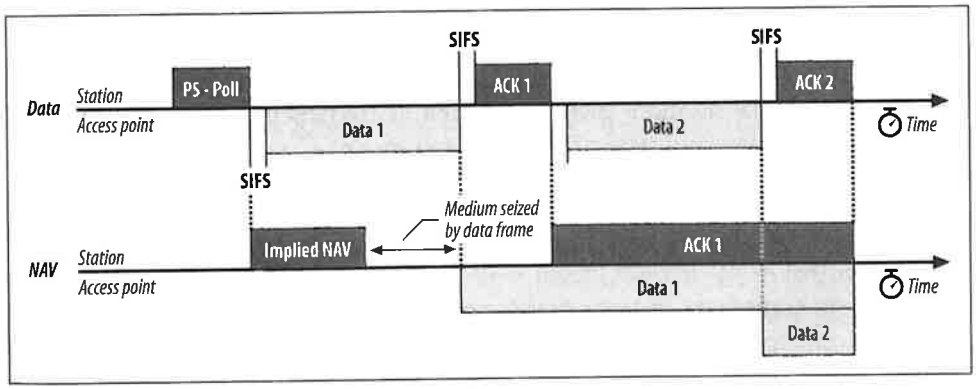

Figure 3-20. Immediate PS-Poll response with fragmentation

Deferred response

Instead of an immediate response, access points can also respond with a simple acknowledgment. This is called a *deferred response* because the access point acknowledges the request for the buffered frame but does not act on it immediately. A station requesting a frame with a PS-Poll must stay awake until it is delivered. Under contention-based service, however, the access point can deliver a frame at any point. A station cannot return to a low-power mode until it receives a Beacon frame in which its bit in the traffic indication map (TIM) is clear.

Figure 3-21 illustrates this process. In this figure, the station has recently changed from a low-power mode to an active mode, and it notes that the access point has buffered frames for it. It transmits a PS-Poll to the access point to retrieve the buffered frames. However, the access point may choose to defer its response by transmitting only an ACK. At this point, the access point has acknowledged the station's request for buffered frames and promised to deliver them at some point in the future. The station must wait in active mode, perhaps through several atomic frame exchanges, before the access point delivers the data. A buffered frame may be subject to fragmentation, although Figure 3-21 does not illustrate this case.

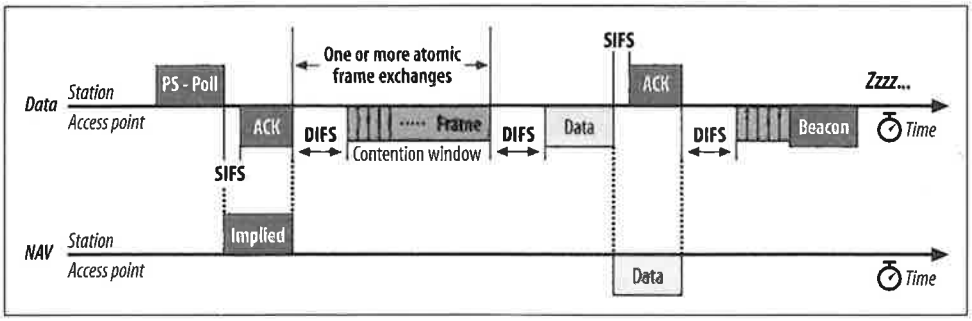

Figure 3-21. Deferred PS-Poll response example

After receiving a data frame, the station must remain awake until the next Beacon is transmitted. Beacon frames only note whether frames are buffered for a station and have no way of indicating the number of frames. Once the station receives a Beacon frame indicating that no more traffic is buffered, it can conclude that it has received the last buffered frame and return to a low-power mode.

CHAPTER 4

802.11 Framing in Detail

Chapter 3 presented the basic frame structure and the fields that comprise it, but it did not go into detail about the different frame types. Ethernet framing is a simple matter: add a preamble, some addressing information, and tack on a frame check at the end. 802.11 framing is much more involved because the wireless medium requires several management features and corresponding frame types not found in wired networks.

Three major frame types exist. Data frames are the pack horses of 802.11, hauling data from station to station. Several different data frame flavors can occur, depending on the network. Control frames are used in conjunction with data frames to perform area clearing operations, channel acquisition and carrier-sensing maintenance functions, and positive acknowledgment of received data. Control and data frames work in conjunction to deliver data reliably from station to station. Management frames perform supervisory functions; they are used to join and leave wireless networks and move associations from access point to access point.

This chapter is intended to be a reference. There is only so much life any author can breathe into framing details, no matter how much effort is expended to make the details interesting. Please feel free to skip this chapter in its entirety and flip back when you need in-depth information about frame structure. With rare exception, detailed framing relationships generally do not fall into the category of "something a network administrator needs to know." This chapter tends to be a bit acronym-heavy as well, so refer to the glossary at the back of the book if you do not recognize an acronym.

Data Frames

Data frames carry higher-level protocol data in the frame body. Figure 4-1 shows a generic data frame. Depending on the particular type of data frame, some of the fields in the figure may not be used.

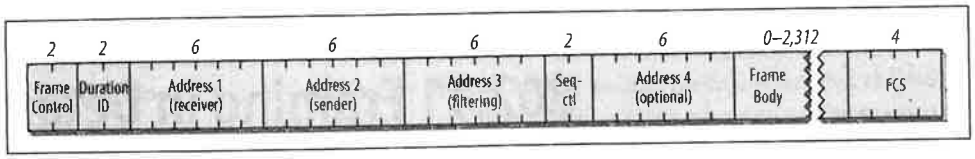

Figure 4-1. Generic data frame

The different data frame types can be categorized according to function. One such distinction is between data frames used for contention-based service and those used for contention-free service. Any frames that appear only in the contention-free period can never be used in an IBSS. Another possible division is between frames that carry data and frames that perform management functions. Table 4-1 shows how frames may be divided along these lines. Frames used in contention-free service are discussed in detail in Chapter 8.

Table 4-1. Categorization of data frame types

Frame type	Contention-based service	Contention-free service	Carries data	Does not carry data
Data	✓		✓	
Data+CF-Ack		✓	✓	
Data+CF-Poll		AP only	✓	
Data+CF-Ack+CF-Poll		AP only	✓	
Null	✓	✓		✓
CF-Ack		✓		✓
CF-Poll		AP only		✓
CF-Ack+CF-Poll		AP only		✓

Frame Control

All the bits in the Frame Control field are used according to the rules described in Chapter 3. Frame Control bits may affect the interpretation of other fields in the MAC header, though. Most notable are the address fields, which depend on the value of the ToDS and FromDS bits.

Duration

The Duration field carries the value of the Network Allocation Vector (NAV). Access to the medium is restricted for the time specified by the NAV. Four rules specify the setting for the Duration field in data frames:

1. Any frames transmitted during the contention-free period set the Duration field to 32,768. Naturally, this applies to any data frames transmitted during this period.

2. Frames transmitted to a broadcast or multicast destination (Address 1 has the group bit set) have a duration of 0. Such frames are not part of an atomic exchange and are not acknowledged by receivers, so contention-based access to the medium can begin after the conclusion of a broadcast or multicast data frame. The NAV is used to protect access to the transmission medium for a frame exchange sequence. With no link-layer acknowledgment following the transmission of a broadcast or multicast frame, there is no need to lock access to the medium for subsequent frames.
3. If the More Fragments bit in the Frame Control field is 0, no more fragments remain in the frame. The final fragment need only reserve the medium for its own ACK, at which point contention-based access resumes. The Duration field is set to the amount of time required for one short interframe space and the fragment acknowledgment. Figure 4-2 illustrates this process. The penultimate fragment's Duration field locks access to the medium for the transmission of the last fragment.

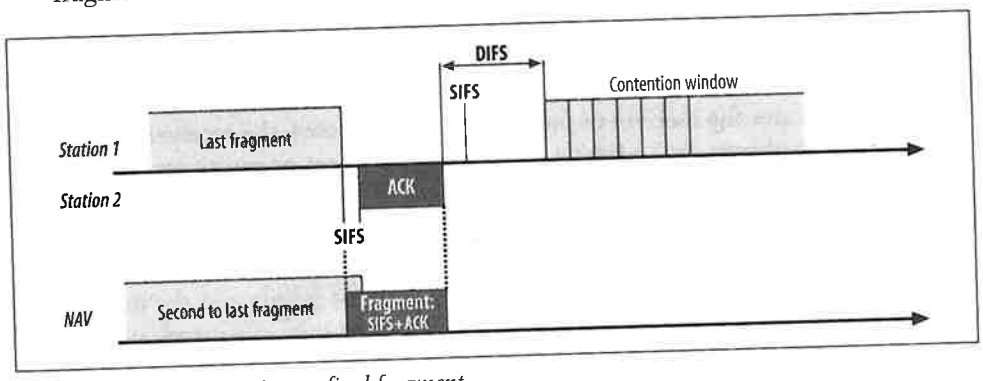

Figure 4-2. Duration setting on final fragment

4. If the More Fragments bit in the Frame Control field is set to 1, more fragments remain. The Duration field is set to the amount of time required for transmission of two acknowledgments, plus three short interframe spaces, plus the time required for the next fragment. In essence, nonfinal fragments set the NAV just like an RTS would (Figure 4-3); for this reason, they are referred to as a *virtual RTS*.

Addressing and DS Bits

The number and function of the address fields depends on which of the distribution system bits are set, so the use of the address fields indirectly depends on the type of network deployed. Table 4-2 summarizes the use of the address fields in data frames.

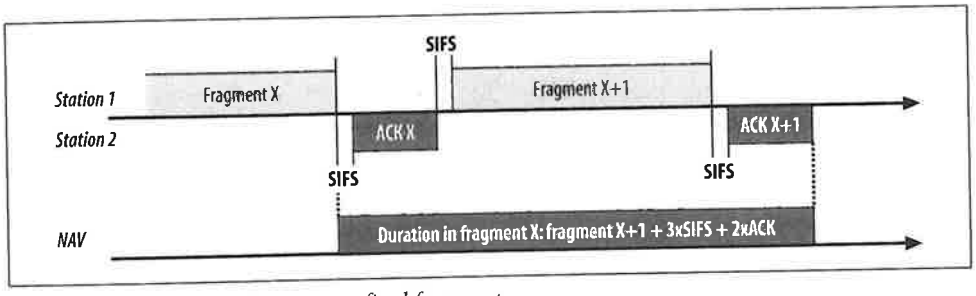

Figure 4-3. Duration settings on nonfinal fragment

Table 4-2. Use of the address fields in data frames

Function	ToDS	FromDS	Address 1 (receiver)	Address 2 (transmitter)	Address 3	Address 4
IBSS	0	0	DA	SA	BSSID	not used
To AP (infra.)	1	0	BSSID	SA	DA	not used
From AP (infra.)	0	1	DA	BSSID	SA	not used
WDS (bridge)	1	1	RA	TA	DA	SA

Address 1 indicates the receiver of the frame. In many cases, the receiver is the destination, but not always. If Address 1 is set to a broadcast or multicast address, the BSSID is also checked. Stations respond only to broadcasts and multicasts originating in the same basic service set (BSS); they ignore broadcasts and multicasts from different BSSs. Address 2 is the transmitter address and is used to send acknowledgments. The Address 3 field is used for filtering by access points and the distribution system, but the use of the field depends on the particular type of network used.

In the case of an IBSS, no access points are used, and no distribution system is present. The transmitter is the source, and the receiver is the destination. All frames carry the BSSID so that stations may check broadcasts and multicasts; only stations that belong to the same BSS will process broadcasts and multicasts. In an IBSS, the BSSID is created by a random-number generator.

802.11 draws a distinction between the source and transmitter and a parallel distinction between the destination and the receiver. The transmitter sends a frame on to the wireless medium but does not necessarily create the frame. A similar distinction holds for destination addresses and receiver addresses. A receiver may be an intermediate destination, but frames are processed by higher protocol levels only when they reach the destination.

To expand on these distinctions, consider the use of the address fields in infrastructure networks. Figure 4-4 shows a simple network in which a wireless client is connected to a server through an 802.11 network. Frames sent by the client to the server use the address fields as specified in the second line of Table 4-2.

The BSSID

Each BSS is assigned a BSSID, a 48-bit binary identifier that distinguishes it from other BSSs throughout the network. The major advantage of the BSSID is filtering. Several distinct 802.11 networks may overlap physically, and there is no reason for one network to receive link-layer broadcasts from a physically overlapping network.

In an infrastructure BSS, the BSSID is the MAC address of the wireless interface in the access point creating the BSS. IBSSs must create BSSIDs for networks brought into existence. To maximize the probability of creating a unique address, 46 random bits are generated for the BSSID. The Universal/Local bit for the new BSSID is set to 1, indicating a local address, and the Individual/Group bit is set to 0. For two distinct IBSSs to create the same BSSID, they would need to generate an identical random 46 bits.

One BSSID is reserved. The all-1s BSSID is the *broadcast BSSID*. Frames that use the broadcast BSSID pass through any BSSID filtering in the MAC. BSSID broadcasts are used only when mobile stations try to locate a network by sending probe requests. In order for probe frames to detect the existence of a network, they must not be filtered by the BSSID filter. Probe frames are the only frames allowed to use the broadcast BSSID.

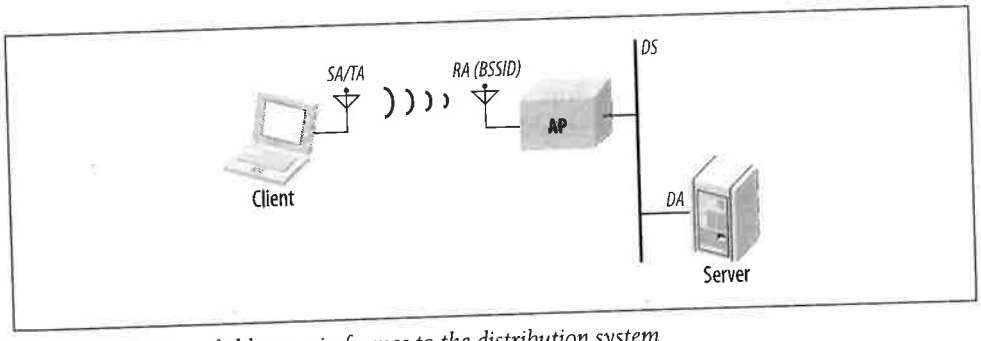

Figure 4-4. Address field usage in frames to the distribution system

In the case of frames bound for a destination on the distribution system, the client is both source and transmitter. The receiver of the wireless frame is the access point, but the access point is only an intermediate destination. When the frame reaches the access point, it is relayed to the distribution system to reach the server. Thus, the access point is the receiver, and the (ultimate) destination is the server. In infrastructure networks, access points create associated BSSs with the address of their wireless interfaces, which is why the receiver address (Address 1) is set to the BSSID.

When the server replies to the client, frames are transmitted to the client through the access point, as in Figure 4-5. This scenario corresponds to the third line in Table 4-2.

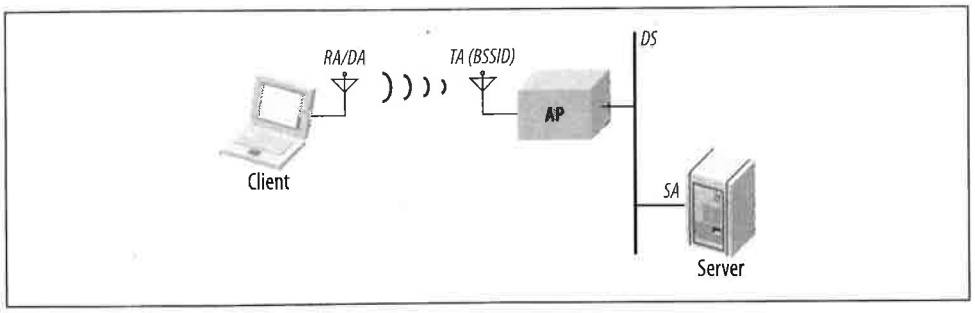

Figure 4-5. Address field usage in frames from the distribution system

Frames are created by the server, so the server's MAC address is the source address for frames. When frames are relayed through the access point, the access point uses its wireless interface as the transmitter address. As in the previous case, the access point's interface address is also the BSSID. Frames are ultimately sent to the client, which is both the destination and receiver.

The fourth line in Table 4-2 shows the use of the address fields in a *wireless distribution system (WDS)*, which is sometimes called a *wireless bridge*. In Figure 4-6, two wired networks are joined by access points acting as wireless bridges. Frames bound from the client to the server traverse the 802.11 WDS. The source and destination addresses of the wireless frames remain the client and server addresses. These frames, however, also identify the transmitter and receiver of the frame on the wireless medium. For frames bound from the client to the server, the transmitter is the client-side access point, and the receiver is the server-side access point. Separating the source from the transmitter allows the server-side access point to send required 802.11 acknowledgments to its peer access point without interfering with the wired link layer.

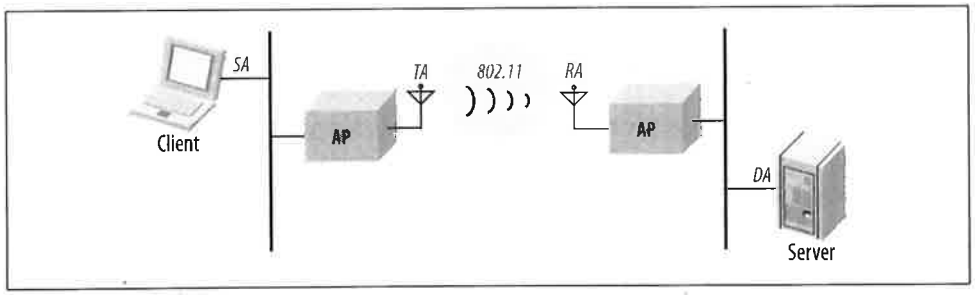

Figure 4-6. Wireless distribution system

Variations on the Data Frame Theme

802.11 uses several different data frame types. Variations depend on whether the service is contention-based or contention-free. Contention-free frames can incorporate

several functions for the sake of efficiency. Data may be transmitted, but by changing the frame subtype, data frames in the contention-free period may be used to acknowledge other frames, saving the overhead of interframe spaces and separate acknowledgments. Here are the different data frame types that are commonly used:

Data

Frames of the Data subtype are transmitted only during the contention-based access periods. They are simple frames with the sole purpose of moving the frame body from one station to another.

Null

Null frames* are a bit of an oddity. They consist of a MAC header followed by the FCS trailer. In a traditional Ethernet, empty frames would be extraneous overhead; in 802.11 networks, they are used by mobile stations to inform the access point of changes in power-saving status. When stations sleep, the access point must begin buffering frames for the sleeping station. If the mobile station has no data to send through the distribution system, it can use a Null frame with the Power Management bit in the Frame Control field set. Access points never enter power-saving mode and do not transmit Null frames. Usage of Null frames is shown in Figure 4-7.

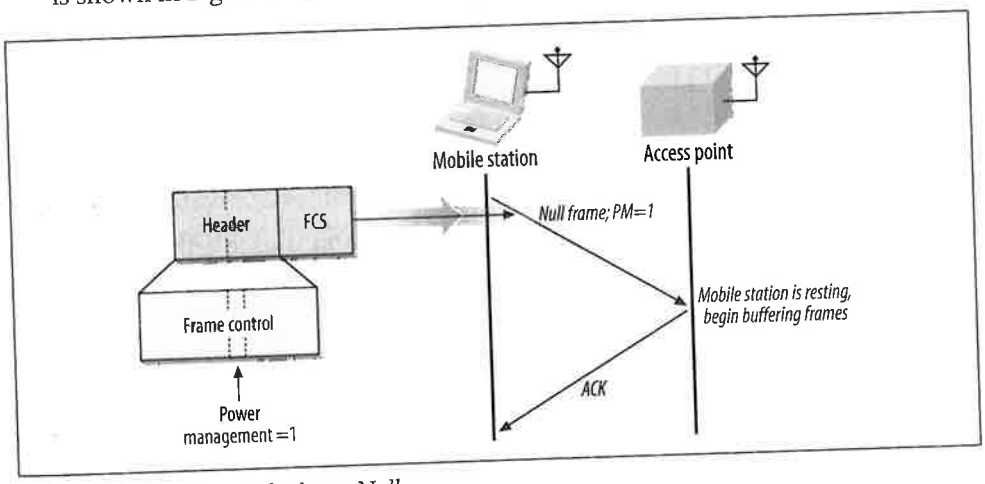

Figure 4-7. Data frame of subtype Null

Several other frame types exist for use within the contention-free period. However, contention-free service is not widely implemented, so the discussion of the contention-free frames (Data+CF-Ack, Data+CF-Poll, Data+CF-Ack+CF-Poll, CF-Ack, CF-Poll, and CF-Ack+CF-Poll) can be found in Chapter 8.

* To indicate that Null is used as the frame type from the specification rather than the English word, it is capitalized. This convention will be followed throughout the chapter.

Applied Data Framing

The form of a data frame can depend on the type of network. The actual subtype of the frame is determined solely by the subtype field, not by the presence or absence of other fields in the frame.

IBSS frames

In an IBSS, three address fields are used, as shown in Figure 4-8. The first address identifies the receiver, which is also the destination address in an IBSS. The second address is the source address. After the source and destination addresses, data frames in an IBSS are labeled with the BSSID. When the wireless MAC receives a frame, it checks the BSSID and passes only frames in the station's current BSSID to higher protocol layers.

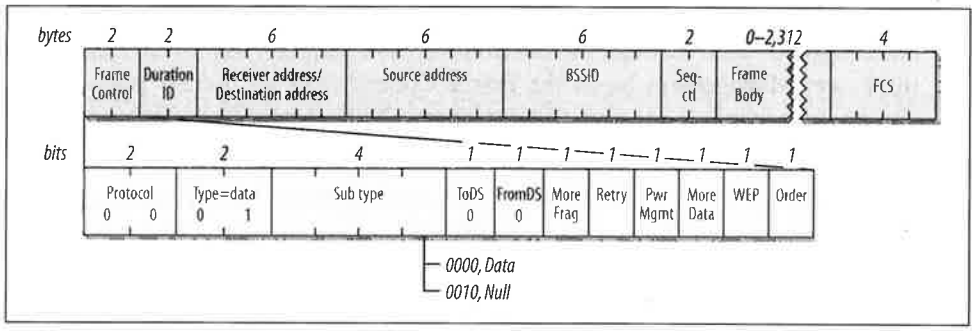

Figure 4-8. IBSS data frame

IBSS data frames have the subtype data or Null; the latter is used only to communicate power management state.

Frames from the AP

Figure 4-9 shows the format of a frame sent from an access point to a mobile station. As in all data frames, the first address field indicates the receiver of the frame on the wireless network, which is the frame's destination. The second address holds the transmitter address. On infrastructure networks, the transmitter address is the address of the station in the access point, which is also the BSSID. Finally, the frame indicates the source MAC address of the frame. The split between source and transmitter is necessary because the 802.11 MAC sends acknowledgments to the frame's transmitter (the access point), but higher layers send replies to the frame's source.

Nothing in the 802.11 specification forbids an access point from transmitting Null frames, but there is no reason to transmit them. Access points are forbidden from using the power-saving routines, and they can acknowledge Null frames from stations without using Null frames in response. In practice, access points send Data

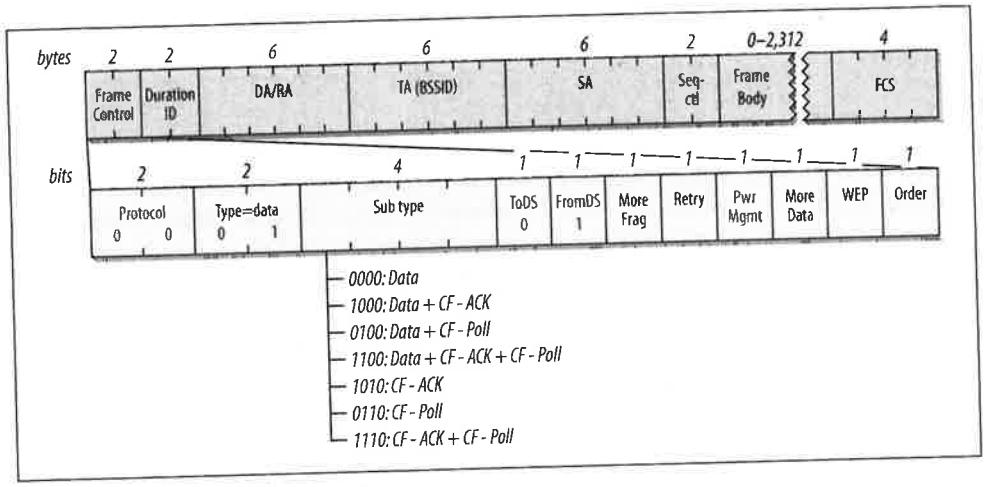

Figure 4-9. Data frames from the AP

frames during the contention-based access period, and they send frames incorporating the CF-Poll feature during the contention-free period.

Frames to the AP

Figure 4-10 shows the format of a frame sent from a mobile station in an infrastructure network to the access point currently serving it. The receiver address is the BSSID. In infrastructure networks, the BSSID is taken from the MAC address of the network station in the access point. Frames destined for an access point take their source/transmitter address from the network interface in the wireless station. Access points do not perform filtering, but instead use the third address to forward data to the appropriate location in the distribution system.

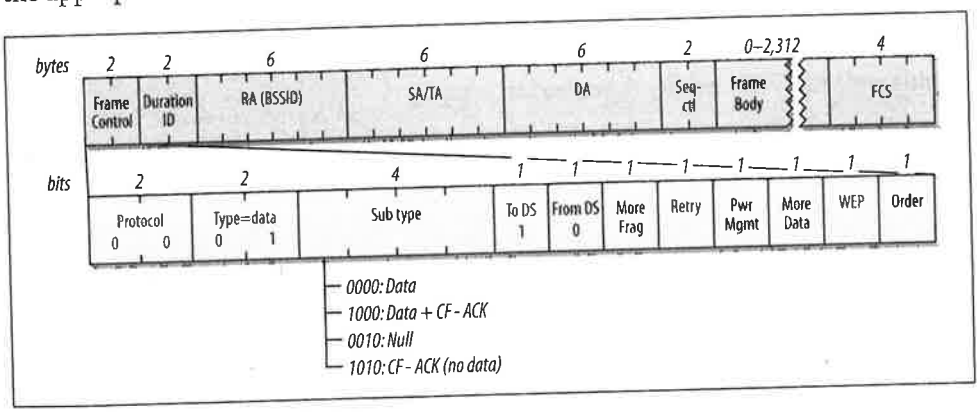

Figure 4-10. Data frames to the AP

Frames from the distribution system have the ToDS bit set, but the FromDS bit is 0. Mobile stations in an infrastructure network cannot become the point coordinator,

and thus never send frames that incorporate the contention-free polling (CF-Poll) functions.

Frames in a WDS

When access points are deployed in a wireless bridge, or WDS, topology, all four address fields are used, as shown in Figure 4-11. Like all other data frames, WDS frames use the first address for the receiver of the frame and the second address for the transmitter. The MAC uses these two addresses for acknowledgments and control traffic, such as RTS, CTS, and ACK frames. Two more address fields are necessary to indicate the source and destination of the frame and distinguish them from the addresses used on the wireless link.

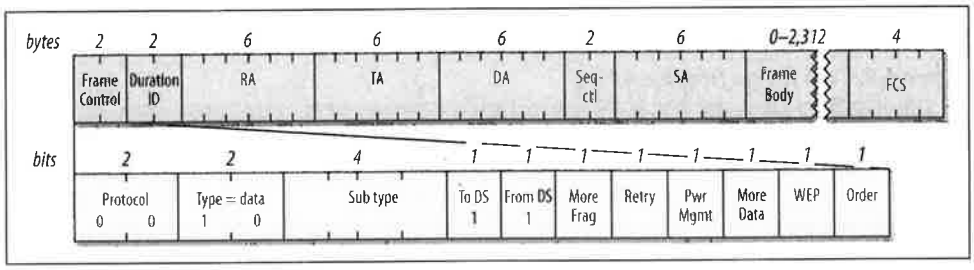

Figure 4-11. WDS frames

On a wireless bridging link, there are usually no mobile stations, and the contention-free period is not used. Access points are forbidden to enter power-saving modes, so the power management bit is always set to 0.

Frames using WEP

Frames protected by WEP are not new frame types. When a frame is handled by WEP, the WEP bit in the Frame Control field is set to 1, and the Frame Body field begins with the WEP header described in Chapter 5.

Control Frames

Control frames assist in the delivery of data frames. They administer access to the wireless medium (but not the medium itself) and provide MAC-layer reliability functions.

Common Frame Control Field

All control frames use the same Frame Control field, which is shown in Figure 4-12.

Protocol version

The protocol version is shown as 0 in Figure 4-12 because that is currently the only version. Other versions may exist in the future.

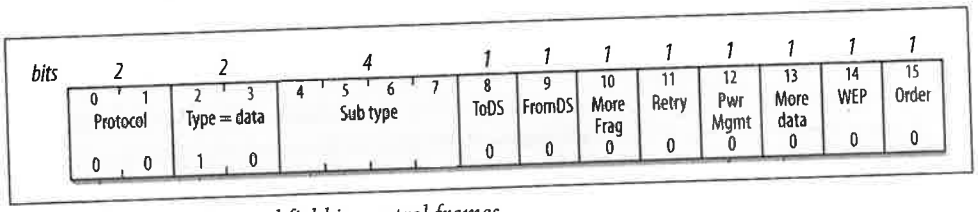

Figure 4-12. Frame Control field in control frames

Type
: Control frames are assigned the Type identifier 01. By definition, all control frames use this identifier.

Subtype
: This field indicates the subtype of the control frame that is being transmitted.

ToDS and FromDS bits
: Control frames arbitrate access to the wireless medium and thus can only originate from wireless stations. The distribution system does not send or receive control frames, so these bits are always 0.

More Fragments bit
: Control frames are not fragmented, so this bit is always 0.

Retry bit
: Control frames are not queued for retransmission like management or data frames, so this bit is always 0.

Power Management bit
: This bit is set to indicate the power management state of the sender after conclusion of the current frame exchange.

More Data bit
: The More Data bit is used only in management and data frames, so this bit is set to 0 in control frames.

WEP bit
: Control frames may not be encrypted by WEP, which may be used only for data frames and association requests. Thus, for control frames, the WEP bit is always 0.

Order bit
: Control frames are used as components of atomic frame exchange operations and thus cannot be transmitted out of order. Therefore, this bit is set to 0.

Request to Send (RTS)

RTS frames are used to gain control of the medium for the transmission of "large" frames, in which "large" is defined by the RTS threshold in the network card driver. Access to the medium can be reserved only for unicast frames; broadcast and multicast frames are simply transmitted. The format of the RTS frame is shown in

Figure 4-13. Like all control frames, the RTS frame is all header. No data is transmitted in the body, and the FCS immediately follows the header.

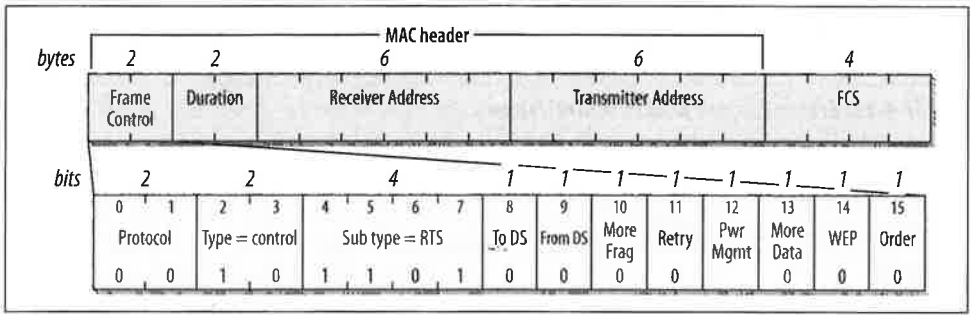

Figure 4-13. RTS frame

Four fields make up the MAC header of an RTS:

Frame Control

There is nothing special about the Frame Control field. The frame subtype is set to 1011 to indicate an RTS frame, but otherwise, it has all the same fields as other control frames. (The most significant bits in the 802.11 specification come at the end of fields, so bit 7 is the most significant bit in the subtype field.)

Duration

An RTS frame attempts to reserve the medium for an entire frame exchange, so the sender of an RTS frame calculates the time needed for the frame exchange sequence after the RTS frame ends. The entire exchange, which is depicted in Figure 4-14, requires three SIFS periods, the duration of one CTS, the final ACK, plus the time needed to transmit the frame or first fragment. (Fragmentation bursts use subsequent fragments to update the Duration field.) The number of microseconds required for the transmission is calculated and placed in the Duration field. If the result is fractional, it is rounded up to the next microsecond.

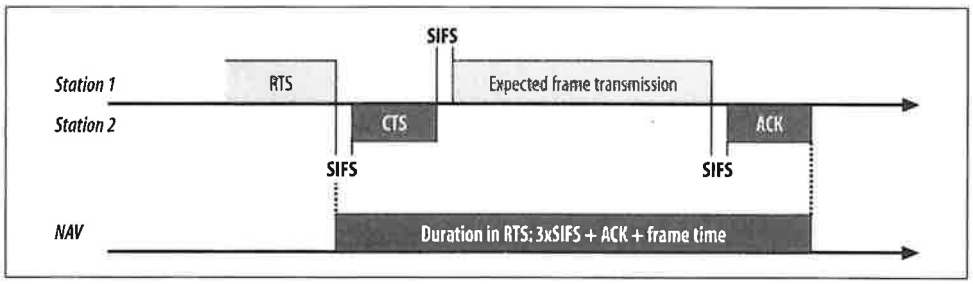

Figure 4-14. Duration field in RTS frame

Address 1: Receiver Address
The address of the station that is the intended recipient of the large frame.

Address 2: Transmitter Address
The address of the sender of the RTS frame.

Clear to Send (CTS)

The CTS frame answers the RTS frame. Its format is shown in Figure 4-15.

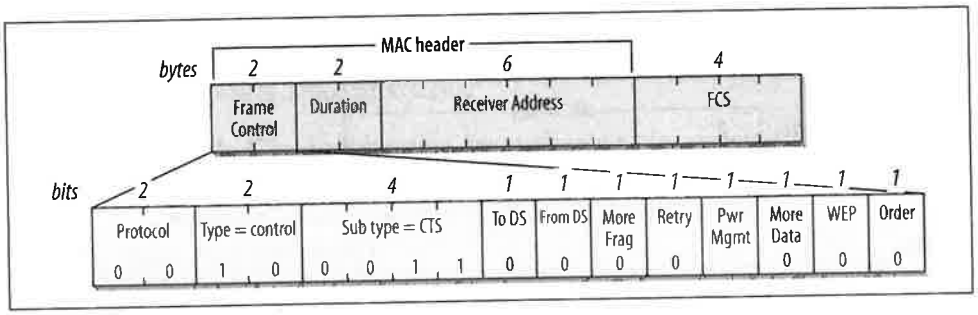

Figure 4-15. CTS frame

Three fields make up the MAC header of a CTS frame:

Frame Control
The frame subtype is set to 1100 to indicate a CTS frame.

Duration
The sender of a CTS frame uses the duration from the RTS frame as the basis for its duration calculation. RTS frames reserve the medium for the entire RTS-CTS-frame-ACK exchange. By the time the CTS frame is transmitted, though, only the pending frame or fragment and its acknowledgment remain. The sender of a CTS frame subtracts the time required for the CTS frame and the short inter-frame space that preceded the CTS from the duration in the RTS frame, and places the result of that calculation in the Duration field. Figure 4-16 illustrates the relationship between the CTS duration and the RTS duration.

Address 1: Receiver Address
The receiver of a CTS frame is the transmitter of the previous RTS frame, so the MAC copies the transmitter address of the RTS frame into the receiver address of the CTS frame.

Acknowledgment (ACK)

ACK frames are used to send the positive acknowledgments required by the MAC and are used with any data transmission, including plain transmissions; frames preceded by an RTS/CTS handshake; and fragmented frames (see Figure 4-17).

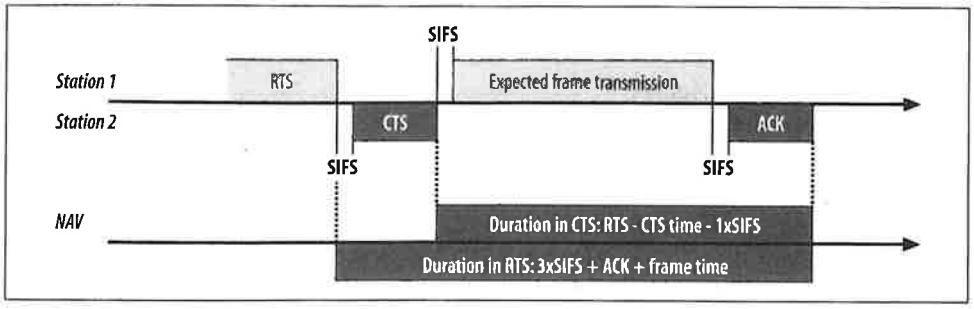

Figure 4-16. CTS duration

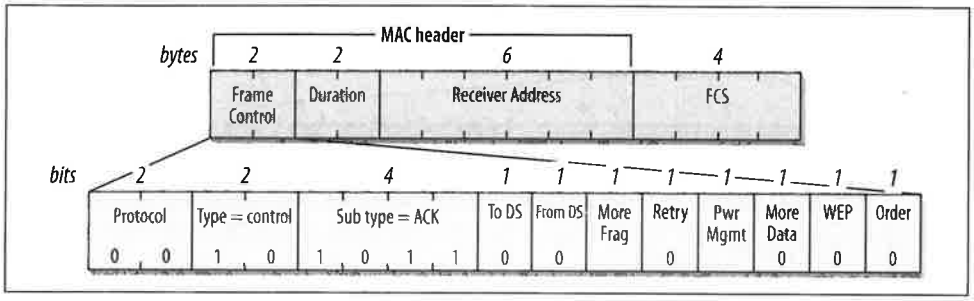

Figure 4-17. ACK frame

Three fields make up the MAC header of an ACK frame:

Frame Control
: The frame subtype is set to 1101 to indicate an ACK frame.

Duration
: The duration may be set in one of two ways, depending on the position of the ACK within the frame exchange. ACKs for complete data frames and final fragments in a fragment burst set the duration to 0. The data sender indicates the end of a data transmission by setting the More Fragments bit in the Frame Control header to 0. If the More Fragments bit is 0, the transmission is complete, and there is no need to extend control over the radio channel for additional transmissions. Thus, the duration is set to 0.

  If the More Fragments bit is 1, a fragment burst is in progress. The Duration field is used like the Duration field in the CTS frame. The time required to transmit the ACK and its short interframe space is subtracted from the duration in the most recent fragment (Figure 4-18). The duration calculation in nonfinal ACK frames is similar to the CTS duration calculation. In fact, the 802.11 specification refers to the duration setting in the ACK frames as a *virtual CTS*.

Address 1: Receiver Address
: The receiver address is copied from the transmitter of the frame being acknowledged. Technically, it is copied from the Address 2 field of the frame being

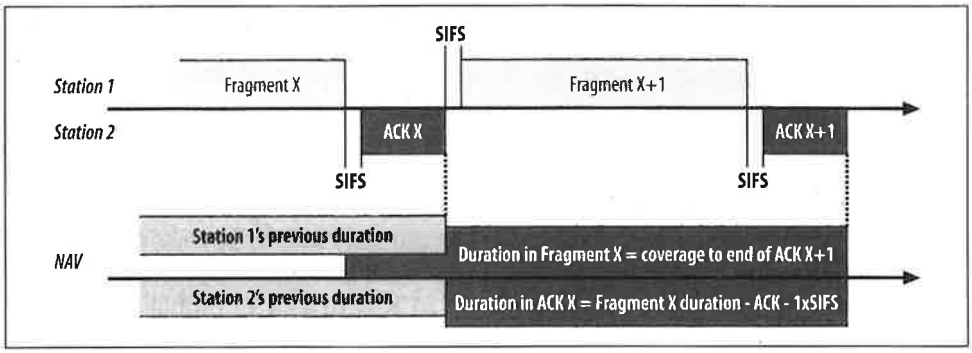

Figure 4-18. Duration in non-final ACK frames

acknowledged. Acknowledgments are transmitted in response to directed data frames, management frames, and PS-Poll frames.

Power-Save Poll (PS-Poll)

When a mobile station wakes from a power-saving mode, it transmits a PS-Poll frame to the access point to retrieve any frames buffered while it was in power-saving mode. The format of the PS-Poll frame is shown in Figure 4-19.

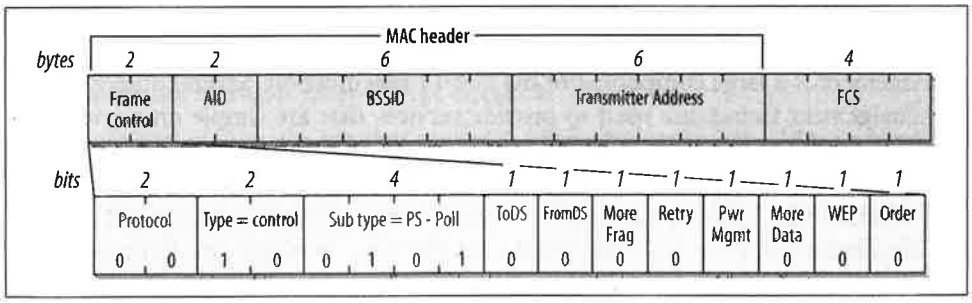

Figure 4-19. PS-Poll frame

Four fields make up the MAC header of a PS-Poll frame:

Frame Control
: The frame subtype is set to 1010 to indicate a PS-Poll frame.

Association ID (AID)
: Instead of a Duration field, the PS-Poll frame uses the third and fourth bytes in the MAC header for the association ID. This is a numeric value assigned by the access point to identify the association. Including this ID in the frame allows the access point to find any frames buffered for the now-awakened mobile station.

Address 1: BSSID

This field contains the BSSID of the BSS created by the access point that the sender is currently associated with.

Address 2: Transmitter Address

This is the address of the sender of the PS-Poll frame.

The PS-Poll frame does not include duration information to update the NAV. However, all stations receiving a PS-Poll frame update the NAV by the short interframe space plus the amount of time required to transmit an ACK. The automatic NAV update allows the access point to transmit an ACK with a small probability of collision with a mobile station.

Association ID (AID)

In the PS-Poll frame, the Duration/ID field is an association ID rather than a value used by the virtual carrier-sensing function. When mobile stations associate with an access point, the access point assigns a value called the Association ID (AID) from the range 1–2,007. The AID is used for a variety of purposes that appear throughout this book.

Management Frames

Management is a large component of the 802.11 specification. Several different types of management frames are used to provide services that are simple on a wired network. Establishing the identity of a network station is easy on a wired network because network connections require dragging wires from a central location to the new workstation. In many cases, patch panels in the wiring closet are used to speed up installation, but the essential point remains: new network connections can be authenticated by a personal visit when the new connection is brought up.

Wireless networks must create management features to provide similar functionality. 802.11 breaks the procedure up into three components. Mobile stations in search of connectivity must first locate a compatible wireless network to use for access. With wired networks, this step is typically finding the appropriate data jack on the wall. Next, the network must authenticate mobile stations to establish that the authenticated identity is allowed to connect to the network. The wired-network equivalent is provided by the network itself. If signals cannot leave the wire, obtaining physical access is at least something of an authentication process. Finally, mobile stations must associate with an access point to gain access to the wired backbone, a step equivalent to plugging the cable into a wired network.

The Structure of Management Frames

802.11 management frames share the structure shown in Figure 4-20.

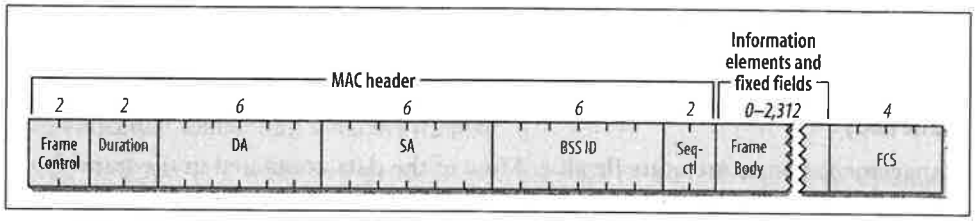

Figure 4-20. Generic management frame

The MAC header is the same in all management frames; it does not depend on the frame subtype. Some management frames use the frame body to transmit information specific to the management frame subtype.

Address fields

As with all other frames, the first address field is used for the frame's destination address. Some management frames are used to maintain properties within a single BSS. To limit the effect of broadcast and multicast management frames, stations inspect the BSSID after receiving a mangement frame. Only broadcast and multicast frames from the BSSID that a station is currently associated with are passed to MAC management layers. The one exception to this rule is Beacon frames, which are used to announce the existence of an 802.11 network.

BSSIDs are assigned in the familiar manner. Access points use the MAC address of the wireless network interface as the BSSID. Mobile stations adopt the BSSID of the access point they are currently associated with. Stations in an IBSS use the randomly generated BSSID from the BSS creation. One exception to the rule: frames sent by the mobile station seeking a specific network may use the BSSID of the network they are seeking, or they may use the broadcast BSSID to find all networks in the vicinity.

Duration calculations

Management frames use the Duration field in the same manner that other frames do:

1. Any frames transmitted in the contention-free period set the duration to 32,768.
2. Frames transmitted during the contention-based access periods using only the DCF use the Duration field to block access to the medium to allow any atomic frame exchanges to complete.
    a. If the frame is a broadcast or multicast (the destination address is a group address), the duration is 0. Broadcast and multicast frames are not acknowledged, so the NAV is not needed to block access to the medium.

b. If a nonfinal fragment is part of a multiframe exchange, the duration is set to the number of microseconds taken up by three SIFS intervals plus the next fragment and its acknowledgment.

c. Final fragments use a duration that is the time required for one acknowledgment plus one SIFS.

Frame body

Management frames are quite flexible. Most of the data contained in the frame body uses fixed-length fields called *fixed fields* and variable-length fields called *information elements*. Information elements are blobs of data of varying size. Each data blob is tagged with a type number and a size, and it is understood that an information element of a certain type has its data field interpreted in a certain way. New information elements can be defined by newer revisions to the 802.11 specification; implementations that predate the revisions can ignore newer elements. Old implementations depend on backward-compatible hardware and frequently can't join networks based on the newer standards. Fortunately, new options usually can be easily turned off for compatibility.

This section presents the fixed fields and information elements as building blocks and shows how the building blocks are assembled into management frames. 802.11 mandates the order in which information elements appear, but not all elements are mandatory. This book shows all the frame building blocks in the specified order, and the discussion of each subtype notes which elements are rare and which are mutually exclusive.

Fixed-Length Management Frame Components

Ten fixed-length fields may appear in management frames. Fixed-length fields are often referred to simply as *fields* to distinguish them from the variable-length information elements.

Authentication Algorithm Number

Two bytes are used for the Authentication Algorithm Number field, shown in Figure 4-21. This field identifies the type of authentication used in the authentication process. (The authentication process is discussed more thoroughly in Chapter 7.) The values permitted for this field are shown in Table 4-3. Only two values are currently defined. Other values are reserved for future standardization work.

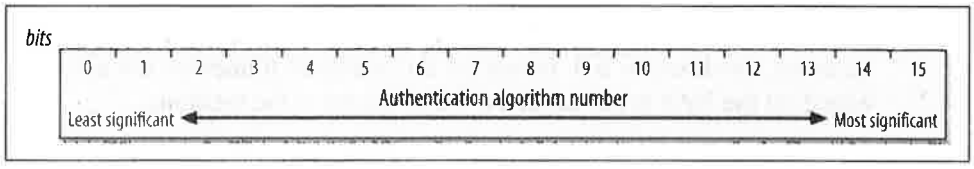

Figure 4-21. Authentication Algorithm Number field

Table 4-3. Values of the Authentication Algorithm Number field

Value	Meaning
0	Open System authentication
1	Shared Key authentication
2–65,535	Reserved

Authentication Transaction Sequence Number

Authentication is a multistep process that consists of a challenge from the access point and a response from the mobile station attempting to associate. The Authentication Transaction Sequence Number, shown in Figure 4-22, is a two-byte field used to track progress through the authentication exchange. It takes values from 1 to 65,535; it is never set to 0. Use of this field is discussed in Chapter 7.

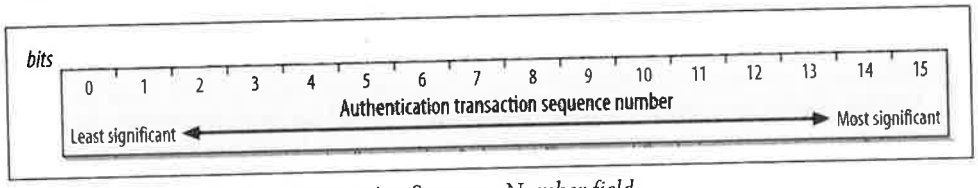

Figure 4-22. Authentication Transaction Sequence Number field

Beacon interval

Beacon transmissions announce the existence of an 802.11 network at regular intervals. Beacon frames carry information about the BSS parameters and the frames buffered by access points, so mobile stations must listen to Beacons. The Beacon Interval, shown in Figure 4-23, is a 16-bit field set to the number of *time units* between Beacon transmissions. One time unit, which is often abbreviated TU, is 1,024 microseconds (μs), which is about 1 millisecond. Time units may also be called kilo-microseconds in various documentation (Kμs or kμs). It is common for the Beacon interval to be set to 100 time units, which corresponds to an interval between Beacon transmissions of approximately 100 milliseconds or 0.1 seconds.

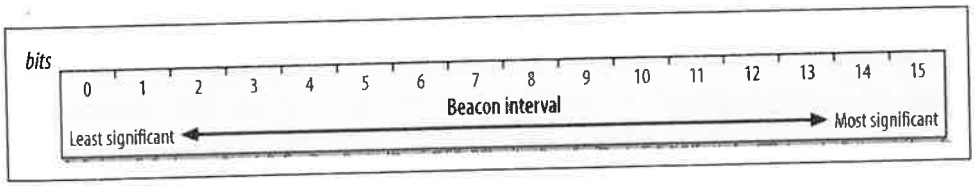

Figure 4-23. Beacon Interval field

Capability Information

The 16-bit Capability Information field, shown in Figure 4-24, is used in Beacon transmissions to advertise the network's capabilities. Capability Information is also used in Probe Request and Probe Response frames. In this field, each bit is used as a

flag to advertise a particular function of the network. Stations use the capability advertisement to determine whether they can support all the features in the BSS. Stations that do not implement all the features in the capability advertisement are not allowed to join.

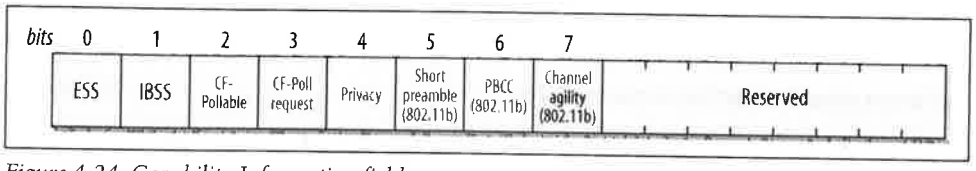

Figure 4-24. Capability Information field

ESS/IBSS

These two bits are mutually exclusive. Access points set the ESS field to 1 and the IBSS field to 0 to indicate that the access point is part of an infrastructure network. Stations in an IBSS set the ESS field to 0 and the IBSS field to 1.

Privacy

Setting the Privacy bit to 1 requires the use of WEP for confidentiality. In infrastructure networks, the transmitter is an access point. In IBSSs, Beacon transmission must be handled by a station in the IBSS.

Short Preamble

This field was added to 802.11b to support the high-rate DSSS PHY. Setting it to 1 indicates that the network is using the short preamble as described in Chapter 10. Zero means the option is not in use and is forbidden in the BSS.

PBCC

This field was added to 802.11b to support the high-rate DSSS PHY. When it is set to 1, it indicates that the network is using the packet binary convolution coding modulation scheme described in Chapter 10. Zero means that the option is not in use and is forbidden in the BSS.

Channel Agility

This field was added to 802.11b to support the high rate DSSS PHY. When it is set to one, it indicates that the network is using the Channel Agility option described in Chapter 10. Zero means the option is not in use and is forbidden in the BSS.

Contention-free polling bits

Stations and access points use these two bits as a label. The meanings of the labels are shown in Table 4-4.

Table 4-4. Interpretation of polling bits in Capability Information

CF-Pollable Station usage	CF-Poll Request	Interpretation
0	0	Station does not support polling
0	1	Station supports polling but does not request to be put on the polling list

Table 4-4. Interpretation of polling bits in Capability Information (continued)

CF-Pollable	CF-Poll Request	Interpretation
1	0	Station supports polling and requests a position on the polling list
1	1	Station supports polling and requests that it never be polled (results in station treated as if it does not support contention-free operation)
Access point usage		
0	0	Access point does not implement the point coordination function
0	1	Access point uses PCF for delivery but does not support polling
1	0	Access point uses PCF for delivery and polling
1	1	Reserved; unused

Current AP Address

Mobile stations use the Current AP Address field, shown in Figure 4-25, to indicate the MAC address of the access point with which they are associated. This field is used to ease associations and reassociations. Stations transmit the address of the access point that handled the last association with the network. When an association is established with a different access point, this field can be used to transfer the association and retrieve any buffered frames.

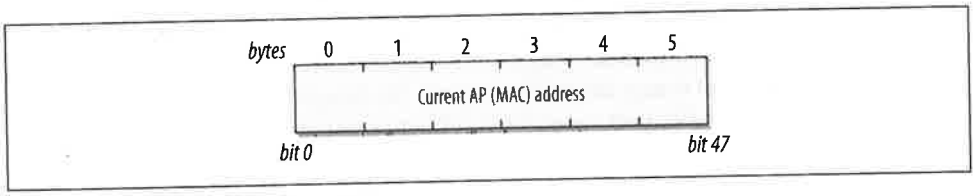

Figure 4-25. Current AP Address field

Listen interval

To save battery power, stations may shut off the antenna units in 802.11 network interfaces. While stations are sleeping, access points must buffer frames for them. Dozing stations periodically wake up to listen to traffic announcements to determine whether the access point has any buffered frames. When stations associate with an access point, part of the saved data is the *Listen Interval*, which is the number of Beacon intervals that stations wait between listening for Beacon frames. The Listen Interval, shown in Figure 4-26, allows mobile stations to indicate how long the access point must retain buffered frames. Higher listen intervals require more access point memory for frame buffering. Access points may use this feature to estimate the resources that will be required and may refuse resource-intensive associations. The Listen Interval is described in Chapter 7.

Association ID

The Association ID, shown in Figure 4-27, is a 16-bit field. When stations associate with an access point, they are assigned an Association ID to assist with control and

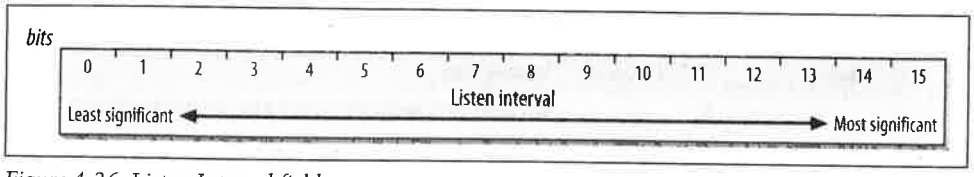

Figure 4-26. Listen Interval field

management functions. Even though 14 bits are available for use in creating Association IDs, they range only from 1–2,007. To maintain compatibility with the Duration/ID field in the MAC header, the two most significant bits are set to 1.

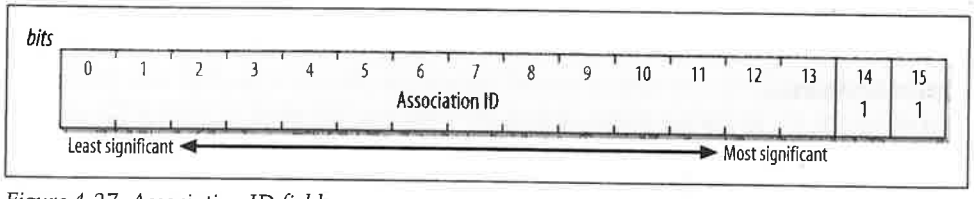

Figure 4-27. Association ID field

Timestamp

The Timestamp field, shown in Figure 4-28, allows synchronization between the stations in a BSS. The master timekeeper for a BSS periodically transmits the number of microseconds it has been active. When the counter reaches its maximum value, it wraps around. (Counter wraps are unlikely given the length of time it takes to wrap a 64-bit counter. At over 580,000 years, I would bet on a required patch or two before the counter wrap.)

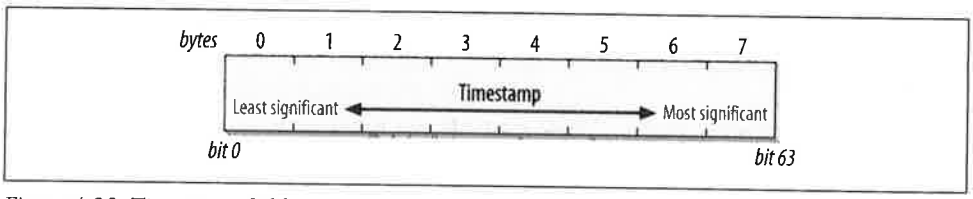

Figure 4-28. Timestamp field

Reason Code

Stations may send Disassociation or Deauthentication frames in response to traffic when the sender has not properly joined the network. Part of the frame is a 16-bit Reason Code field, shown in Figure 4-29, to indicate what the sender has done incorrectly. Table 4-5 shows why certain reason codes are generated. Fully understanding the use of reason codes requires an understanding of the different classes of frames and states of the 802.11 station, which is discussed in the section "Frame Transmission and Association and Authentication States."

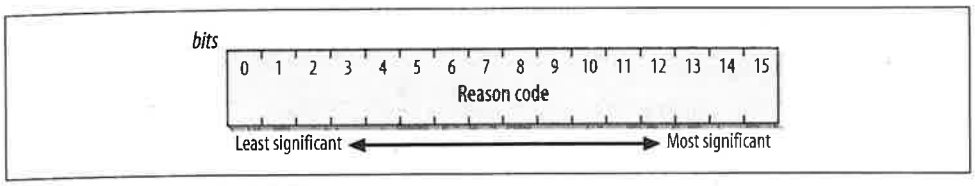

Figure 4-29. Reason Code field

Table 4-5. Reason codes

Code	Explanation
0	Reserved; unused
1	Unspecified
2	Prior authentication is not valid
3	Station has left the basic service area or extended service area and is deauthenticated
4	Inactivity timer expired and station was disassociated
5	Disassociated due to insufficient resources at the access point
6	Incorrect frame type or subtype received from unauthenticated station
7	Incorrect frame type or subtype received from unassociated station
8	Station has left the basic service area or extended service area and is disassociated
9	Association or reassociation requested before authentication is complete
10–65,535	Reserved; unused

Status Code

Status codes indicate the success or failure of an operation. The Status Code field, shown in Figure 4-30, is 0 when an operation succeeds and nonzero on failure. Table 4-6 shows the status codes that have been standardized.

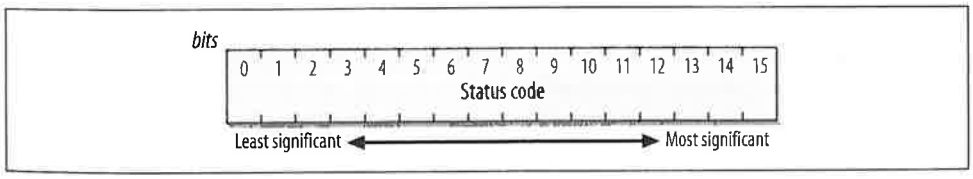

Figure 4-30. Status Code field

Table 4-6. Status codes

Code	Explanation
0	Operation completed successfully
1	Unspecified failure
2–9	Reserved; unused
10	Requested capability set is too broad and cannot be supported
11	Reassociation denied; prior association cannot be identified and transferred

Table 4-6. Status codes (continued)

Code	Explanation
12	Association denied for a reason not specified in the 802.11 standard
13	Requested authentication algorithm not supported
14	Unexpected authentication sequence number
15	Authentication rejected; the response to the challenge failed
16	Authentication rejected; the next frame in the sequence did not arrive in the expected window
17	Association denied; the access point is resource-constrained
18	Association denied; the mobile station does not support all of the data rates required by the BSS
19 (802.11b)	Association denied; the mobile station does not support the Short Preamble option
20 (802.11b)	Association denied; the mobile station does not support the PBCC modulation option
21 (802.11b)	Association denied; the mobile station does not support the Channel Agility option
22–65,535	Reserved for future standardization work

Management Frame Information Elements

Information elements are variable-length components of management frames. A generic information element has an ID number, a length, and a variable-length component, as shown in Figure 4-31. Standardized values for the element ID number are shown in Table 4-7.

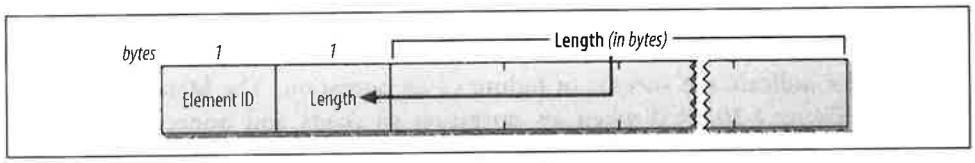

Figure 4-31. Generic management frame information element

Table 4-7. Information elements

Element ID	Name
0	Service Set Identity (SSID)
1	Supported Rates
2	FH Parameter Set
3	DS Parameter Set
4	CF Parameter Set
5	Traffic Indication Map (TIM)
6	IBSS Parameter Set
7–15	Reserved; unused
16	Challenge text
17–31	Reserved for challenge text extension
32–255	Reserved; unused

Service Set Identity (SSID)

Network managers are only human, and they usually prefer to work with letters, numbers, and names rather than 48-bit identifiers. 802.11 networks, in the broadest sense, are either extended service sets or independent BSSs. The SSID, shown in Figure 4-32, allows network managers to assign an identifier to the service set. Stations attempting to join a network may scan an area for available networks and join the network with a specified SSID. The SSID is the same for all the basic service areas composing an extended service area.

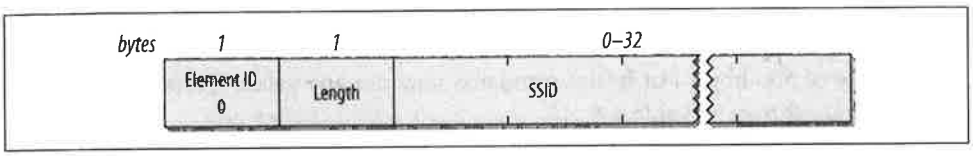

Figure 4-32. Service Set Identity information element

Some documentation refers to the SSID as the *network name* because network administrators frequently assign a character string to it. Most products require that the string be a garden variety, null-terminated ASCII string. In all cases, the length of the SSID ranges between 0 and 32 bytes. The zero-byte case is a special case called the *broadcast SSID*; it is used only in Probe Request frames when a station attempts to discover all the 802.11 networks in its area.

Supported Rates

Several data rates have been standardized for wireless LANs. The Supported Rates information element allows an 802.11 network to specify the data rates it supports. When mobile stations attempt to join the network, they check the data rates used in the network. Some rates are mandatory and must be supported by the mobile station, while others are optional.

The Supported Rates information element is shown in Figure 4-33. It consists of a string of bytes. Each byte uses the seven low-order bits for the data rate; the most significant bit indicates whether the data rate is mandatory. Mandatory rates are encoded with the most significant bit set to 1 and optional rates have a 0. Up to eight rates may be encoded in the information element.

In the initial revision of the 802.11 specification, the seven bits encoded the data rate as a multiple of 500 kbps. New technology, especially ETSI's HIPERLAN efforts, required a change to the interpretation. When seven bits are used to have a multiple of 500 kbps, the maximum data rate that can be encoded is 63.5 Mbps. Research and development on wireless LAN technology has made this rate achievable in the near future. As a result, the IEEE changed the interpretation from a multiple of 500 kbps to a simple label in 802.11b. Previously standardized rates were given labels corresponding to

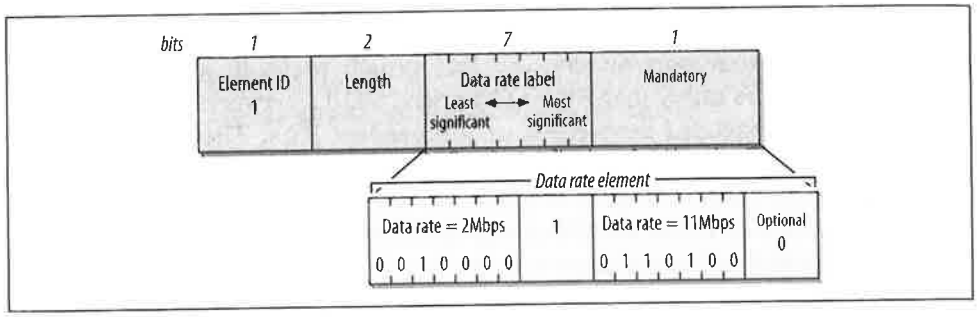

Figure 4-33. Supported Rates information element

the multiple of 500 kbps, but future standards may use any value. Currently standardized values are shown in Table 4-8.

Table 4-8. Supported Rate labels

Binary value	Corresponding rate
2	1 Mbps
4	2 Mbps
11	5.5 Mbps
22	11 Mbps

As an example, Figure 4-33 shows the encoding of two data rates. Two-Mbps service is mandatory, and 11-Mbps service is supported. This is encoded as a mandatory 2-Mbps rate and an optional 11-Mbps rate.

FH Parameter Set

The FH Parameter Set information element, shown in Figure 4-34, contains all parameters necessary to join a frequency-hopping 802.11 network.

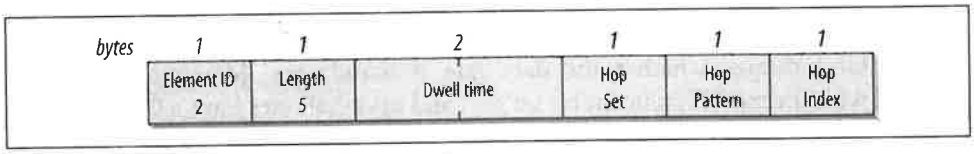

Figure 4-34. FH Parameter Set information element

The FH Parameter Set has four fields that uniquely specify an 802.11 network based on frequency hopping. Chapter 10 describes these identifiers in depth.

Dwell Time

802.11 FH networks hop from channel to channel. The amount of time spent on each channel in the hopping sequence is called the *dwell time*. It is expressed in time units (TUs).

Hop Set
: Several hopping patterns are defined by the 802.11 frequency-hopping PHY. This field, a single byte, identifies the set of hop patterns in use.

Hop Pattern
: Stations select one of the hopping patterns from the set. This field, also a single byte, identifies the hopping pattern in use.

Hop Index
: Each pattern consists of a long sequence of channel hops. This field, a single byte, identifies the current point in the hop sequence.

DS Parameter Set

Direct-sequence 802.11 networks have only one parameter: the channel number used by the network. High-rate, distribution ststem networks use the same channels and thus can use the same parameter set. The channel number is encoded as a single byte, as shown in Figure 4-35.

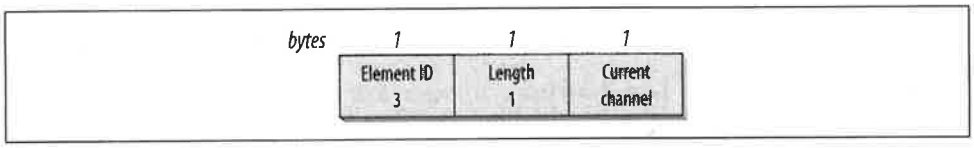

Figure 4-35. DS Parameter Set information element

Traffic Indication Map (TIM)

Access points buffer frames for mobile stations sleeping in low-power mode. Periodically, the access point attempts to deliver buffered frames to sleeping stations. A practical reason for this arrangement is that much more power is required to power up a transmitter than to simply turn on a receiver. The designers of 802.11 envisioned battery-powered mobile stations; the decision to have buffered frames delivered to stations periodically was a way to extend battery life for low-power devices.

Part of this operation is to send the Traffic Indication Map (TIM) information element (Figure 4-36) to the network to indicate which stations have buffered traffic waiting to be picked up.

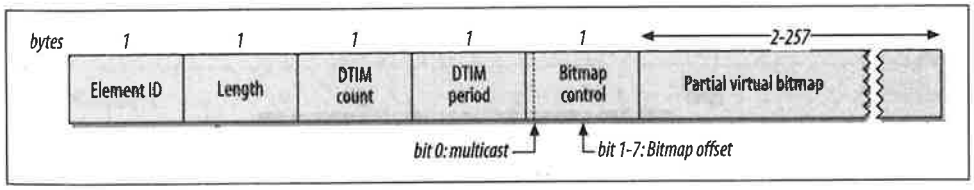

Figure 4-36. Traffic Indication Map information element

The meat of the traffic indication map is the *virtual bitmap*, a logical structure composed of 2,008 bits. Each bit is tied to the Association ID. When traffic is buffered for that Association ID, the bit is 1. If no traffic is buffered, the bit tied to the Association ID is 0.

Four fields make up the body of the TIM information element:

DTIM Count
: This one-byte field is the number of Beacons that will be transmitted before the next DTIM frame. DTIM frames indicate that buffered broadcast and multicast frames will be delivered shortly. Not all Beacon frames are DTIM frames.

DTIM Period
: This one-byte field indicates the number of Beacon intervals between DTIM frames. Zero is reserved and is not used. The DTIM count cycles through from the period down to 0.

Bitmap Control and Partial Virtual Bitmap
: The Bitmap Control field is divided into two subfields. Bit 0 is used for the traffic indication status of Association ID 0, which is reserved for multicast traffic. The remaining seven bits of the Bitmap Control field are used for the Bitmap Offset field.

  To save transmission capacity, the Bitmap Offset field can be used to transmit a portion of the virtual bitmap. The Bitmap Offset is related to the start of the virtual bitmap. By using the Bitmap Offset and the Length, 802.11 stations can infer which part of the virtual bitmap is included.

CF Parameter Set

The CF Parameter Set information element is transmitted in Beacons by access points that support contention-free operation. Contention-free service is discussed in Chapter 8 because of its optional nature.

IBSS Parameter Set

IBSSs currently have only one parameter, the announcement traffic indication map (ATIM) window, shown in Figure 4-37. This field is used only in IBSS Beacon frames. It indicates the number of time units (TUs) between ATIM frames in an IBSS.

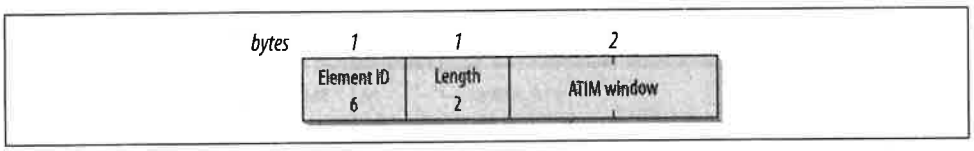

Figure 4-37. IBSS Parameter Set information element

Challenge Text

The shared-key authentication system defined by 802.11 requires that the mobile station successfully decrypt an encrypted challenge. The challenge is sent using the Challenge Text information element, which is shown in Figure 4-38.

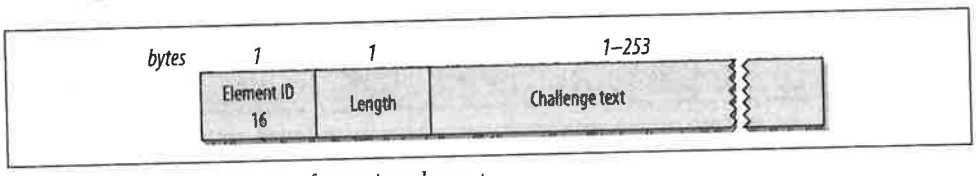

Figure 4-38. Challenge Text information element

Types of Management Frames

The fixed fields and information elements are used in the body of management frames to convey information. Several types of management frames exist and are used for various link-layer maintenance functions.

Beacon

Beacon frames announce the existence of a network and are an important part of many network maintenance tasks. They are transmitted at regular intervals to allow mobile stations to find and identify a network, as well as match parameters for joining the network. In an infrastructure network, the access point is responsible for transmitting Beacon frames. The area in which Beacon frames appear defines the basic service area. All communication in an infrastructure network is done through an access point, so stations on the network must be close enough to hear the Beacons.

Figure 4-39 shows all the fields that can be used in a Beacon frame in the order in which they appear. Not all of the elements are present in all Beacons. Optional fields are present only when there is a reason for them to be used. The FH and DS Parameter Sets are used only when the underlying physical layer is based on frequency hopping or direct-sequence techniques. Only one physical layer can be in use at any point, so the FH and DS Parameter Sets are mutually exclusive.

The CF Parameter Set is used only in frames generated by access points that support the PCF, which is optional. The TIM element is used only in Beacons generated by access points, because only access points perform frame buffering.

Probe Request

Mobile stations use Probe Request frames to scan an area for existing 802.11 networks. The format of the Probe Request frame is shown in Figure 4-40. All fields are mandatory.

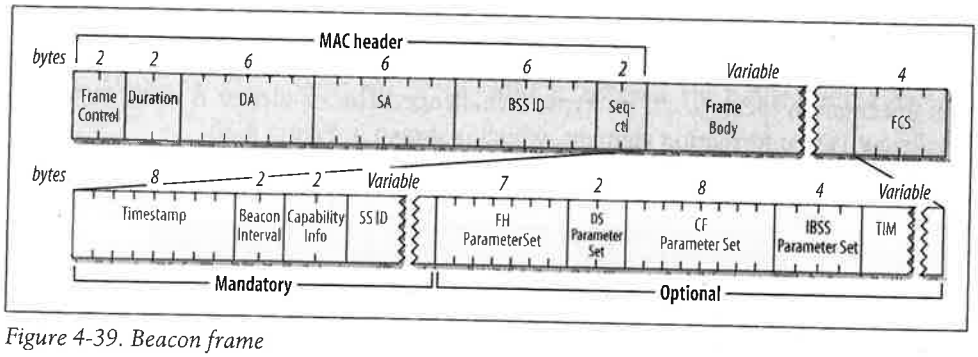

Figure 4-39. Beacon frame

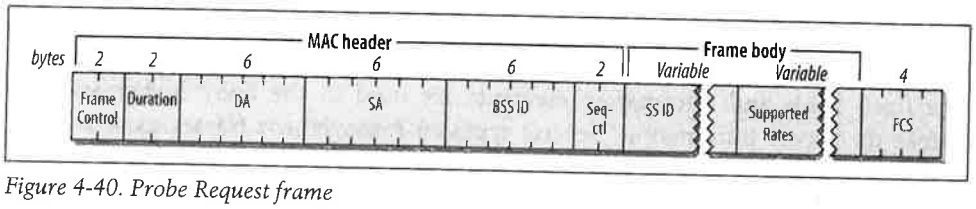

Figure 4-40. Probe Request frame

A Probe Request frame contains two fields: the SSID and the rates supported by the mobile station. Stations that receive Probe Requests use the information to determine whether the mobile station can join the network. To make a happy union, the mobile station must support all the data rates required by the network and must want to join the network identified by the SSID. This may be set to the SSID of a specific network or set to join any compatible network. Drivers that allow cards to join any network use the broadcast SSID in Probe Requests.

Probe Response

If a Probe Request encounters a network with compatible parameters, the network sends a Probe Response frame. The station that sent the last Beacon is responsible for responding to incoming probes. In infrastructure networks, this station is the access point. In an IBSS, responsibility for Beacon transmission is distributed. After a station transmits a Beacon, it assumes responsibility for sending Probe Response frames for the next Beacon interval. The format of the Probe Response frame is shown in Figure 4-41. Some of the fields in the frame are mutually exclusive; the same rules apply to Probe Response frames as to Beacon frames.

The Probe Response frame carries all the parameters in a Beacon frame, which enables mobile stations to match parameters and join the network. Probe Response frames can safely leave out the TIM element because stations sending probes are not yet associated and thus would not need to know which associations have buffered frames waiting at the access point.

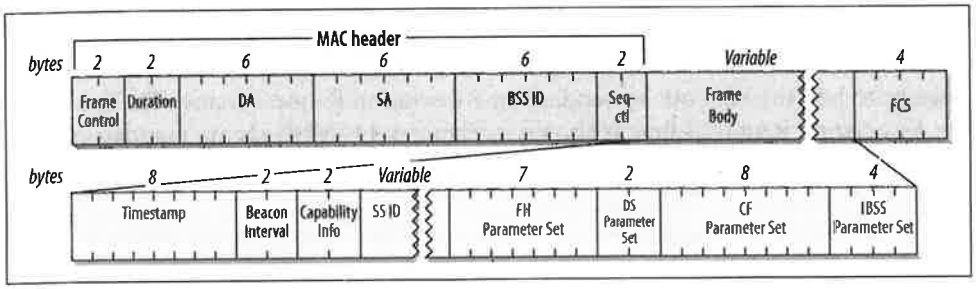

Figure 4-41. Probe Response frame

IBSS announcement traffic indication map (ATIM)

IBSSs have no access points and therefore cannot rely on access points for buffering. When a station in an IBSS has buffered frames for a receiver in low-power mode, it sends an ATIM frame during the delivery period to notify the recipient it has buffered data. See Figure 4-42.

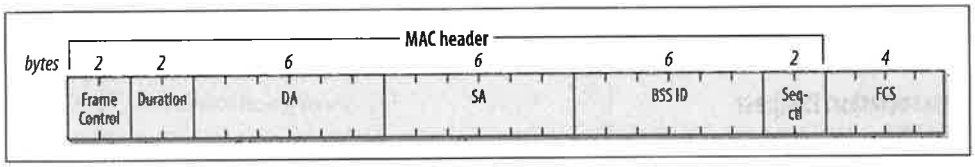

Figure 4-42. ATIM frame

Disassociation and Deauthentication

Disassociation frames are used to end an association relationship, and Deauthentication frames are used to end an authentication relationship. Both frames include a single fixed field, the Reason Code, as shown in Figure 4-43. Of course, the Frame Control fields differ because the subtype distinguishes between the different types of management frames.

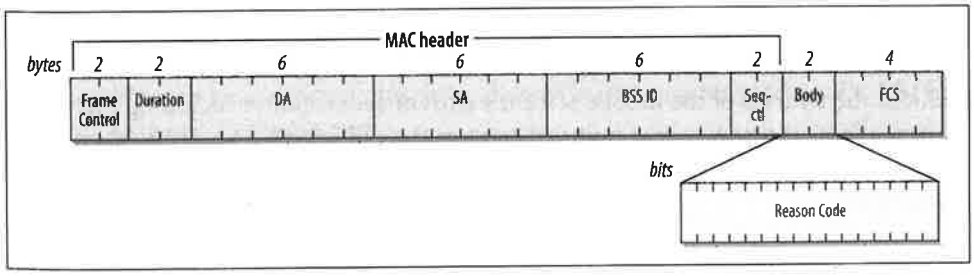

Figure 4-43. Disassociation and Deauthentication frames

Association Request

Once mobile stations identify a compatible network and authenticate to it, they may attempt to join the network by sending an Association Request frame. The format of the Association Request frame is shown in Figure 4-44. All fields are mandatory, and they must appear in the order shown.

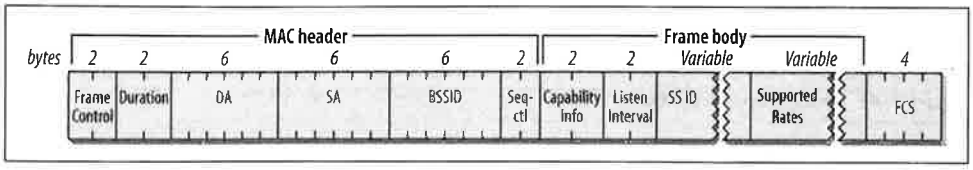

Figure 4-44. Association Request frame

The Capability Information field is used to indicate the type of network the mobile station wants to join. Before an access point accepts an association request, it verifies that the Capability Information, SSID, and Supported Rates all match the parameters of the network. Access points also note the Listen Interval, which describes how often a mobile station listens to Beacon frames to monitor the TIM.

Reassociation Request

Mobile stations moving between basic service areas within the same extended service area need to reassociate with the network before using the distribution system again. Stations may also need to reassociate if they leave the coverage area of an access point temporarily and rejoin it later. See Figure 4-45.

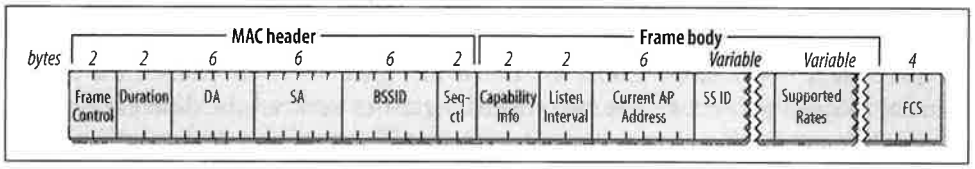

Figure 4-45. Reassociation Request frame

Association and Reassociation Requests differ only in that a Reassociation Request includes the address of the mobile station's current access point. Including this information allows the new access point to contact the old access point and transfer the association data. The transfer may include frames that were buffered at the old access point.

Association Response and Reassociation Response

When mobile stations attempt to associate with an access point, the access point replies with an Association Response or Reassociation Response frame, shown in Figure 4-46. The two differ only in the subtype field in the Frame Control field. All

fields are mandatory. As part of the response, the access point assigns an Association ID. How an access point assigns the association ID is implementation-dependent.

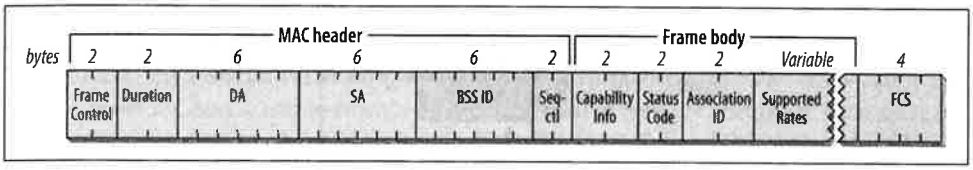

Figure 4-46. (Re)Association Response frame

Authentication

To authenticate to the access point, mobile stations exchange Authentication frames, which are shown in Figure 4-47.

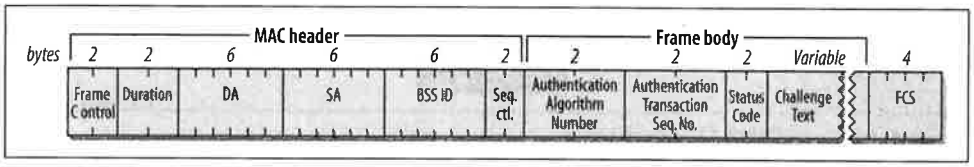

Figure 4-47. Authentication frames

Different authentication algorithms may co-exist. The Authentication Algorithm Number field is used for algorithm selection. The authentication process may involve a number of steps (depending on the algorithm), so there is a sequence number for each frame in the authentication exchange. The Status Code and Challenge Text are used in different ways by different algorithms; details are discussed in Chapter 7.

Frame Transmission and Association and Authentication States

Allowed frame types vary with the association and authentication states. Stations are either authenticated or unauthenticated and can be associated or unassociated. These two variables can be combined into three allowed states, resulting in the 802.11 Hierarchy of Network Development:

1. Initial state; not authenticated and not associated
2. Authenticated but not yet associated
3. Authenticated and associated

Each state is a successively higher point in the development of an 802.11 connection. All mobile stations start in State 1, and data can be transmitted through a distribution system only in State 3. (IBSSs do not have access points or associations and

thus only reach Stage 2.) Figure 4-48 is the overall state diagram for frame transmission in 802.11.

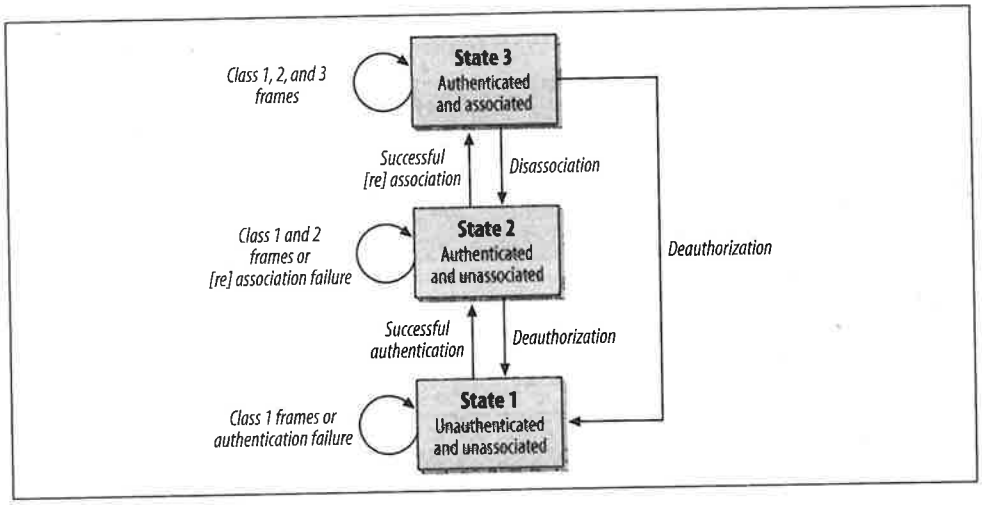

Figure 4-48. Overall 802.11 state diagram

Frame Classes

Frames are also divided into different classes. Class 1 frames can be transmitted in State 1; Class 1 and 2 frames in State 2; and Class 1, 2, and 3 frames in State 3.

Class 1 frames

Class 1 frames may be transmitted in any state and are used to provide the basic operations used by 802.11 stations. Control frames are received and processed to provide basic respect for the CSMA/CA "rules of the road" and to transmit frames in an IBSS. Class 1 frames also allow stations to find an infrastructure network and authenticate to it. Table 4-9 shows a list of the frames that belong to the Class 1 group.

Table 4-9. Class 1 frames

Control	Management	Data
Request to Send (RTS)	Probe Request	Any frame with ToDS and FromDS false (0)
Clear to Send (CTS)	Probe Response	
Acknowledgment (ACK)	Beacon	
CF-End	Authentication	
CF-End+CF-Ack	Deauthentication	
	Announcement Traffic Indication Message (ATIM)	

Class 2 frames

Class 2 frames can be transmitted only after a station has successfully authenticated to the network, and they can be used only in States 2 and 3. Class 2 frames manage associations. Successful association or reassociation requests move a station to State 3; unsuccessful association attempts cause the station to stay in State 2. When a station receives a Class 2 frame from a nonauthenticated peer, it responds with a Deauthentication frame, dropping the peer back to State 1.* Table 4-10 shows the Class 2 frames.

Table 4-10. Class 2 frames

Control	Management	Data
None	Association Request/Response	None
	Reassociation Request/Response	
	Disassociation	

Class 3 frames

Class 3 frames are used when a station has been successfully authenticated and associated with an access point. Once a station has reached State 3, it is allowed to use distribution system services and reach destinations beyond its access point. Stations may also use the power-saving services provided by access points in State 3 by using the PS-Poll frame. Table 4-11 lists the different types of Class 3 frames.

Table 4-11. Class 3 frames

Control	Management	Data
PS-Poll	Deauthentication	Any frames, including those with either the ToDS or FromDS bits set

If an access point receives frames from a mobile station that is authenticated but not associated, the access point responds with a Disassociation frame to bump the mobile station back to State 2. If the mobile station is not even authenticated, the access point responds with a Deauthentication frame to force the mobile station back into State 1.

* This rejection action takes place only for frames that are not filtered. Filtering prevents frames from a different BSS from triggering a rejection.

CHAPTER 5

Wired Equivalent Privacy (WEP)

Anyone who is not shocked by quantum theory has not understood it.

—Niels Bohr

In wireless networks, the word "broadcast" takes on an entirely new meaning. Security concerns have haunted 802.11 deployments since the standardization effort began. IEEE's attempt to address snooping concerns culminated in the optional Wired Equivalent Privacy (WEP) standard, which is found in clause 8.2 of 802.11. WEP can be used by stations to protect data as it traverses the wireless medium, but it provides no protection past the access point.

Many of the headlines about 802.11 over the past year were due to WEP. As networks become important to doing business, security has become an increasingly prominent worry. WEP was initially marketed as the security solution for wireless LANs, though its design was so flawed as to make that impossible.

WEP is so flawed that it is not worth using in many cases. Some of the flaws are severe design flaws, and the complete break of WEP in late 2001 was caused by a latent problem with the cryptographic cipher used by WEP. To understand WEP and its implications for the security of your network, this chapter presents some background on WEP's cryptographic heritage, lists the design flaws, and discusses the final straw. It closes with recommendations on the use of WEP. To make a long chapter much shorter, the basic recommendation is to think very, very carefully before relying on WEP because it has been soundly defeated.

Cryptographic Background to WEP

Before discussing the design of WEP, it's necessary to cover some basic cryptographic concepts. I am not a cryptographer, and a detailed discussion of the cryptog-

raphy involved would not be appropriate in this book, so this chapter is necessarily brief.*

To protect data, WEP requires the use of the RC4 cipher, which is a symmetric (secret-key) stream cipher. RC4 shares a number of properties with all stream ciphers. Generally speaking, a stream cipher uses a stream of bits, called the *keystream*. The keystream is then combined with the message to produce the *ciphertext*. To recover the original message, the receiver processes the ciphertext with an identical keystream. RC4 uses the exclusive OR (XOR) operation to combine the keystream and the ciphertext. Figure 5-1 illustrates the process.

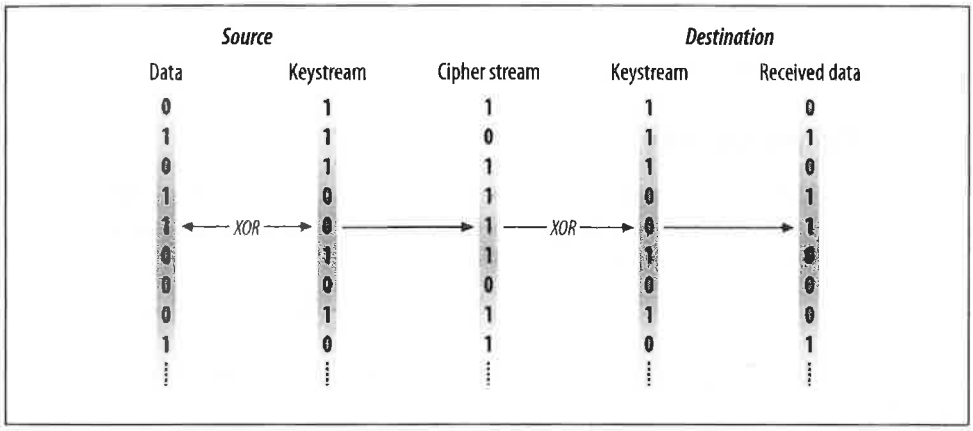

Figure 5-1. Generic stream cipher operation

Most stream ciphers operate by taking a relatively short secret key and expanding it into a pseudorandom keystream the same length as the message. This process is illustrated in Figure 5-2. The pseudorandom number generator (PRNG) is a set of rules used to expand the key into a keystream. To recover the data, both sides must share the same secret key and use the same algorithm to expand the key into a pseudorandom sequence.

Because the security of a stream cipher rests entirely on the randomness of the keystream, the design of the key-to-keystream expansion is of the utmost importance. When RC4 was selected by the 802.11 working group, it appeared to be quite secure. But once RC4 was selected as the ciphering engine of WEP, it spurred research that ultimately found an exploitable flaw in the RC4 cipher that will be discussed later.

* Readers interested in more detailed explanations of the cryptographic algorithms involved should consult *Applied Cryptography* by Bruce Schneier (Wiley, 1996).

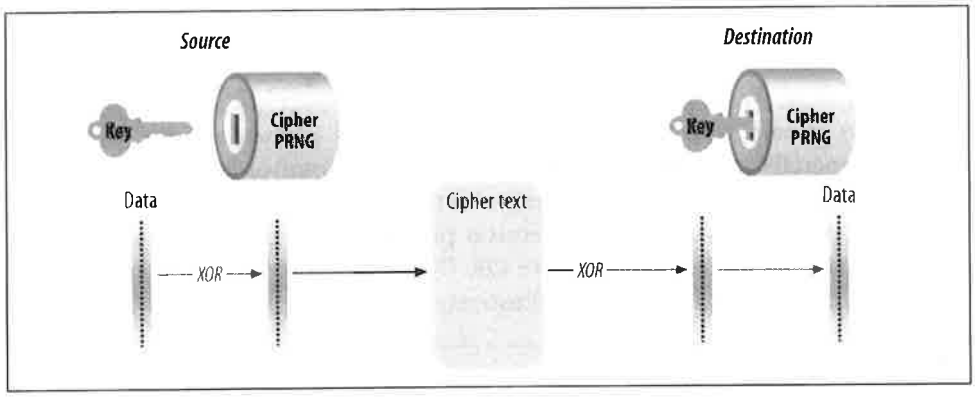

Figure 5-2. Keyed stream cipher operation

Stream Cipher Security

A totally random keystream is called a *one-time pad* and is the only known encryption scheme that is mathematically proven to protect against certain types of attacks. One-time pads are not commonly used because the keystream must be perfectly random and the same length as the data that will be protected, and it can never be reused.

Attackers are not limited to attacking the underlying cipher. They can choose to exploit any weak point in a cryptographic system. One famous Western intelligence effort, code-named VENONA, broke Soviet messages encrypted with one-time pads that were reused. The National Security Agency has made some information on the project public at *http://www.nsa.gov/docs/venona*. It is easy to understand the temptation to reuse the one-time pads. Huge volumes of keying material are necessary to protect even a small amount of data, and those keying pads must be securely distributed, which in practice proves to be a major challenge.

Stream ciphers are a compromise between security and practicality. The perfect randomness (and perfect security) of a one-time pad is attractive, but the practical difficulties and cost incurred in generating and distributing the keying material is worthwhile only for short messages that require the utmost security. Stream ciphers use a less random keystream but one that is random enough for most applications.

Cryptographic Politics

Three major nontechnical concerns may impact the use of WEP:

1. RC4 is the intellectual property of RSA Security, Inc., and must be licensed. RSA would almost certainly file suit against any unlicensed RC4 implementation. For

most end users, this is a minor point because wireless LAN equipment vendors would need to license RC4. In the past, this has been a problem for Linux users because some early wireless cards didn't include WEP on the card, and patents prevented open source developers from implementing it in the device driver. The latest generation of wireless cards solves this problem by implementing WEP on the card itself; all the device driver has to do is load the card with the keys.

2. Products must be exportable from U.S. locations to compete across the world. The 802.11 project committee specifically designed WEP to meet with approval from the U.S. export regulations at the time; as a consequence, WEP implementations were restricted to a maximum key length of 40 bits. Rules have been relaxed since then, and longer keys are allowed. Unfortunately, longer key lengths were never formally specified and may not be interoperable between products from different vendors.
3. Some governments impose restrictions on the importation of cryptographic hardware and software, which may prevent the use of encryption to protect the wireless LAN link. Without even the minimal protection provided by WEP, it may not be worth the risk to use wireless LAN technology in such locations.

WEP Cryptographic Operations

Communications security has three major objectives. Any protocol that attempts to secure data as it travels across a network must help network managers to achieve these goals. *Confidentiality* is the term used to describe data that is protected against interception by unauthorized parties. *Integrity* means that the data has not been modified. *Authentication* underpins any security strategy because part of the reliability of data is based on its origin. Users must ensure that data comes from the source it purports to come from. Systems must use authentication to protect data appropriately. Authorization and access control are both implemented on top of authentication. Before granting access to a piece of data, systems must find out who the user is (authentication) and whether the access operation is allowed (authorization).

WEP provides operations that attempt to help meet these objectives. Frame body encryption supports confidentiality. An integrity check sequence protects data in transit and allows receivers to validate that the received data was not altered in transit. WEP also enables stronger shared-key authentication of stations for access points, a feature discussed in Chapter 7. In practice, WEP falls short in all of these areas. Confidentiality is compromised by flaws in the RC4 cipher; the integrity check was poorly designed; and authentication is of users' MAC addresses, not users themselves.

WEP also suffers from the approach it takes. It encrypts frames as they traverse the wireless medium. Nothing is done to protect frames on a wired backbone, where

they are subject to any attack. Furthermore, WEP is designed to secure the network from external intruders. Once an intruder discovers the WEP key, though, the wireless medium becomes the equivalent of a big shared wired network.

WEP Data Processing

Confidentiality and integrity are handled simultaneously, as illustrated in Figure 5-3. Before encryption, the frame is run through an integrity check algorithm, generating a hash called an integrity check value (ICV). The ICV protects the contents against tampering by ensuring that the frame has not changed in transit. The frame and the ICV are both encrypted, so the ICV is not available to casual attackers.

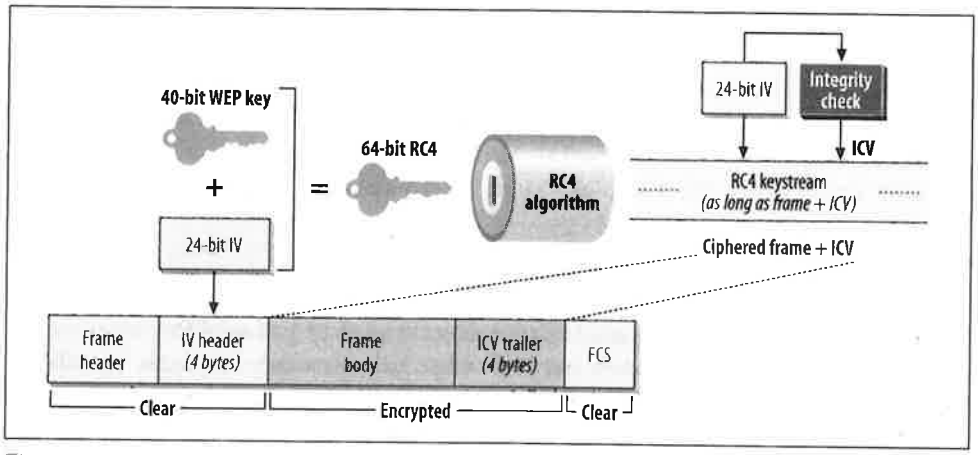

Figure 5-3. WEP operations

WEP specifies the use of a 40-bit secret key. The secret WEP key is combined with a 24-bit initialization vector (IV) to create a 64-bit RC4 key; the first 24 bits of the RC4 key are the IV, followed by the 40-bit WEP key. RC4 takes the 64 input bits and generates a keystream equal to the length of the frame body plus the ICV. The keystream is then XORed with the frame body and the ICV to cipher it. To enable the receiver to decrypt the frame, the IV is placed in the header of the frame.

WEP keying

To protect traffic from brute-force decryption attacks, WEP uses a set of up to four *default keys*, and it may also employ pairwise keys, called *mapped keys*, when allowed. Default keys are shared among all stations in a service set. Once a station has obtained the default keys for its service set, it may communicate using WEP.

Key reuse is often a weakness of cryptographic protocols. For this reason, WEP has a second class of keys used for pairwise communications. These keys are shared only

WEP Key Lengths

Standardized WEP implementations use 64-bit shared RC4 keys. Of the 64 bits, 40 are a shared secret. Vendors use a variety of names for the standard WEP mode: "standard WEP," "802.11-compliant WEP," "40-bit WEP," "40+24-bit WEP," or even "64-bit WEP." I personally feel that the last term is a stretch, based on hoodwinking the consumer with the length of the shared key and not the size of the shared secret, but it has become somewhat standard throughout the industry.

Concerns about the key length used in WEP have dogged it since its inception. Products that use 40-bit secret keys have always been exportable from the United States, which has served to cast doubt on the security provided by such a key. In a well-designed cryptographic system, additional security can be obtained by using a longer key. Each additional bit doubles the number of potential keys and, in theory, doubles the amount of time required for a successful attack.

To buy time for the standardization of a better solution than WEP, most of the industry moved to a 128-bit shared RC4 key. After subtracting 104 bits for the shared secret component of the RC4 key, only 104 bits are secret. Even though only 104 bits are secret, vendors refer to this as "128-bit WEP." Longer key-length implementations are not guaranteed to be compatible because no standard for them exists. At least one vendor uses 128 secret bits, plus the 24 in the initialization vector, for a total of 152 bits.

WEP, however, is not a well-designed cryptographic system, and the extra bits in the key buy you very little. The best publicly disclosed attack against WEP can recover the key in seconds, no matter what its length is. This book explores the use of the AirSnort tool to recover WEP keys in Chapter 16.

between the two stations communicating. The two stations sharing a key have a *key mapping relationship*; the key mapping relationship is part of the 802.11 MIB, which is presented in detail in Appendix A.

WEP framing

When WEP is in use, the frame body expands by eight bytes. Four bytes are used for a frame body IV header, and four are used for the ICV trailer. See Figure 5-4.

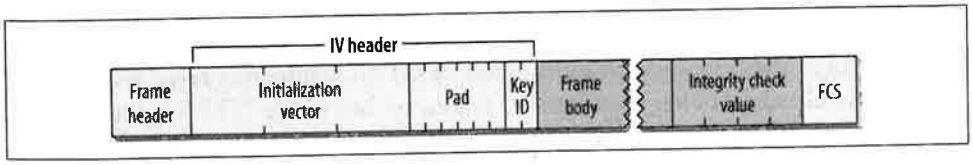

Figure 5-4. WEP frame extensions

The IV header uses 3 bytes for the 24-bit IV, with the fourth byte used for padding and key identification. When a default key is used, the Key ID subfield identifies the default key that was used to encrypt the frame. If a key mapping relationship is used, the Key ID subfield is 0. The 6 padding bits of the last byte must be 0. The integrity check is a 32-bit CRC of the data frame; it is appended to the frame body and protected by RC4.

Cryptographic Properties

Reuse of the keystream is the major weakness in any stream cipher–based cryptosystem. When frames are encrypted with the same RC4 keystream, the XOR of the two encrypted packets is equivalent to the XOR of the two plaintext packets. By analyzing differences between the two streams in conjunction with the structure of the frame body, attackers can learn about the contents of the plaintext frames themselves. To help prevent the reuse of the keystream, WEP uses the IV to encrypt different packets with different RC4 keys. However, the IV is part of the packet header and is not encrypted, so eavesdroppers are tipped off to packets that are encrypted with the same RC4 key.

Implementation problems can contribute to the lack of security. 802.11 admits that using the same IV for a large number of frames is insecure and should be avoided. The standard allows for using a different IV for each frame, but it is not required.

WEP incorporates an integrity check, but the algorithm used is a cyclic redundancy check (CRC). CRCs can catch single-bit alterations with high probability, but they are not *cryptographically secure*. Cryptographically secure integrity checks are based on hash functions, which are unpredictable. With unpredictable functions, if the attacker changes even one bit of the frame, the integrity check will change in unpredictable ways. The odds of an attacker finding an altered frame with the same integrity check are so slim that it cannot be done in real time. CRCs are not cryptographically secure. CRC calculations are straightforward mathematics, and it is easy to predict how changing a single bit will affect the result of the CRC calculation. (This property is often used by compressed data files for repair! If just a few bits are bad, they can sometimes be identified and corrected based on a CRC value.)

Key Distribution

Like so many other cryptographic protocols based on symmetric keys, WEP suffers from the Achilles heel of key distribution. The secret bits of the WEP key must be distributed to all stations participating in an 802.11 service set secured by WEP. The 802.11 standard, however, fails to specify the key distribution mechanism. The result is that vendors haven't done anything; you typically type keys into your device drivers

or access points by hand. Unfortunately, manual configuration by the system administrator is the most nonscalable "protocol" in use.

Setting aside the system management headaches for a minute, consider the difficulties inherent in a cryptographic system requiring manual key distribution:

- Keys cannot be considered secret: all keys must be statically entered into either the driver software or the firmware on the wireless card. Either way, the key cannot be protected from a local user who wants to discover it.*
- If keys are accessible to users, then all keys must be changed whenever staff members leave the organization. Knowledge of WEP keys allows a user to set up an 802.11 station and passively monitor and decrypt traffic using the secret key for the network. WEP cannot protect against authorized insiders who also have the key.
- Organizations with large numbers of authorized users must publish the key to the user population, which effectively prevents it from being a secret. In the course of doing research for this book, I found network documentation at one major research university that described how to use the campus wireless network, including the WEP key.

Problems with WEP

Cryptographers have identified many flaws in WEP. The designers specified the use of RC4, which is widely accepted as a strong cryptographic cipher. Attackers, however, are not limited to a full-frontal assault on the cryptographic algorithms—they can attack any weak point in the cryptographic system. Methods of defeating WEP have come from every angle. One vendor shipped access points that exposed the secret WEP keys through SNMP, allowing an attacker to ask for just the key. Most of the press, though, has been devoted to flaws beyond implementation errors, which are much harder to correct.

Design Flaws

WEP's design flaws initially gained prominence when the Internet Security, Applications, Authentication and Cryptography (ISAAC) group at the University of California, Berkeley, published preliminary results based on their analysis of the WEP standard.† None of the problems identified by researchers depend on breaking RC4.

* Anecdotal evidence suggests that this may be commonplace. Power users who prefer to use Linux or FreeBSD may attempt to recover the key simply to allow access to the network from an otherwise unsupported operating system.

† The report is available on the Web at *http://www.isaac.cs.berkeley.edu/isaac/wep-faq.html*. Items 3–6 on the following list are summarized from that report.

Here's a summary of the problems they found; I've already touched on some of them:

1. Manual key management is a minefield of problems. Setting aside the operational issues with distributing identical shared secrets to the user population, the security concerns are nightmarish. New keying material must be distributed on a "flag day" to all systems simultaneously, and prudent security practices would lean strongly toward rekeying whenever anybody using WEP leaves the company (the administrative burden may, however, preclude doing this). Widely distributed secrets tend to become public over time. Passive sniffing attacks require obtaining only the WEP keys, which are likely to be changed infrequently. Once a user has obtained the WEP keys, sniffing attacks are easy. Market-leading sniffers are now starting to incorporate this capability for system administrators, claiming that after entering the network's WEP keys, all the traffic is readable!
2. In spite of vendor claims to the contrary, standardized WEP offers a shared secret of only 40 bits. Security experts have long questioned the adequacy of 40-bit private keys, and many recommend that sensitive data be protected by at least 128-bit keys.* Unfortunately, no standard has been developed for longer keys, so interoperability on multivendor networks with long WEP keys is not guaranteed without future work by the IEEE.
3. Stream ciphers are vulnerable to analysis when the keystream is reused. WEP's use of the IV tips off an attacker to the reuse of a keystream. Two frames that share the same IV almost certainly use the same secret key and keystream. This problem is made worse by poor implementations, which may not pick random IVs. The Berkeley team identified one implementation that started with an IV of 0 when the card was inserted and simply incremented the IV for each frame. Furthermore, the IV space is quite small (less than 17 million), so repetitions are guaranteed on busy networks.
4. Infrequent rekeying allows attackers to assemble what the Berkeley team calls *decryption dictionaries*—large collections of frames encrypted with the same keystreams. As more frames with the same IV pile up, more information is available about the unencrypted frames even if the secret key is not recovered. Given how overworked the typical system and network administration staff is, infrequent rekeying is the rule.
5. WEP uses a CRC for the integrity check. Although the value of the integrity check is encrypted by the RC4 keystream, CRCs are not cryptographically

* To be fair, WEP was originally developed with the goal of being exportable under the then current U.S. regulations for export of cryptographic systems. A longer key could not have been used without jeopardizing the commercial viability of U.S.-built 802.11 products.

secure. Use of a weak integrity check does not prevent determined attackers from transparently modifying frames.*

6. The access point is in a privileged position to decrypt frames. Conceptually, this feature can be attacked by tricking the access point into retransmitting frames that were encrypted by WEP. Frames received by the access point would be decrypted and then retransmitted to the attacker's station. If the attacker is using WEP, the access point would helpfully encrypt the frame using the attacker's key.

The Complete Break

In August 2001, Scott Fluhrer, Itsik Mantin, and Adi Shamir published a paper titled "Weaknesses in the Key Scheduling Algorithm of RC4." At the end of the paper, the authors describe a theoretical attack on WEP. At the heart of the attack is a weakness in the way that RC4 generates the keystream. All that is assumed is the ability to recover the first byte of the encrypted payload. Unfortunately, 802.11 uses LLC encapsulation, and the cleartext value of the first byte is known to be 0xAA (the first byte of the SNAP header). Because the first cleartext byte is known, the first byte of the keystream can be easily deduced from a trivial XOR operation with the first encrypted byte.

The paper's attacks are focused on a class of weak keys written in the form (B+3):ff:N. Each weak IV is used to attack a particular byte of the secret portion of the RC4 key. Key bytes are numbered from zero. Therefore, the weak IV corresponding to byte zero of the secret key has the form 3:FF:N. The second byte must be 0xFF; knowledge of the third byte in the key is required, but it need not be any specific value.

A standard WEP key is 40 secret bits, or 5 bytes numbered consecutively from 0 to 4. Weak IVs on a network protected by standard WEP must have a first byte that ranges from 3 (B=0) to 7 (B=4) and a second byte of 255. The third byte must be noted but is not constrained to any specific value. There are $5 \times 1 \times 256=1,280$ weak IVs in a standard WEP network.

It is interesting to note that the number of weak keys depends partly on the length of the RC4 key used. If the WEP key size is increased for added protection, the weak key net pulls in more data for use in the attack. Most commercial products use a 128-bit shared RC4 key, so there are more than twice as many weak IVs. Table 5-1 shows the number of weak IVs as a function of the secret key length.

* 802.11 requires frame retransmissions in the case of loss, so it may be possible for an attacker to retransmit a frame and cause a replacement injected frame to be accepted as legitimate.

Table 5-1. Number of weak IVs as a function of key length

Secret key length	Values of B+3 in weak IV (B+3:FF:N)	Number of weak IVs	Fraction of IV space
40 bits	3 <= B+3 < 8 (0 <= B < 5)	1,280	0.008%
104 bits	3 <= B+3 < 16 (0 <= B < 13)	3,328	0.020%
128 bits	3 <= B+3 < 19 (0 <= B < 16)	4,096	0.024%

Applying probability theory, Fluhrer, Mantin, and Shamir predict that about 60 resolved cases are needed to determine a key byte. Furthermore, and perhaps worst of all, the attack gains speed as more key bytes are determined; overall, it works in linear time. Doubling the key length only doubles the time for the attack to succeed.

With such a tantalizing result, it was only a matter of time before it was used to attack a real system. In early August 2001, Adam Stubblefield, John Ioannidis, and Avi Rubin applied the Fluhrer/Mantin/Shamir attack to an experimental, but real, network with devastating effect.* In their testing, 60 resolved cases usually determined a key byte, and 256 resolved cases always yielded a full key. It took less than a week to implement the attack, from the ordering of the wireless card to the recovery of the first full key. Coding the attack took only a few hours. Key recovery was accomplished between five and six million packets, which is a small number for even a moderately busy network.

Reporting on a successful attack, however, is nothing compared to having a public code base available to use at will. The hard part of the Fluhrer/Mantin/Shamir attack was finding the RC4 weakness. Implementing their recommendations is not too difficult. In late August 2001, Jeremy Bruestle and Blake Hegerle released AirSnort, an open source WEP key recovery program. Use of AirSnort is discussed in Chapter 16.

Conclusions and Recommendations

WEP was designed to provide relatively minimal protection to frames in the air. It was not designed for environments demanding a high level of security and therefore offers a comparatively smaller level of protection. The IEEE 802.11 working group has devoted an entire task group to security. The task group is actively working on a revised security standard. In the meantime, some vendors are offering proprietary approaches that allow stronger public-key authentication and random session keys, but these approaches are a single-vendor solution and only a stopgap. Better solutions can be built from off-the-shelf standardized components. Specific topology

* This work is described more fully in AT&T Labs Technical Report TD-4ZCPZZ.

deployments are discussed in Chapter 15. To close, I offer the following list of conclusions and recommendations.

1. WEP is not useful for anything other than protecting against casual traffic capture attacks. With the total break in August 2001 and the subsequent release of public implementation code, security administrators should assume that WEP on its own offers no confidentiality. Furthermore, 802.11 networks announce themselves to the world. On a recent trip through San Francisco, I configured a laptop to scan for area networks and found half a dozen. I was not making a serious effort to do this, either. My laptop was placed on the front passenger seat of my car, and I was using a PC Card 802.11 interface, which does not have particularly high gain. Had I been serious, I would have used a high-gain antenna to pick up fainter Beacon frames, and I would have mounted the antenna higher up so the radio signals were not blocked by the steel body of the car. Obscurity plus WEP may meet some definition of "wired equivalent" because frames on wired networks may be delivered to a number of users other than the intended recipient. However, defining "wired equivalent" is a semantic argument that is not worth getting bogged down in.
2. Manual key management is a serious problem. Peer-to-peer networking systems all have problems in the area of management scalability, and WEP is no different. Deploying pairwise keys is a huge burden on system administrators and does not add much, if any, security.
3. When a secret is widely shared, it quickly ceases to become a secret. WEP depends on widely sharing a secret key. Users may come and go, and WEP keys must be changed with every departure to ensure the protection provided by WEP.
4. Data that must be kept confidential should use strong cryptographic systems designed from the ground up with security in mind. The obvious choices are IPSec or SSH. The choice can be based on technical evaluation, product availability, user expertise, and nontechnical factors (institutional acceptance, pricing and licensing, and so on).
5. Varying levels of concern are appropriate for different locations. When using 802.11 for LAN extension, greater threats are likely to be found in large offices.
    a. Remote teleworkers should be protected by strong VPN systems such as IPSec. Using 802.11 in remote locations may increase the risk of interception, but any transmissions from a client to a central site should already be protected using a strong VPN system. Attackers may be able to capture packets traveling over a wireless network more easily, but IPSec was designed to operate in an environment where attackers had large amounts of encrypted traffic to analyze.

b. Large offices pose a much greater concern. VPNs to branch offices are typically site-to-site, protecting only from the edge of the branch office to the access link at the headquarters offices. Anything inside the remote office is not protected by IPSec and is vulnerable to sniffing if other measures are not taken.

6. Stopping anything more casual than packet sniffing requires client software that implements strong cryptographic protection. However, it requires extra system integration work and testing.
   a. A higher-security, point-to-point tunneling technology may be all that is required for your organization. Unix systems can run PPP over SSH tunnels, and some IPSec solutions can be used to create point-to-point tunnels across the access point.
   b. IPSec also protects across the LAN, which may be important. It is possible that a determined attacker can obtain access to the wired backbone LAN where traffic is no longer protected by WEP.
7. WEP does not protect users from each other. When all users have the WEP key, any traffic can be decrypted easily. Wireless networks that must protect users from each other should use VPN solutions or applications with strong built-in security.

It is dangerous to assume that protocols such as IPSec and SSH are magic bullets that can solve your security problems. But the bottom line for wireless networks is that you can't count on WEP to provide even minimal security, and using IPSec or SSH to encrypt your traffic goes a long way to improve the situation.

CHAPTER 6

Security, Take 2: 802.1x

If at first you don't succeed, try again.
—Anonymous
(from the motivational poster in breakrooms everywhere)

Security is a common thread linking many of the wireless LAN stories in the news throughout the past year, and several polls have shown that network managers consider security to be a significant obstacle to wider deployment of wireless LANs. Many of the security problems that have prevented stronger acceptance of 802.11 are caused by flaws in the design of WEP. WEP attempts to serve as both an authentication mechanism and a privacy mechanism. I hope Chapter 5 showed that it effectively serves as neither.

To address the shortcomings of WEP for authentication, the industry is working towards solutions based on the 802.1x specification, which is itself based on the IETF's Extensible Authentication Protocol (EAP). EAP was designed with flexibility in mind, and it has been used as the basis for a number of network authentication extensions. (Cisco's lightweight EAP, LEAP, also is based on EAP.)

802.1x is not without its problems, however. A recent research report identified several problems with the specification.* The first major problem is that 802.11 does not provide a way to guarantee the authenticity and integrity of any frames on the wireless network. Frames on wireless networks can easily be tampered with or forged outright, and the protocol does not provide a way to easily stop or even detect such attacks. The second major problem is that 802.1x was designed to allow the network to authenticate the user. Implicit in the protocol design is the assumption that users will connect to only the "right" network. On wireline networks, connecting to the right network can be as simple as following the wires. Access to the wiring helps

* "An Initial Analysis of the 802.1x Standard" by Arunesh Mishra and Bill Arbaugh; available at *http://www.cs.umd.edu/~waa/1x.pdf*.

the users identify the "right" network. On a wireless network, clear physical connections do not exist, so other mechanims must be designed for networks to prove their identity (or, more precisely, the identity of their owners) to users. 802.1x was designed to collect authentication information from users and grant or deny access based on that information. It was not designed to help networks provide credentials to users, so that function is not addressed by the 802.1x. The specter for rogue access points will not be put to rest by 802.1x.

How 802.1x will be applied to wireless networks is a matter for task group I (TGi) of the 802.11 working group. With no standard available, I have elected to describe how 802.1x works on LANs to provide a basic understanding of how the future 802.11i specification is likely to work. Some modifications will undoubtedly be made to adapt 802.1x to the wireless world, but the fundamental ideas will remain the same. Before talking about 802.1x, though, it is best to gain a solid understanding of the protocol that started it all: EAP.

The Extensible Authentication Protocol

802.1x is based on EAP. EAP is formally specified in RFC 2284 and was initially developed for use with PPP. When PPP was first introduced, there were two protocols available to authenticate users, each of which required the use of a PPP protocol number. Authentication is not a "one size fits all" problem, and it was an active area of research at the time. Rather than burn up PPP protocol numbers for authentication protocols that might become obsolete, the IETF standardized EAP. EAP used a single PPP protocol number while supporting a wide variety of authentication mechanisms. EAP is a simple encapsulation that can run over any link layer, but it has been most widely deployed on PPP links. Figure 6-1 shows the basic EAP architecture, which is designed to run over any link layer and use any number of authentication methods.

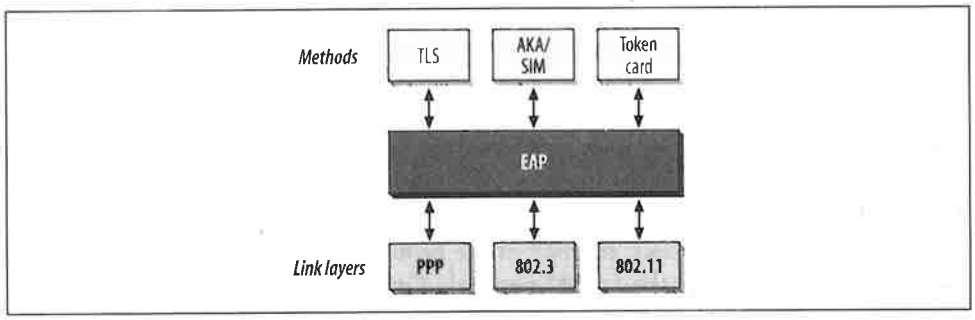

Figure 6-1. EAP architecture

EAP Packet Format

Figure 6-2 shows the format of an EAP packet. When used on PPP links, EAP is carried in PPP frames with a protocol number of 0xC227. There is no strict requirement that EAP run on PPP; the packet shown in Figure 6-2 can be carried in any type of frame. The fields in an EAP packet are:

Code
: The Code field, the first field in the packet, is one byte long and identifies the type of EAP packet. It is used to interpret the Data field of the packet.

Identifier
: The Identifier field is one byte long. It contains an unsigned integer used to match requests with responses to them. Retransmissions reuse the same identifier numbers, but new transmissions use new identifier numbers.

Length
: The Length field is two bytes long. It is the number of bytes in the entire packet, which includes the Code, Identifier, Length, and Data fields. On some link layer protocols, padding may be required. EAP assumes that any data in excess of the Length field is link-layer padding and can be ignored.

Data
: The last field is the variable-length Data field. Depending on the type of packet, the Data field may be zero bytes long. Interpretation of the Data field is based on the value of the Code field.

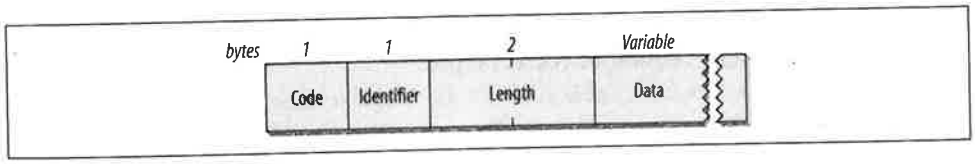

Figure 6-2. EAP packet format

EAP Requests and Responses

EAP exchanges are composed of requests and responses. The authenticator sends requests to the system seeking access, and based on the responses, access may be granted or denied. The format of request and response packets is shown in Figure 6-3.

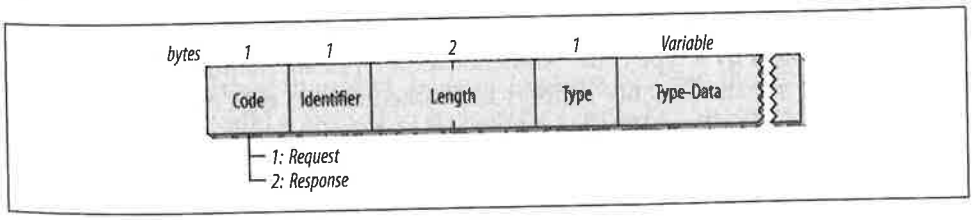

Figure 6-3. EAP Request and EAP Response packets

The Code field is set to 1 for requests and 2 for responses. The Identifier and Length fields are used as described in the previous section on the generic format. The Data field carries the data used in requests and responses. Each Data field carries one type of data, broken down into a type identifier code and the associated data:

Type
: The Type field is a one-byte field that indicates the type of request or response. Only one type is used in each packet. With one exception, the Type field of the response matches the corresponding request. That exception is that when a request is unacceptable, the peer may send a NAK to suggest an alternative type. Types greater than or equal to 4 indicate authentication methods.

Type-Data
: The Type-Data field is a variable field that must be interpreted according to the rules for each type.

Type code 1: Identity

The authenticator generally uses the Identity type as the initial request. After all, identifying the user is the first step in authentication. Naturally, most implementations of EAP prompt the user for input to determine the user identity. The Type-Data field may contain text used to prompt the user; the length of the string is computed from the Length field in the EAP packet itself.

Some EAP implementations may attempt to look up the user identity in a Response even before issuing the authentication challenge. If the user does not exist, the authentication can fail without further processing. Most implementations automatically reissue the identity request to correct typos.

Type Code 2: Notification

The authenticator can use the Notification type to send a message to the user. The user's system can then display the message for the user's benefit. Notification messages are used to provide messages to the user from the authentication system, such as a password about to expire. Responses must be sent in reply to Notification requests. However, they serve as simple acknowledgments, and the Type-Data field has a zero length.

Type code 3: NAK

NAKs are used to suggest a new authentication method. The authenticator issues a challenge, encoded by a type code. Authentication types are numbered 4 and above. If the end user system does not support the authentication type of the challenge, it can issue a NAK. The Type-Data field of a NAK message includes a single byte corresponding to the suggested authentication type.

Type code 4: MD-5 Challenge

The MD-5 Challenge is used to implement the EAP analog of the CHAP protocol, specified in RFC 1994. Requests contain a challenge to the end user. For successful authentication, CHAP requires that the challenge be successfully encoded with a shared secret. All EAP implementations must support the MD-5 Challenge, but they are free to NAK it in favor of another authentication method.

Type code 5: One-time password (OTP)

The one-time password system used by EAP is defined in RFC 1938. The Request issued to the user contains the OTP challenge string. In an OTP (type 5) response, the Type-Data field contains the words from the OTP dictionary in RFC 1938. Like all authentication types, responses may be NAKs (type 3).

Type code 6: Generic Token Card

Token cards such as RSA's SecurID and Secure Computing's Safeword are popular with many institutions because they offer the security of "random" one-time passwords without the hassle of an OTP rollout. The Request contains the Generic Token Card information necessary for authentication. The Type-Data field of the request must be greater than zero bytes in length. In the Response, the Type-Data field is used to carry the information copied from the token card by the user. In both Request and Response packets, the Length field of the EAP packet is used to compute the length of the Type-Data request.

Type code 13: TLS

In its initial form, EAP does not protect transmissions from eavesdropping. In a way, this is an understandable posture, given the origins of the protocol. When EAP is used over dial-up or dedicated links, there is a small chance of interception, but many administrators feel comfortable that the link is reasonably protected against eavesdropping.

For some links, however, assuming the existence of security may not be appropriate. RFC 2716 describes the use of Transport Layer Security (TLS) for authentication. TLS is the standardized successor to the widely deployed Secure Socket Layer (SSL), and TLS authentication inherits a number of useful characteristics from SSL. Most notably, mutual authentication is possible with TLS. Rather than issuing a one-sided challenge to the client ("Who are you?"), EAP-TLS can ensure that the client is communicating with a legitimate authenticator. In addition to mutual authentication, TLS provides a method to protect the authentication between the client and authenticator. It also provides a method to exchange a session key securely between the client and authenticator.

EAP-TLS is likely to become popular on wireless networks. In the current 802.11 authentication regime, access points are implicitly trusted by the clients. EAP-TLS could ensure clients that they are sending sensitive authentication data to legitimate access points instead of "rogue" access points set up by attackers seeking to collect data for a later attack on the network. EAP-TLS also enables the exchange of session keys, which limits the impact of a compromised WEP key.

Additional type codes in draft format

Several additional authentication types are currently in Internet-Draft form and may be standardized after this book is printed. Two of the most notable concepts are Kerberos authentication and cell-phone authentication (SIM cards on second-generation networks and AKA on third-generation networks).

EAP Success and Failure

At the conclusion of an EAP exchange, the user has either authenticated successfully or has failed to authenticate (Figure 6-4). Once the authenticator determines that the exchange is complete, it can issue a Success (code 3) or Failure (code 4) frame to end the EAP exchange. Implementations are allowed to send multiple requests before failing the authentication to allow a user to get the correct authentication data.

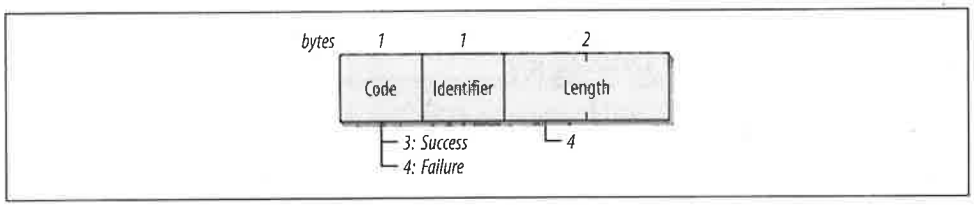

Figure 6-4. EAP Success and Failure frames

A Sample EAP Exchange

A sample EAP exchange is shown in Figure 6-5. It is unnecessarily complex to illustrate several features of the protocol. The EAP exchange is a series of steps beginning with a request for identity and ending with a success or failure message:

1. The authenticator issues a Request/Identity packet to identify the user.
2. The end user system prompts for input, collects the user identifier, and sends the user identifier in a Response/Identity message.
3. With the user identified, the authenticator can issue authentication challenges. In step 3 in the figure, the authenticator issues an MD-5 Challenge to the user with a Request/MD-5 Challenge packet.
4. The user system is configured to use a token card for authentication, so it replies with a Response/NAK, suggesting the use of Generic Token Card authentication.

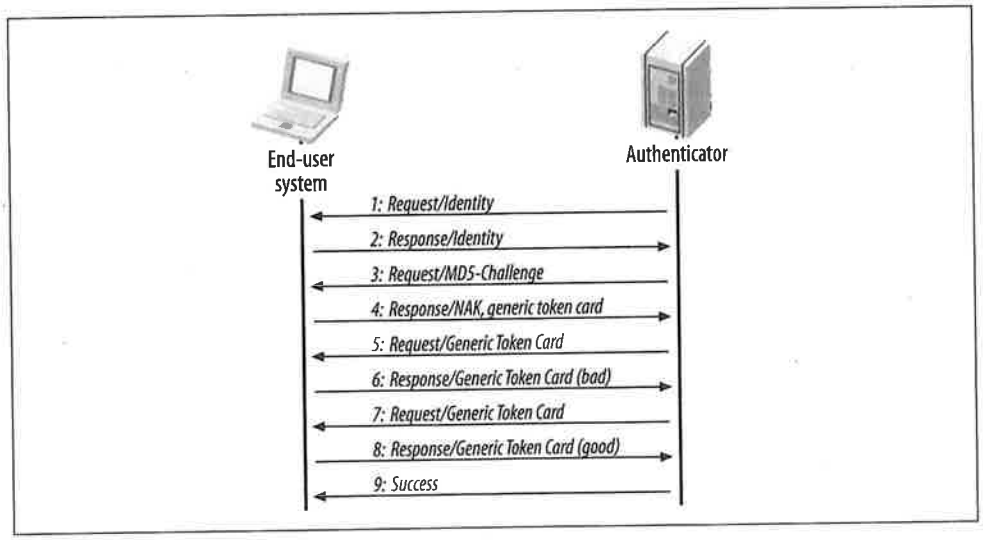

Figure 6-5. Sample EAP exchange

5. The authenticator issues a Request/Generic Token Card challenge, prompting for the numerical sequence on the card.
6. The user types a response, which is passed along in a Response/Generic Token Card.
7. The user response was not correct, so authentication is not possible. However, the authenticator EAP implementation allows for multiple authentication Requests, so a second Request/Generic Token Card is issued.
8. Once again, the user types a response, which is passed along in a Response/Generic Token Card.
9. On the second try, the response is correct, so the authenticator issues a Success message.

802.1x: Network Port Authentication

As LAN acceptance mushroomed in the 1990s, LAN ports popped up everywhere. Some types of organizations, such as universities, were further hampered by a need for openness. Network resources must be made available to a user community, but that community is fluid. Students are not like many network users. They frequently move from computer to computer and do not have a fixed network address; they may also graduate, transfer, enroll, leave campus, work on staff, or undergo any number of changes that may require changes in access privileges. Although network access must be extended to this fluid community, academic budgets are frequently tight, so it is important to prevent unauthorized use by outsiders.

In short, a generic network sign-on was required. Academic environments would not be the sole beneficiaries, however. Authentication to access network resources is common among Internet service providers, and corporations found the idea attractive because of the increasing flexibility of staffing plans.

Authentication to network devices at the link layer is not new. Network port authentication has been required by dial-up access servers for years. Most institutions already have a wide range of deployed infrastructure to support user authentication, such as RADIUS servers and LDAP directories. PPP over Ethernet (PPPoE) could conceivably be used to require user authentication to access an Ethernet, but it would add an unacceptable level of encapsulation overhead and complexity. Instead, the IEEE took the PPP authentication protocols and developed LAN-based versions. The resulting standard was 802.1x, "Port-Based Network Access Control."

802.1x Architecture and Nomenclature

802.1x defines three components to the authentication conversation, which are all shown in Figure 6-6. The *supplicant* is the end user machine that seeks access to network resources. Network access is controlled by the *authenticator*; it serves the same role as the access server in a traditional dial-up network. Both the supplicant and the authenticator are referred to as *Port Authentication Entities* (PAEs) in the specification. The authenticator terminates only the link-layer authentication exchange. It does not maintain any user information. Any incoming requests are passed to an *authentication server*, such as a RADIUS server, for actual processing.

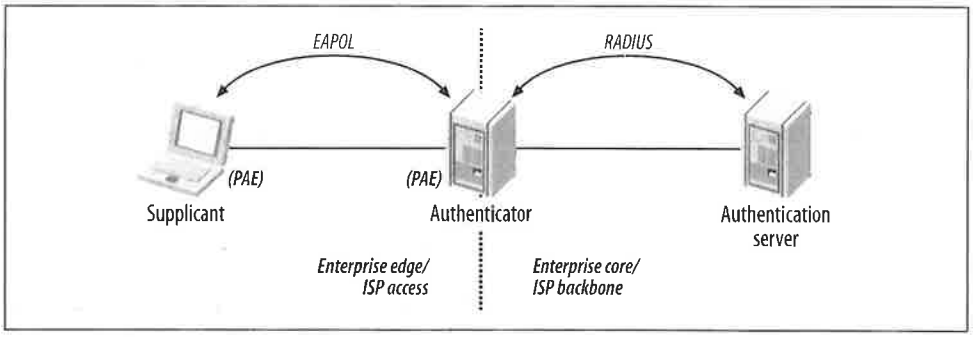

Figure 6-6. 802.1x architecture

Ports on an 802.1x-capable device are in an *authorized* state, in which the port is enabled, or an *unauthorized* state, in which it is disabled. Even while in the unauthorized state, however, the specification allows DHCP and other initialization traffic if permitted by a network manager.

The authentication exchange is logically carried out between the supplicant and the authentication server, with the authenticator acting only as a bridge. A derivation of EAP is used by the authenticator to pass challenges and responses back and forth.

From the supplicant to the authenticator (the "front end"), the protocol is EAP over LANs (EAPOL) or EAP over wireless (EAPOW). On the "back end," the protocol used is RADIUS. Some documentation may refer to it as "EAP over RADIUS." Figure 6-6 can be read as two different scenarios. In the enterprise scenario, the supplicant is a corporate host on the edge of the enterprise network, and the RADIUS server is located in the enterprise core. The figure also depicts an ISP using 802.1x to authenticate users, in which case the lefthand side of the figure is an ISP access area, and the righthand side is the ISP backbone.

802.1x is a framework, not a complete specification in and of itself. The actual authentication mechanism is implemented by the authentication server. 802.1x supplies a mechanism for issuing challenges and confirming or denying access, but it does not pass judgment on the offered credentials. Changes to the authentication method do not require complex changes to the end user devices or the network infrastructure. The authentication server can be reconfigured to "plug in" a new authentication service without changes to the end user driver software or switch firmware.

EAPOL Encapsulation

The basic format of an EAPOL frame is shown in Figure 6-7. EAPOL encapsulation is now analyzed by many popular network analyzers, including Ethereal. The frame's components are:

MAC header
: Figure 6-7 shows the encapsulation on a wired Ethernet, so the MAC header consists of the destination MAC address and the source MAC address. On a wireless network, the MAC header would be the 24- to 30-byte header described in Chapter 3.

Ethernet Type
: As with any other Ethernet frame, the Ethernet Type field contains the two-byte type code assigned to EAPOL: 88-8e.

Version
: At this point, only Version 1 is standardized.

Packet Type
: EAPOL is an extension of EAP. In addition to the EAP messages described in the previous section, EAPOL adds some messages to adapt EAP to the port-based LAN environment. Table 6-1 lists the packet types and their descriptions.

Table 6-1. EAPOL message types

Packet type	Name	Description
0000 0000	EAP-Packet	Contains an encapsulated EAP frame. Most frames are EAP-Packet frames.
0000 0001	EAPOL-Start	Instead of waiting for a challenge from the authenticator, the supplicant can issue an EAPOL-Start frame. In response, the authenticator sends an EAP-Request/Identity frame.

Table 6-1. EAPOL message types (continued)

Packet type	Name	Description
0000 0010	EAPOL-Logoff	When a system is done using the network, it can issue an EAPOL-Logoff frame to return the port to an unauthorized state.
0000 0011	EAPOL-Key	EAPOL can be used to exchange cryptographic keying information.
0000 0100	EAPOL-Encapsulated-ASF-Alert	The Alerting Standards Forum (ASF) has defined a way of allowing alerts, such as SNMP traps, to be sent to an unauthorized port using this frame type.

Packet Body Length
: This two-byte field is the length of the Packet Body field in bytes. It is set to 0 when no packet body is present.

Packet Body
: This variable-length field is present in all EAPOL frames except the EAPOL-Start and EAPOL-Logoff messages. It encapsulates one EAP packet in EAP-Packet frames, one key descriptor in EAPOL-Key frames, and one alert in EAPOL-Encapsulated-ASF-Alert frames.

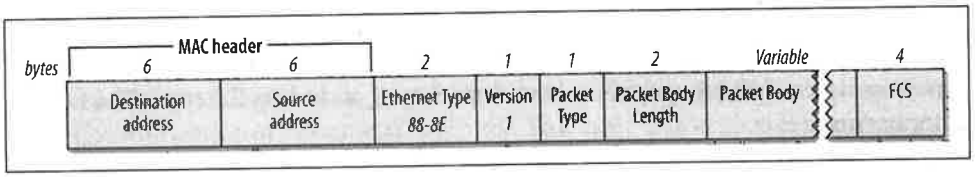

Figure 6-7. EAPOL frame format

Addressing

In shared-media LANs such as Ethernet, supplicants send EAPOL messages to the group address of 01:80:C2:00:00:03. On 802.11 networks, ports do not exist as such, and EAPOL can proceed only after the association process has allowed both the supplicant (mobile wireless station) and the authenticator (access point) to exchange MAC addresses. In environments such as 802.11, EAPOL requests use station addresses.

Sample 802.1x Exchange

EAPOL exchanges look almost exactly like EAP exchanges. The main difference is that supplicants can issue EAPOL-Start frames to trigger the EAP exchange, and they can use EAPOL-Logoff messages to deauthorize the port when the station is done using the network. The examples in this section assume that a RADIUS server is used as the back-end authentication server, and therefore they show the authenticator performing translation from EAP on the front end to RADIUS on the back end. EAP authentication in RADIUS packets is specified in RFC 2869.

The most common case, successful authentication, is shown in Figure 6-8. In the beginning, the port is unauthorized, so access to the network is blocked. The steps in this typical EAPOL exchange are:

1. The supplicant starts the 802.1x exchange with an EAPOL-Start message.
2. The "normal" EAP exchange begins. The authenticator (network switch) issues an EAP-Request/Identity frame.
3. The supplicant replies with an EAP-Response/Identity frame, which is passed on to the RADIUS server as a Radius-Access-Request packet.
4. The RADIUS server replies with a Radius-Access-Challenge packet, which is passed on to the supplicant as an EAP-Request of the appropriate authentication type containing any relevant challenge information.
5. The supplicant gathers the reply from the user and sends an EAP-Response in return. The response is translated by the authenticator into a Radius-Access-Request with the response to the challenge as a data field.
6. The RADIUS server grants access with a Radius-Access-Accept packet, so the authenticator issues an EAP-Success frame. The port is authorized, and the user can begin accessing the network. DHCP configuration may take place at this point.
7. When the supplicant is done accessing the network, it sends an EAPOL-Logoff message to put the port back into an authorized station.

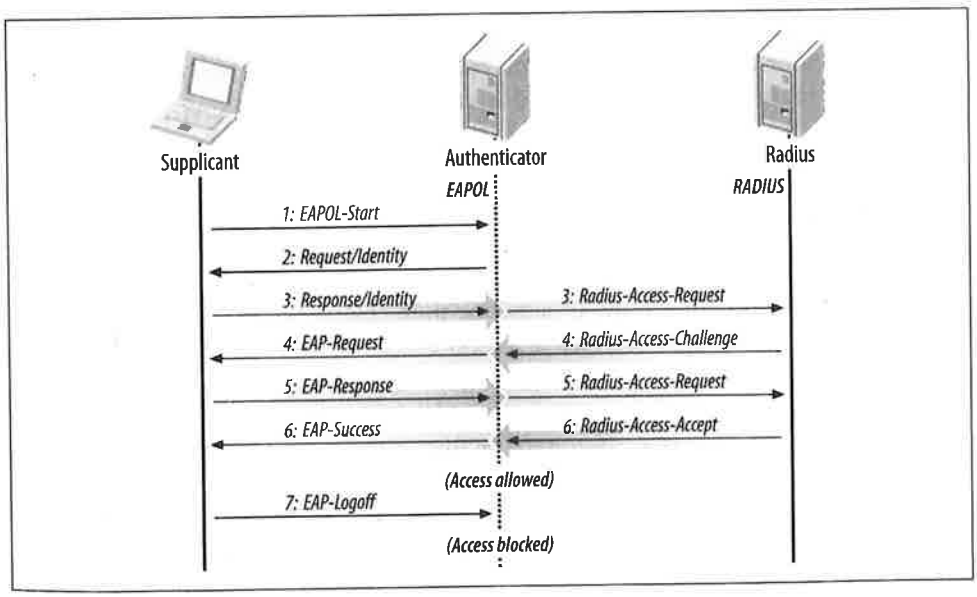

Figure 6-8. Typical EAPOL exchange

Exchanges similar to Figure 6-8 may be used at any point. It is not necessary for the user to begin an EAPOL exchange with the EAPOL-Start message. At any point, the authenticator can begin an EAPOL exchange by issuing an EAP-Request/Identity frame to refresh the authentication data.

802.1x Implementation and the User Experience

802.1x is naturally implemented in the device driver. No change to the hardware itself is necessary. Microsoft has implemented 802.1x as an operating system feature in Windows XP. With no additional software required, XP is likely to be a widely deployed 802.1x client. Cisco has made 802.1x clients available for Windows 9x, NT, 2000, MacOS, and Linux.

EAP is a generic authentication framework, so it does not require any particular authentication method. The first task is to decide on the type of authentication for wireless users. EAP-TLS is attractive because it enables mutual authentication and protects against rogue access points, but it requires that the RADIUS server support EAP-TLS. Use of EAP-TLS also requires that a certificate authority be deployed to manage certificates for wireless network users.

On the client side, authentication is configured in the driver. On XP, 802.1x authentication is configured along with other interface properties. Other implementations allow configuration in the driver. Certificate authentication may take place without any user interaction or may require the user to input a passphrase to unlock the certificate. EAP messages can result in pop ups to the user. Notification messages from the authenticator are immediately displayed, and the identity request allows for a display prompt. If the user is challenged to enter credentials, a pop up is displayed to enter the credentials that will be passed to the authentication server.

802.1x on Wireless LANs

802.1x provides a framework for user authentication over wireless LANs. The only minor change is to define how a "network port" exists within 802.11. The IEEE decided that an association between a station and an access point would be considered a "logical port" for the purpose of interpreting 802.1x. The successful exchange of Association Request and Association Response frames is reported to the 802.1x state engine as the link layer becoming active. 802.11 association must complete before the 802.1x negotiation begins because the 802.1x state machine requires an active link. Prior to a successful 802.1x authentication, the access point drops all non-802.1x traffic. Once the authentication succeeds, the access point removes the filter and allows traffic to flow normally.

The second change made possible by 802.1x is that the EAPOL-Key frame can be used to distribute keying information dynamically for WEP. Figure 6-9 shows a sample EAPOL exchange on an 802.11 network. The only differences from the previous figure

are the requirements of the 802.11 Association Request and Response before beginning the EAPOL exchange and the EAPOL-Key frame that follows the EAP-Success frame. The figure also omits the closing EAPOL-Logoff message.

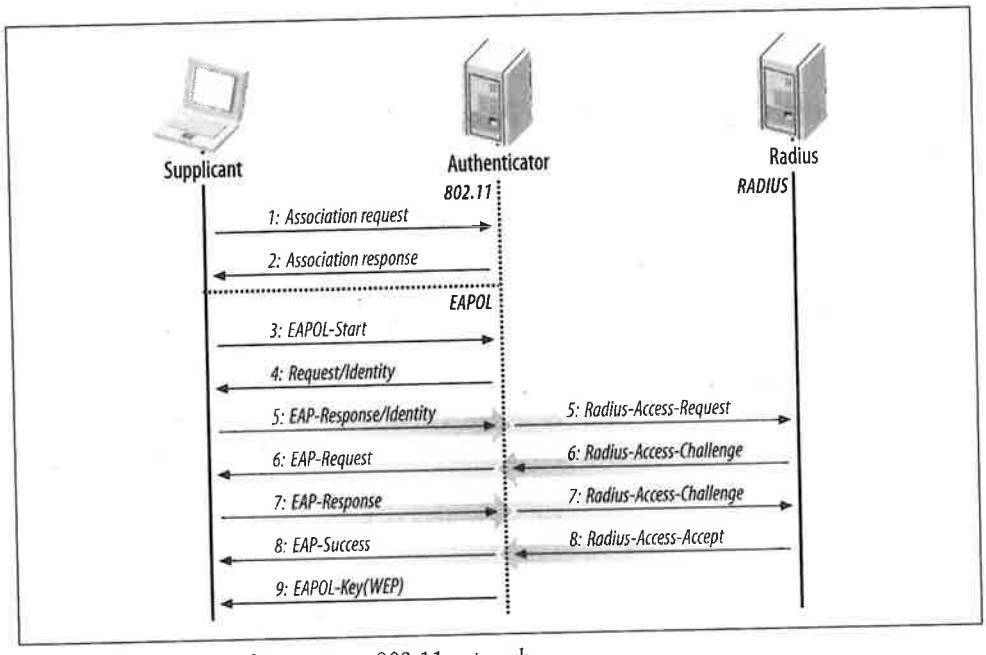

Figure 6-9. EAPOL exchange on an 802.11 network

Association Transfer

802.1x was designed for stations that attach to a port, use it, and then finish. 802.11 poses additional complexities for 802.1x because the association can move from access point to access point. Nothing in 802.1x describes how to move an authentication relationship from one port to another, because 802.1x was designed for a wired world. Transfer of the association and authentication information is an active area of standards development, and it should be addressed in the final 802.11i specification. Several proposals for a "fast handoff" between access points have been submitted, but their development and standardization will require cooperating with the task group developing the Inter-Access Point Protocol.

Keying

The EAPOL-Key frame allows keys to be sent from the access point to the client and vice versa. One commercial implementation uses two WEP keys for each associated station. One key encrypts downstream traffic to the client, and the other encrypts traffic from the client to the access point. Key exchange frames are sent only if the authentication succeeds; this prevents the compromise of key information. EAPOL-Key

frames can be used periodically to update keys dynamically as well. Several of the weaknesses in WEP stem from the long lifetime of the keys. When it is difficult to rekey every station on the network, keys tend to be used for long periods of time. Several experts have recommended changing WEP keys on a regular basis, but no practical mechanism to do so has existed until now, with the development of 802.1x.

Enhancements to 802.11 Made Possible by 802.1x

802.1x brings a number of enhancements to 802.11 networks. Instead of deploying of Mobile IP, it might be possible for user-specific VLANs to be dynamically assigned using the RADIUS tunnel attributes specified in RFC 2868. Dynamic VLANs are clearly not an Internet-scale solution, but they may be a feasible enterprise solution.

Public Ethernet ports are much more attractive in an 802.1x world because they can tie in to a centralized billing infrastructure using RADIUS accounting, defined in RFCs 2866 and 2867 (Figure 6-10). When a user attempts to use a public Ethernet port, he can be authenticated by a corporate RADIUS server, and the company can be billed based on accounting information generated by the service provider's RADIUS server. "Routing" of RADIUS requests to the correct home server can be done if the user identifiers are formatted according to RFC 2486.

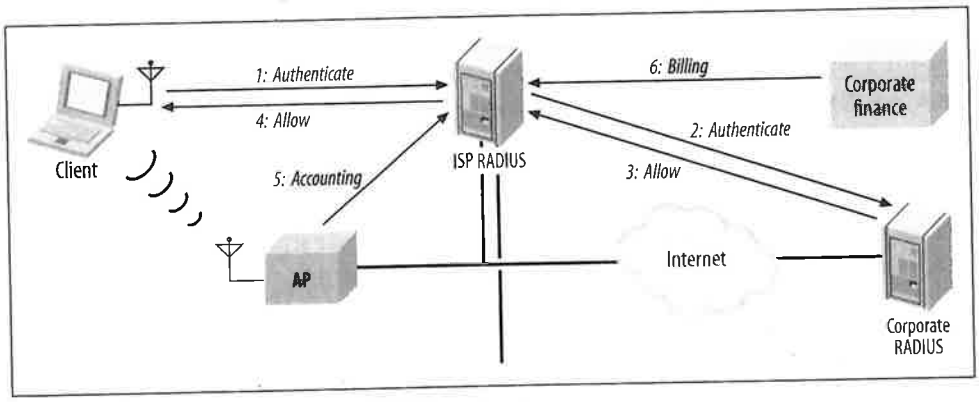

Figure 6-10. 802.11x supporting public Ethernet ports

One further advantage that 802.1x could bring to the build-out of 802.11 networks is a multiprovider access point. The idea is simple. Wireless networks are difficult and costly to build because of the extensive survey work required to avoid sources of interference. If the infrastructure could be built once and shared by a number of service providers, the overall network costs would be much lower.

As an example, airports are a popular wireless LAN "hot spot." Several service providers have attempted to build nationwide networks to allow roaming, and the airlines themselves may want to access an airport-wide wireless LAN. Multiprovider access points could be shared between several users, such as a few different wireless

service providers plus airline kiosks and baggage workers. Multiprovider access points would need to support multiple wireless networks (multiple SSIDs) and multiple VLANs. 802.1x plays a key role in allowing the access point to implement only polices configured on external servers. Airlines are likely to have much more stringent access controls for their internal networks than a wireless service provider.

CHAPTER 7

Management Operations

While being untethered from a wired network can be an advantage, it can lead to problems: the medium is unreliable, unauthorized users can take advantage of the lack of physical boundaries, and power consumption is critical when devices are running on batteries. The management features of the 802.11 protocol were designed to reduce the effect of these problems.

Some device drivers allow you to customize the management features discussed in this chapter. Keep in mind, though, that the capabilities of the device driver vary from one product to another, and the state of wireless networking is such that some vendors are trying to produce the most feature-rich products possible, while others are aiming at the home market and trying to produce the simplest products. The only way to know what's possible is to understand the capabilities that have been built into the protocol. Then you'll be in a good position to work with whatever hardware drops in your lap.

Management Architecture

Conceptually, the 802.11 management architecture is composed of three components: the MAC layer management entity (MLME), a physical-layer management entity (PLME), and a system management entity (SME). The relation between the different management entities and the related parts of 802.11 is shown in Figure 7-1.

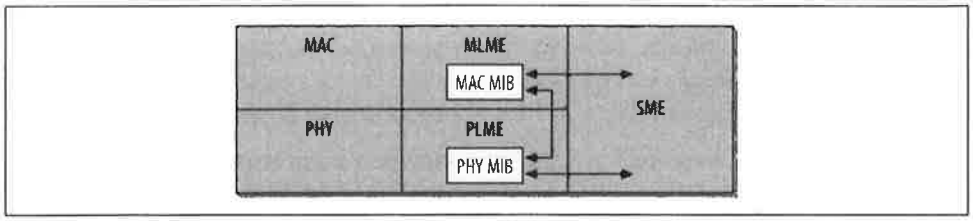

Figure 7-1. Relationship between management entities and components of the 802.11 specification

802.11 does not formally specify the SME. It is the method by which users and device drivers interact with the 802.11 network interface and gather information about its status. Both the MAC and PHY layers have access to a management information base (MIB). The MIB has objects that can be queried to gain status information, as well as objects that can cause certain actions to take place. A full description of the 802.11 MIB is found in Appendix A.

There are three defined interfaces between the management components. The station management entity may alter both the MAC and PHY MIBs through the MLME and PLME service interfaces. Additionally, changes to the MAC may require corresponding changes in the PHY, so an additional interface between the MLME and PLME allows the MAC to make changes to the PHY.

Scanning

Before using any network, you must first find it. With wired networks, finding the network is easy: look for the cable or a jack on the wall. In the wireless world, stations must identify a compatible network before joining it. The process of identifying existing networks in the area is called *scanning*.

Several parameters are used in the scanning procedure. These parameters may be specified by the user; many implementations have default values for these parameters in the driver.

BSSType (independent, infrastructure, or both)
: Scanning can specify whether to seek out independent ad hoc networks, infrastructure networks, or all networks.

BSSID (individual or broadcast)
: The device can scan for a specific network to join (individual) or for any network that is willing to allow it to join (broadcast). When 802.11 devices are moving, setting the BSSID to broadcast is a good idea because the scan results will include all BSSs in the area.

SSID ("network name")
: The SSID assigns a string of bits to an extended service set. Most products refer to the SSID as the network name because the string of bits is commonly set to a human-readable string. Clients wishing to find any network should set this to the broadcast SSID.

ScanType (active or passive)
: Active scanning uses the transmission of Probe Request frames to identify networks in the area. Passive scanning saves battery power by listening for Beacon frames.

ChannelList
: Scans must either transmit a Probe Request or listen on a channel for the existence of a network. 802.11 allows stations to specify a list of channels to try.

Products allow configuration of the channel list in different ways. What exactly constitutes a channel depends on the physical layer in use. With direct-sequence products, it is a list of channels. With frequency-hopping products, it is a hop pattern.

ProbeDelay
: This is the delay, in microseconds, before the procedure to probe a channel in active scanning begins. This delay ensures that an empty or lightly loaded channel does not completely block the scan.

MinChannelTime and MaxChannelTime
: These values, specified in time units (TUs), specify the minimum and maximum amount of time that the scan works with any particular channel.

Passive Scanning

Passive scanning saves battery power because it does not require transmitting. In passive scanning, a station moves to each channel on the channel list and waits for Beacon frames. Any Beacons received are buffered to extract information about the BSS that sent them.

In the passive scanning procedure, the station sweeps from channel to channel and records information from any Beacons it receives. Beacons are designed to allow a station to find out everything it needs to match parameters with the basic service set (BSS) and begin communications. In Figure 7-2, the mobile station uses a passive scan to find BSSs in its area; it hears Beacon frames from the first three access points. If it does not hear Beacons from the fourth access point, it reports that only three BSSs were found.

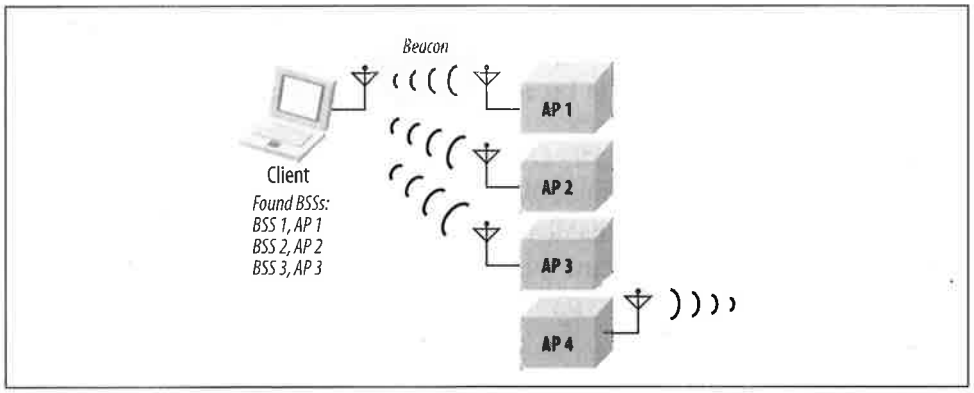

Figure 7-2. Passive scanning

Active Scanning

In active scanning, a station takes a more assertive role. On each channel, Probe Request frames are used to solicit responses from a network with a given name.

Rather than listening for that network to announce itself, an active scan attempts to find the network. Stations using active scanning employ the following procedure for each channel in the channel list:

1. Move to the channel and wait for either an indication of an incoming frame or for the ProbeDelay timer to expire. If an incoming frame is detected, the channel is in use and can be probed. The timer prevents an empty channel from blocking the entire procedure; the station won't wait indefinitely for incoming frames.
2. Gain access to the medium using the basic DCF access procedure and send a Probe Request frame.
3. Wait for the minimum channel time, MinChannelTime, to elapse.
   a. If the medium was never busy, there is no network. Move to the next channel.
   b. If the medium was busy during the MinChannelTime interval, wait until the maximum time, MaxChannelTime, and process any Probe Response frames.

Probe Response frames are generated by networks when they hear a Probe Request that is searching for the extended service set to which the network belongs. At a party, you might look for a friend by wandering around the dance floor shouting out his or her name. (It's not polite, but if you really want to find your friend, you may not have much choice.) If your friend hears you, he or she will respond—others will (you hope) ignore you. Probe Request frames function similarly, but they can also use a broadcast SSID, which triggers a Probe Response from all 802.11 networks in the area. (It's like shouting "Fire!" at the party—that's sure to result in a response from everybody!)

One station in each BSS is responsible for responding to Probe Requests. The station that transmitted the last Beacon frame is also responsible for transmitting any necessary Probe Response frames. In infrastructure networks, the access points transmit Beacons and thus are also responsible for responding to itinerant stations searching the area with Probe Requests. IBSSs may pass around the responsibility of sending Beacon frames, so the station that transmits Probe Response frames may vary. Probe Responses are unicast management frames and are therefore subject to the positive acknowledgment requirement of the MAC.

It is common for multiple Probe Responses to be transmitted as a result of a single Probe Request. The purpose of the scanning procedure is to find every basic service area that the scanning station can join, so a broadcast Probe Request results in a response from every access point within range. Any overlapping independent BSSs may also respond.

Figure 7-3 shows the relationship between the transmission of Probe frames and the various timing intervals that can be configured as part of a scan.

In Figure 7-3a, a mobile station transmits a probe request to which two access points respond. The activity on the medium is shown in Figure 7-3b. The scanning station transmits the Probe Request after gaining access to the medium. Both access points

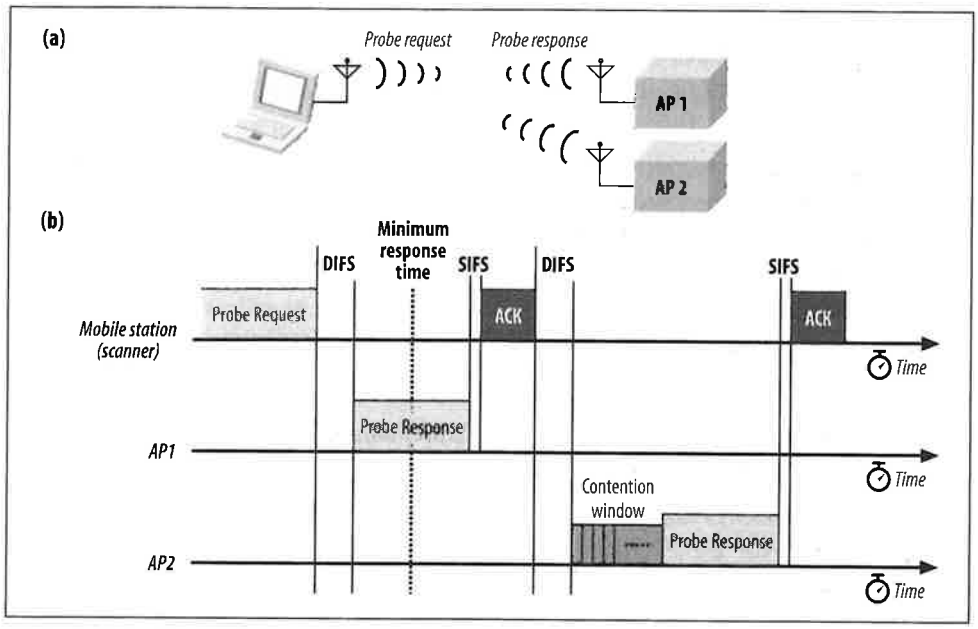

Figure 7-3. Active scanning procedure and medium access

respond with a Probe Response that reports their network's parameters. Note that the second Probe Response is subject to the rules of the distributed coordination function and must wait for the congestion window to elapse before transmitting. The first response is transmitted before the minimum response time elapses, so the station waits until the maximum response time has elapsed before collating the results. In areas with a large number of networks, it may be necessary to adjust the maximum channel time so the responses from all the access points in the area can be processed.

Scan Report

A scan report is generated at the conclusion of a scan. The report lists all the BSSs that the scan discovered and their parameters. The complete parameter list enables the scanning station to join any of the networks that it discovered. In addition to the BSSID, SSID, and BSSType, the parameters also include:*

Beacon interval (integer)
: Each BSS can transmit Beacon frames at its own specific interval, measured in TUs.

DTIM period (integer)
: DTIM frames are used as part of the power-saving mechanism.

* The items actually exposed by any particular software vary.

Timing parameters
: Two fields assist in synchronizing the station's timer to the timer used by a BSS. The Timestamp field indicates the value of the timer received by the scanning station; the other field is an offset to enable a station to match timing information to join a particular BSS.

PHY parameters, CF parameters, and IBSS parameters
: These three facets of the network have their own parameter sets, each of which was discussed in detail in Chapter 4. Channel information is included in the physical-layer parameters.

BSSBasicRateSet
: The basic rate set is the list of data rates that must be supported by any station wishing to join the network. Stations must be able to receive data at all the rates listed in the set. The basic rate set is composed of the mandatory rates in the Supported Rates information element of management frames, as presented in Chapter 4.

Joining

After compiling the scan results, a station can elect to *join* one of the BSSs. Joining is a precursor to association; it is analogous to aiming a weapon. It does not enable network access. Before this can happen, both authentication and association are required.

Choosing which BSS to join is an implementation-specific decision and may even involve user intervention. BSSs that are part of the same ESS are allowed to make the decision in any way they choose; common criteria used in the decision are power level and signal strength. Observers cannot tell when a station has joined a network because the joining process is internal to a node; it involves matching local parameters to the parameters required by the selected BSS. One of the most important tasks is to synchronize timing information between the mobile station and the rest of the network, a process discussed in much more detail in the section "Timer Synchronization."

The station must also match the PHY parameters, which guarantees that any transmissions with the BSS are on the right channel. (Timer synchronization also guarantees that frequency-hopping stations hop at the correct time, too.) Using the BSSID ensures that transmissions are directed to the correct set of stations and ignored by stations in another BSS.* Capability information is also taken from the scan result, which matches the use of WEP and any high-rate capabilities. Stations must also adopt the Beacon interval and DTIM period of the BSS, though these parameters are not as important as the others for enabling communication.

* Technically, this is true only for stations obeying the filtering rules for received frames. Malicious attackers intent on compromising network security can easily choose to disobey these rules and capture frames.

Authentication

On a wired network, authentication is implicitly provided by physical access; if you're close enough to the network to plug in a cable, you must have gotten by the receptionist at the front door. While this is a weak definition of authentication, and one that is clearly inappropriate for high-security environments, it works reasonably well as long as the physical access control procedures are strong. Wireless networks are attractive in large part because physical access is not required to use network resources. Therefore, a major component of maintaining network security is ensuring that stations attempting to associate with the network are allowed to do so. Two major approaches are specified by 802.11: *open-system* authentication and *shared-key* authentication. Shared-key authentication is based on WEP and requires that both stations implement WEP.

802.11 does not restrict authentication to any particular scenario. Any station can authenticate with any other station. In practice, authentication is most useful in infrastructure networks. The usefulness of authentication for infrastructure networks is due in part to the design of the authentication methods, which do not really result in mutual authentication. As a matter of design, the authentication process really only proves the identity of one station. 802.11 implicitly assumes that access points are in a privileged position by virtue of the fact that they are typically under control of network administrators. Network administrators may wish to authenticate mobile stations to ensure that only authorized users access the 802.11 network, but mobile stations can't authenticate the access point. For this reason, the examples in this section assume that a mobile station such as an 802.11-equipped PC is attempting to authenticate to an access point. The standard, however, does not restrict authentication to infrastructure networks.

802.11 authentication is currently a one-way street. Stations wishing to join a network must authenticate to it, but networks are under no obligation to authenticate themselves to a station. The designers of 802.11 probably felt that access points were part of the network infrastructure and thus in a more privileged position, but this curious omission makes a man-in-the-middle attack possible. A rogue access point could certainly send Beacon frames for a network it is not a part of and attempt to steal authentication credentials.

Open-System Authentication

Open-system authentication is the only method required by 802.11. Calling it authentication is stretching the meaning of the term a great deal. In open-system authentication, the access point accepts the mobile station at face value without verifying its identity. (Imagine a world where similar authentication applied to bank withdrawals!) An open-system authentication exchange consists of two frames, shown in Figure 7-4.

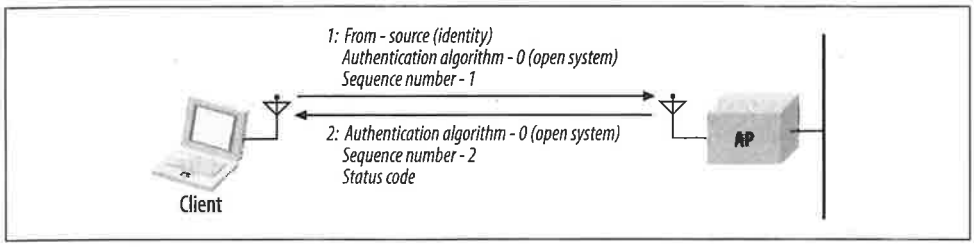

Figure 7-4. Open-system authentication exchange

The first frame from the mobile station is a management frame of subtype authentication. 802.11 does not formally refer to this frame as an authentication request, but that is its practical purpose. In 802.11, the identity of any station is its MAC address. Like Ethernet networks, MAC addresses must be unique throughout the network and can readily double as station identifiers. Access points use the source address of frames as the identity of the sender; no fields within the frame are used to further identify the sender.

There are two information elements in the body of the authentication request. First, the Authentication Algorithm Identification is set to 0 to indicate that the open-system method is in use. Second, the Authentication Transaction Sequence number is set to 1 to indicate that the first frame is in fact the first frame in the sequence.

The access point then processes the authentication request and returns its response. Like the first frame, the response frame is a management frame of subtype authentication. Three information elements are present: the Authentication Algorithm Identification field is set to 0 to indicate open-system authentication, the Sequence Number is 2, and a Status Code indicates the outcome of the authentication request. Values for the Status Code are shown in Table 4-6.

Shared-Key Authentication

Shared-key authentication makes use of WEP and therefore can be used only on products that implement WEP. Furthermore, 802.11 requires that any stations implementing WEP also implement shared-key authentication. Shared-key authentication, as its name implies, requires that a shared key be distributed to stations before attempting authentication. A shared-key authentication exchange consists of four management frames of subtype authentication, shown in Figure 7-5.

The first frame is nearly identical to the first frame in the open-system authentication exchange. Like the open-system frame, it has information elements to identify the authentication algorithm and the sequence number; the Authentication Algorithm Identification is set to 1 to indicate shared-key authentication.

Instead of blindly allowing admission to the network, the second frame in a shared-key exchange serves as a challenge. Up to four information elements may be present

Address Filtering

WEP is not required by 802.11, and a number of earlier products implement only open-system authentication. To provide more security than straight open-system authentication allows, many products offer an "authorized MAC address list." Network administrators can enter a list of authorized client addresses, and only clients with those addresses are allowed to connect.

While address filtering is better than nothing, it leaves a great deal to be desired. MAC addresses are generally software- or firmware-programmable and can easily be overridden by an attacker wishing to gain network access. Furthermore, distributing lists of allowed addresses to multiple access points is a painful process. Some access points implement trivial file transfer protocol (TFTP) servers that allow administrators to push out the address lists, but TFTP is fraught with its own security perils.

Authorized address filtering may be part of a security solution, but it should not be the linchpin. Shared-key authentication is currently the strongest standardized solution available. Once network administrators have made the effort to distribute the WEP keys, authentication will be as secure as standards provide for, and address filtering will only add complexity without significant additional security benefits.

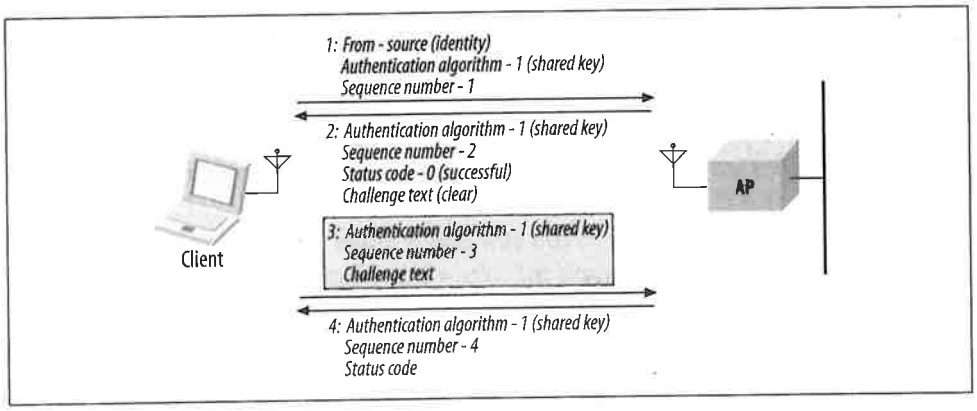

Figure 7-5. Shared-key authentication exchange

in the second frame. Naturally, the Authentication Algorithm Identification, Sequence Number, and Status Code are present. The access point may deny an authentication request in the second frame, ending the transaction. To proceed, however, the Status Code should be set to 0 (success), as shown in Figure 7-5. When the Status Code is successful, the frame also includes a fourth information element, the Challenge Text. The Challenge Text is composed of 128 bytes generated using the WEP keystream generator with a random key and initialization vector.

The third frame is the mobile station's response to the challenge. To prove that it is allowed on the network, the mobile station constructs a management frame with

three information elements: the Authntication Algorithm Identifier, a Sequence Number of 3, and the Challenge Text. Before transmitting the frame, the mobile station processes the frame with WEP. The header identifying the frame as an authentication frame is preserved, but the information elements are hidden by WEP.

After receiving the third frame, the access point attempts to decrypt it and verify the WEP integrity check. If the frame decrypts to the Challenge Text, and the integrity check is verified, the access point will respond with a status code of successful. Successful decryption of the challenge text proves that the mobile station has been configured with the WEP key for the network and should be granted access. If any problems occur, the access point returns an unsuccessful status code.

Preauthentication

Stations must authenticate with an access point before associating with it, but nothing in 802.11 requires that authentication take place immediately before association. Stations can authenticate with several access points during the scanning process so that when association is required, the station is already authenticated. This is called preauthentication. As a result of preauthentication, stations can reassociate with access points immediately upon moving into their coverage area, rather than having to wait for the authentication exchange.

In both parts of Figure 7-6, there is an extended service set composed of two access points. Only one mobile station is shown for simplicity. Assume the mobile station starts off associated with AP1 at the left side of the diagram because it was powered on in AP1's coverage area. As the mobile station moves towards the right, it must eventually associate with AP2 as it leaves AP1's coverage area.

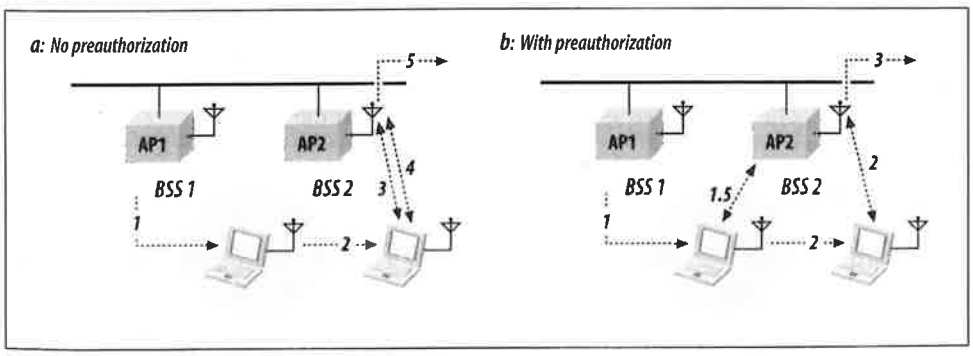

Figure 7-6. Time savings of preauthentication

Preauthentication is not used in the most literal interpretation of 802.11, shown in Figure 7-6a. As the mobile station moves to the right, the signal from AP1 weakens. The station continues monitoring Beacon frames corresponding to its ESS, and will eventually note the existence of AP2. At some point, the station may choose to disassociate

from AP1, and then authenticate and reassociate with AP2. These steps are identified in the figure, in which the numbers are the time values from Table 7-1.

Table 7-1. Chronology for Figure 7-6

Step	Action without preauthentication (Figure 7-6a)	Action with preauthentication (Figure 7-6b)
0	Station is associated with AP1	Station is associated with AP1
1	Station moves right into the overlap between BSS1 and BSS2	Station moves right into the overlap between BSS1 and BSS2 and detects the presence of AP2
1.5		Station preauthenticates to AP2
2	AP2's signal is stronger, so station decides to move association to AP2	AP2's signal is stronger, so station decides to move association to AP2
3	Station authenticates to AP2	Station begins using the network
4	Station reassociates with AP2	
5	Station begins using the network	

Figure 7-6b shows what happens when the station is capable of preauthentication. With this minor software modification, the station can authenticate to AP2 as soon as it is detected. As the station is leaving AP1's coverage area, it is authenticated with both AP1 and AP2. The time savings become apparent when the station leaves the coverage area of AP1: it can immediately reassociate with AP2 because it is already authenticated. Preauthentication makes roaming a smoother operation because authentication can take place before it is needed to support an association. All the steps in Figure 7-6b are identified by time values from Table 7-1.Proprietary Authentication Approaches

The shared-key authentication method has its drawbacks. It is stronger than open-system authentication with address filtering, but it inherits all of WEP's security weaknesses. In response, some vendors have developed proprietary public-key authentication algorithms, many of which are based on 802.1x. Some of these proprietary approaches may serve as the basis for future standards work.

Association

Once authentication has completed, stations can associate with an access point (or reassociate with a new access point) to gain full access to the network. Association is a recordkeeping procedure that allows the distribution system to track the location of each mobile station, so frames destined for the mobile station can be forwarded to the correct access point. After association completes, an access point must register the mobile station on the network so frames for the mobile station are delivered to the access point. One method of registering is to send a gratuitous ARP so the station's MAC address is associated with the switch port connected to the access point.

Association is restricted to infrastructure networks and is logically equivalent to plugging into a wired network. Once the procedure is complete, a wireless station can use the distribution system to reach out to the world, and the world can respond through the distribution system. 802.11 explicitly forbids associating with more than one access point.

Association Procedure

The basic association procedure is shown in Figure 7-7.

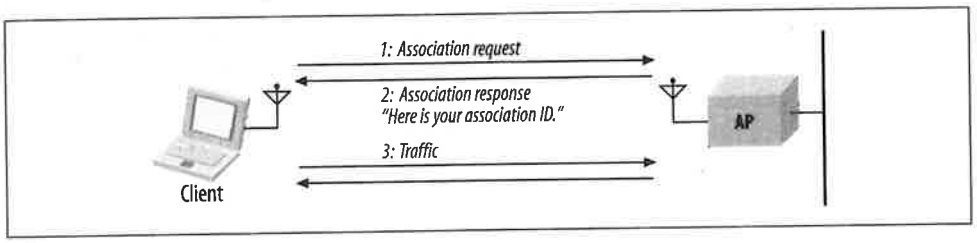

Figure 7-7. Association procedure

Like authentication, association is initiated by the mobile station. No sequence numbers are needed because the association process is a three-step exchange. The two frames are management frame subtypes defined by the specification. As unicast management frames, both steps in the association procedure are composed of an association frame and the required link-layer acknowledgment:

1. Once a mobile station has authenticated to an access point, it can issue an Association Request frame. Stations that have not yet authenticated receive a Deauthentication frame from the access point in response.
2. The access point then processes the association request. 802.11 does not specify how to determine whether an association should be granted; it is specific to the access point implementation. One common consideration is the amount of space required for frame buffering. Rough estimates are possible based on the Listen Interval in the Association Request frame.
   a. When the association request is granted, the access point responds with a status code of 0 (successful) and the Association ID (AID). The AID is a numerical identifier used to logically identify the mobile station to which buffered frames need to be delivered. More detail on the process can be found in the "Power Conservation" section of this chapter.
   b. Unsuccessful association requests include only a status code, and the procedure ends.
3. The access point begins processing frames for the mobile station. In all commonly used products, the distribution system medium is Ethernet. When an access point receives a frame destined for an associated mobile station, that

frame can be bridged from the Ethernet to the wireless medium or buffered if the mobile station is in a power-saving state. In shared Ethernets, the frame will be sent to all the access points and will be bridged by the correct one. In switched Ethernets, the station's MAC address will be associated with a particular switch port. That switch port is, of course, connected to the access point currently providing service for the station.

Reassociation Procedure

Reassociation is the process of moving an association from an old access point to a new one. Over the air, it is almost the same as an association; on the backbone network, however, access points may interact with each other to move frames. When a station moves from the coverage area of one access point to another, it uses the reassociation process to inform the 802.11 network of its new location. The procedure is shown in Figure 7-8.

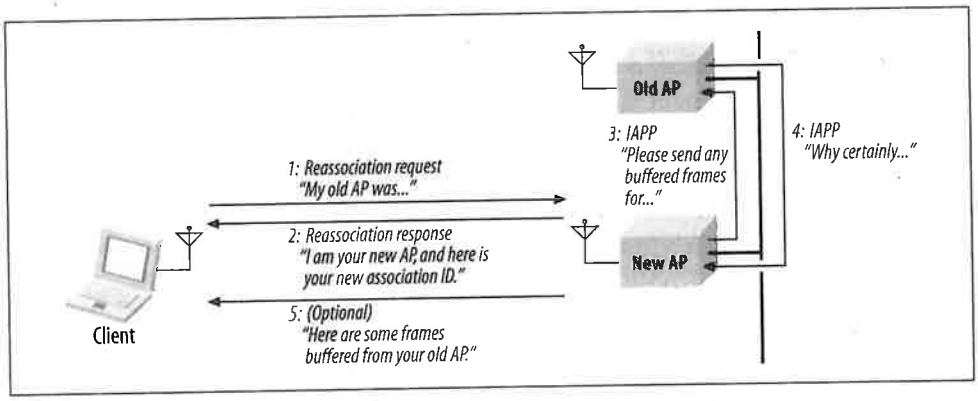

Figure 7-8. Reassociation procedure

The mobile station begins the procedure associated with an access point. The station monitors the quality of the signal it receives from that access point, as well as the signal quality from other access points in the same ESS. When the mobile station detects that another access point would be a better choice, it initiates the reassociation procedure. The factors used to make that decision are product-dependent. Received signal strength can be used on a frame-by-frame basis, and the constant Beacon transmissions provide a good baseline for signal strength from an access point. Before the first step, the mobile station must authenticate to the new access point if it has not done so already.

Figure 7-8 depicts the following steps:

1. The mobile station issues a Reassociation Request to the new access point. Reassociation Requests have content similar to Association Requests. The only difference is that Reassociation Request frames contain a field with the address of the

old access point. The new access point must communicate with the old access point to determine that a previous association did exist. The content of the inter-access point messages is proprietary, though the 802.11 working group is in the process of standardizing the inter-access point protocol. If the new access point cannot verify that the old access point authenticated the station, the new access point responds with a Deauthentication frame and ends the procedure.

2. The access point processes the Reassociation Request. Processing Reassociation Requests is similar to processing Association Requests; the same factors may be used in deciding whether to allow the reassociation:
   a. If the Reassociation Request is granted, the access point responds with a Status Code of 0 (successful) and the AID.
   b. Unsuccessful Reassociation Requests include just a Status Code, and the procedure ends.
3. The new access point contacts the old access point to finish the reassociation procedure. This communication is part of the IAPP.
4. The old access point sends any buffered frames for the mobile station to the new access point. 802.11 does not specify the communication between access points; filling in this omission is one of the major standardization efforts in the 802.11 working group. At the conclusion of the buffered frame transfer:
   a. Any frames buffered at the old access point are transferred to the new access point so they can be delivered to the mobile station.
   b. The old access point terminates its association with the mobile station. Mobile stations are allowed to associate with only one access point at any given time.
5. The new access point begins processing frames for the mobile station. When it receives a frame destined for the mobile station, that frame is bridged from the Ethernet to the wireless medium or buffered for a mobile station in a power-saving mode.

Reassociation is also used to rejoin a network if the station leaves the coverage area and returns later to the same access point. Figure 7-9 illustrates this scenario.

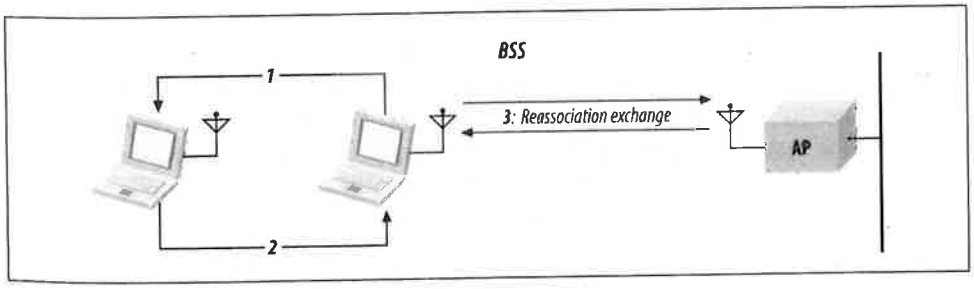

Figure 7-9. Reassociation with the same access point

Power Conservation

The major advantage of wireless networks is that network access does not require nodes to be in any particular location. To take full advantage of mobility, nothing can constrain the location of a node, including the availability of electrical power. Mobility therefore implies that most mobile devices can run on batteries. But battery power is a scarce resource; batteries can run only so long before they need to be recharged. Requiring mobile users to return frequently to commercial power is inconvenient, to say the least. Many wireless applications require long battery life without sacrificing network connectivity.

As with any other network interface, powering down the transceiver can lead to great power savings in wireless networks. When the transceiver is off, it is said to be *sleeping*, *dozing*, or in *power-saving mode* (PS). When the transceiver is on, it is said to be *awake*, *active*, or simply *on*. Power conservation in 802.11 is achieved by minimizing the time spent in the latter stage and maximizing the time in the former. However, 802.11 accomplishes this without sacrificing connectivity.

Power Management in Infrastructure Networks

Power management can achieve the greatest savings in infrastructure networks. All traffic for mobile stations must go through access points, so they are an ideal location to buffer traffic. There is no need to work on a distributed buffer system that must be implemented on every station; the bulk of the work is left to the access point. By definition, access points are aware of the location of mobile stations, and a mobile station can communicate its power management state to its access point. Furthermore, access points must remain active at all times; it is assumed that they have access to continuous power. Combining these two facts allows access points to play a key role in power management on infrastructure networks.

Access points have two power management–related tasks. First, because an access point knows the power management state of every station that has associated with it, it can determine whether a frame should be delivered to the wireless network because the station is active or buffered because the station is asleep. But buffering frames alone does not enable mobile stations to pick up the data waiting for them. An access point's second task is to announce periodically which stations have frames waiting for them. The periodic announcement of buffer status also helps to contribute to the power savings in infrastructure networks. Powering up a receiver to listen to the buffer status requires far less power than periodically transmitting polling frames. Stations only need to power up the transmitter to transmit polling frames after being informed that there is a reason to expend the energy.

Power management is designed around the needs of the battery-powered mobile stations. Mobile stations can sleep for extended periods to avoid using the wireless network interface. Part of the association request is the Listen Interval parameter, which

is the number of Beacon periods for which the mobile station may choose to sleep. Longer listen intervals require more buffer space on the access point; therefore, the Listen Interval is one of the key parameters used in estimating the resources required to support an association. The Listen Interval is a contract with the access point. In agreeing to buffer any frames while the mobile station is sleeping, the access point agrees to wait for at least the listen interval before discarding frames. If a mobile station fails to check for waiting frames after each listen interval, they may be discarded without notification.

Unicast frame buffering and delivery using the Traffic Indication Map (TIM)

When frames are buffered, the destination node's AID provides the logical link between the frame and its destination. Each AID is logically connected to frames buffered for the mobile station that is assigned that AID. Multicast and broadcast frames are buffered and linked to an AID of zero. Delivery of buffered multicast and broadcast frames is treated in the next section.

Buffering is only half the battle. If stations never pick up their buffered frames, saving the frames is a rather pointless exercise. To inform stations that frames are buffered, access points periodically assemble a traffic indication map (TIM) and transmit it in Beacon frames. The TIM is a virtual bitmap composed of 2,008 bits; offsets are used so that the access point needs to transmit only a small portion of the virtual bitmap. This conserves network capacity when only a few stations have buffered data. Each bit in the TIM corresponds to a particular AID; setting the bit indicates that the access point has buffered unicast frames for the station with the AID corresponding to the bit position.

Mobile stations must wake up and enter the active mode to listen for Beacon frames to receive the TIM. By examining the TIM, a station can determine if the access point has buffered traffic on its behalf. To retrieve buffered frames, mobile stations use PS-Poll Control frames. When multiple stations have buffered frames, all stations with buffered data must use the random backoff algorithm before transmitting the PS-Poll.

Each PS-Poll frame is used to retrieve one buffered frame. That frame must be positively acknowledged before it is removed from the buffer. Positive acknowledgment is required to keep a second, retried PS-Poll from acting as an implicit acknowledgment. Figure 7-10 illustrates the process.

If multiple frames are buffered for a mobile station, then the More Data bit in the Frame Control field is set to 1. Mobile stations can then issue additional PS-Poll requests to the access point until the More Data bit is set to 0, though no time constraint is imposed by the standard.

After transmitting the PS-Poll, a mobile station must remain awake until either the polling transaction has concluded or the bit corresponding to its AID is no longer set in the TIM. The reason for the first case is obvious: the mobile station has successfully

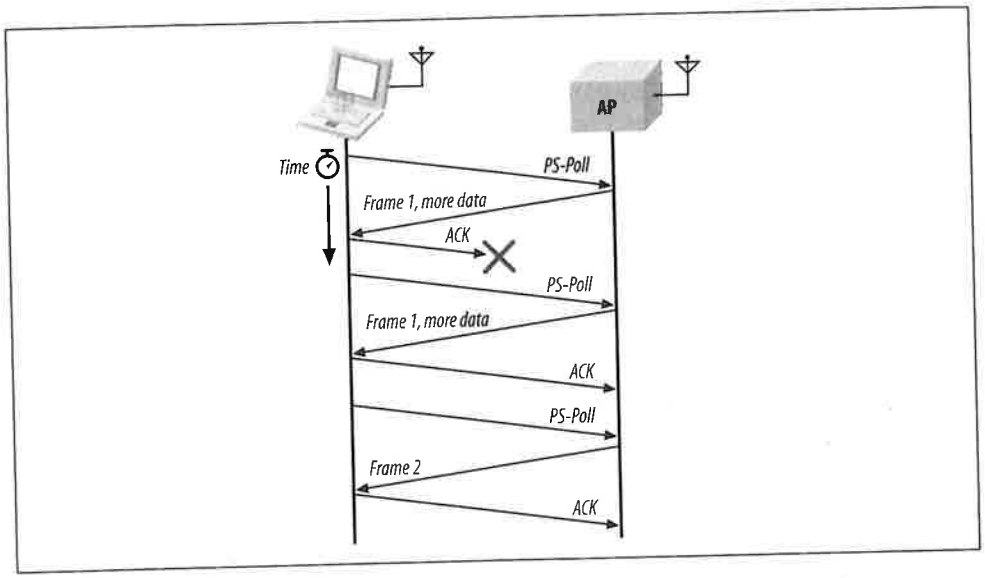

Figure 7-10. PS-Poll frame retrieval

polled the access point; part of that transaction was a notification that the mobile station will be returning to a sleeping mode. The second case allows the mobile station to return to a power conservation mode if the access point discards the buffered frame. Once all the traffic buffered for a station is delivered or discarded, the station can resume sleeping.

The buffering and delivery process is illustrated in Figure 7-11, which shows the medium as it appears to an access point and two associated power-saving stations. The hash marks on the timeline represent the beacon interval. Every beacon interval, the access point transmits a Beacon frame with a TIM information element. (This figure is somewhat simplified. A special kind of TIM is used to deliver multicast traffic; it will be described in the next section.) Station 1 has a listen interval of 2, so it must wake up to receive every other TIM, while station 2 has a listen interval of 3, so it wakes up to process every third TIM. The lines above the station base lines indicate the ramp-up process of the receiver to listen for the TIM.

At the first beacon interval, there are frames buffered for station 1. No frames are buffered for station 2, though, so it can immediately return to sleep. At the second beacon interval, the TIM indicates that there are buffered frames for stations 1 and 2, though only station 1 woke up to listen to the TIM. Station 1 issues a PS-Poll and receives the frame in response. At the conclusion of the exchange, station 1 returns to sleep. Both stations are asleep during the third beacon. At the fourth beacon, both wake up to listen to the TIM, which indicates that there are frames buffered for both. Both station 1 and station 2 prepare to transmit PS-Poll frames after the expiration of a contention window countdown as described in Chapter 3. Station 1 wins because

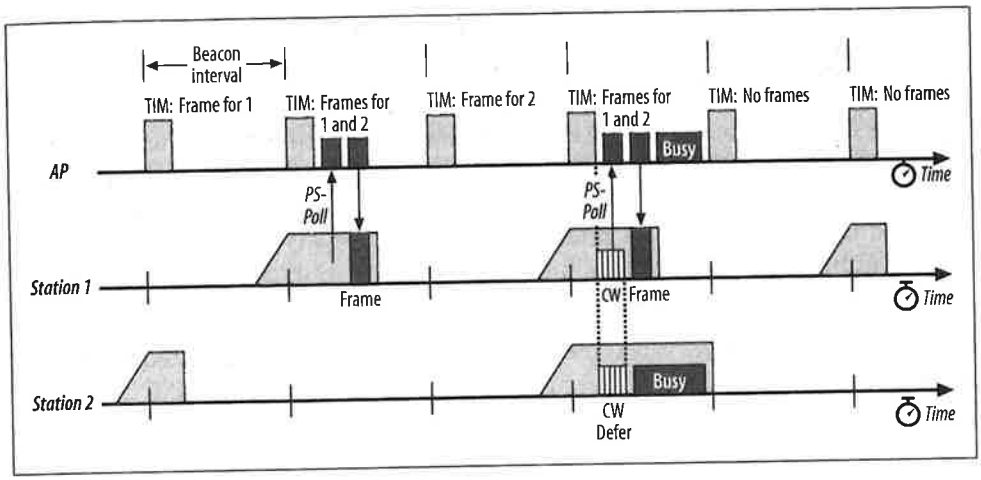

Figure 7-11. Buffered frame retrieval process

its random delay was shorter. Station 1 issues a PS-Poll and receives its buffered frame in response. During the transmission, station 2 defers. If, at the end of that frame transmission, a third station, which is not illustrated, seizes the medium for transmission, station 2 must continue to stay awake until the next TIM. If the access point has run out of buffer space and has discarded the buffered frame for station 2, the TIM at the fifth beacon indicates that no frames are buffered, and station 2 can finally return to a low-power mode.

Stations may switch from a power conservation mode to active mode at any time. It is common for laptop computers to operate with full power to all peripherals when connected to AC power and conserve power only when using the battery. If a mobile station switches to the active mode from a sleeping mode, frames can be transmitted without waiting for a PS-Poll. PS-Poll frames indicate that a power-saving mobile station has temporarily switched to an active mode and is ready to receive a buffered frame. By definition, active stations have transceivers operating continuously. After a switch to active mode, the access point can assume that the receiver is operational, even without receiving explicit notification to that effect.

Access points must retain frames long enough for mobile stations to pick them up, but buffer memory is a finite resource. 802.11 mandates that access points use an *aging function* to determine when buffered frames are old enough to be discarded. The standard leaves a great deal to the discretion of the developer because it specifies only one constraint. Mobile stations depend on access points to buffer traffic for at least the listen interval specified with the association, and the standard forbids the aging function from discarding frames before the listen interval has elapsed. Beyond that, however, there is a great deal of latitude for vendors to develop different buffer management routines.

Delivering multicast and broadcast frames: the Delivery TIM (DTIM)

Frames with a group address cannot be delivered using a polling algorithm because they are, by definition, addressed to a group. Therefore, 802.11 incorporates a mechanism for buffering and delivering broadcast and multicast frames. Buffering is identical to the unicast case, except that frames are buffered whenever any station associated with the access point is sleeping. Buffered broadcast and multicast frames are saved using AID 0. Access points indicate whether any broadcast or multicast frames are buffered by setting the first bit in the TIM to 0; this bit corresponds to AID 0.

Each BSS has a parameter called the DTIM Period. TIMs are transmitted with every Beacon. At a fixed number of Beacon intervals, a special type of TIM, a Delivery Traffic Indication Map (DTIM), is sent. The TIM element in Beacon frames contains a counter that counts down to the next DTIM; this counter is zero in a DTIM frame. Buffered broadcast and multicast traffic is transmitted after a DTIM Beacon. Multiple buffered frames are transmitted in sequence; the More Data bit in the Frame Control field indicates that more frames must be transmitted. Normal channel acquisition rules apply to the transmission of buffered frames. The access point may choose to defer the processing of incoming PS-Poll frames until the frames in the broadcast and multicast transmission buffers have been transmitted.

Figure 7-12 shows an access point and one associated station. The DTIM interval of the access point is set to 3, so every third TIM is a DTIM. Station 1 is operating in a sleep mode with a listen interval of 3. It will wake up on every third beacon to receive buffered broadcast and multicast frames. After a DTIM frame is transmitted, the buffered broadcast and multicast frames are transmitted, followed by any PS-Poll exchanges with associated stations. At the second beacon interval, only broadcast and multicast frames are present in the buffer, and they are transmitted to the BSS. At the fifth beacon interval, a frame has also been buffered for station 1. It can monitor the map in the DTIM and send a PS-Poll after the transmission of buffered broadcast and multicast frames has concluded.

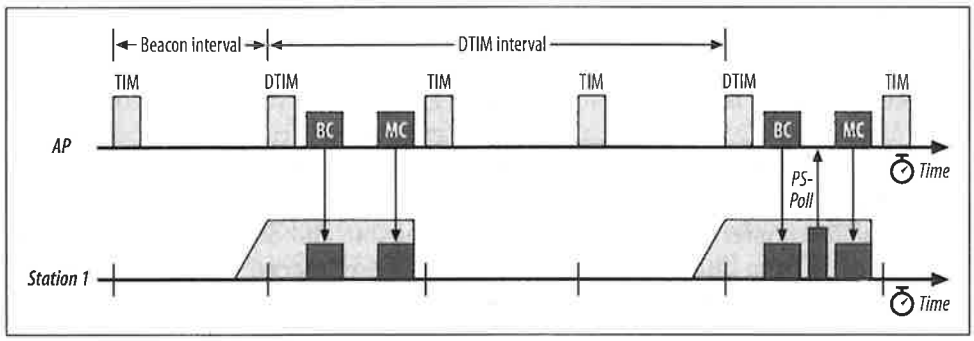

Figure 7-12. Multicast and broadcast buffer transmission after DTIMs

To receive broadcast and multicast frames, a mobile station must be awake for DTIM transmissions. Nothing in the specification, however, keeps power-saving stations in infrastructure networks from waking up to listen to DTIM frames. Some products that implement power-saving modes will attempt to align their awakenings with DTIM transmissions. If the system administrator determines that battery life is more important than receiving broadcast and multicast frames, a station can be configured to sleep for its listen period without regard to DTIM transmissions. Some documentation may refer to this as *extremely low power*, *ultra power-saving mode*, *deep sleep*, or something similar.

Several products allow configuration of the DTIM interval. Lengthening the DTIM interval allows mobile stations to sleep for longer periods and maximizes battery life at the expense of timely delivery. Shorter DTIM intervals emphasize quick delivery at the expense of more frequent power-up and power-down cycles. You can use a longer DTIM when battery life is at a premium and delivery of broadcast and multicast frames is not important. Whether this is appropriate depends on the applications you are using and how they react to long link-layer delays.

IBSS Power Management

Power management in an IBSS is not as efficient as power management in an infrastructure network. In an IBSS, far more of the burden is placed on the sender to ensure that the receiver is active. Receivers must also be more available and cannot sleep for the same lengths of time as in infrastructure networks.

As in infrastructure networks, power management in independent networks is based on traffic indication messages. Independent networks must use a distributed system because there is no logical central coordinator. Stations in an independent network use *announcement traffic indication messages* (ATIMs), which are sometimes called *ad hoc traffic indication messages*, to preempt other stations from sleeping. All stations in an IBSS listen for ATIM frames during specified periods after Beacon transmissions.

If a station has buffered data for another station, it can send an ATIM frame as notification. In effect, the ATIM frame is a message to keep the transceiver on because there is pending data. Stations that do not receive ATIM frames are free to conserve power. In Figure 7-13a, station A has buffered a frame for station C, so it sends a unicast ATIM frame to station C during the ATIM transmission window, which has the effect of notifying station C that it should not enter power-saving mode. Station B, however, is free to power down its wireless interface. Figure 7-13b shows a multicast ATIM frame in use. This frame can be used to notify an entire group of stations to avoid entering low-power modes.

A time window called the *ATIM window* follows the Beacon transmission. This window is the period during which nodes must remain active. No stations are permitted to power down their wireless interfaces during the ATIM window. It starts at the

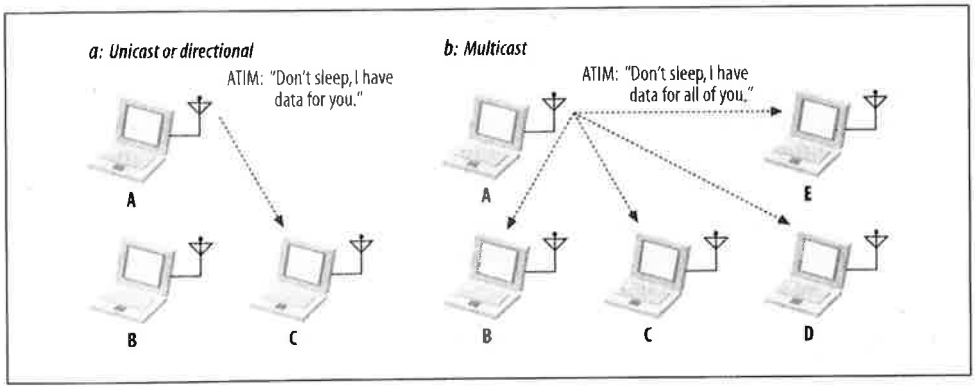

Figure 7-13. ATIM usage

time when the beacon is expected and ends after a period specified when the IBSS is created. If the beacon is delayed due to a traffic overrun, the usable portion of the ATIM window shrinks by the same amount.

The ATIM window is the only IBSS-specific parameter required to create an IBSS. Setting it to 0 avoids using any power management. Figure 7-14 illustrates the ATIM window and its relation to the beacon interval. In the figure, the fourth beacon is delayed due to a busy medium. The ATIM window remains constant, starting at the target beacon interval and extending the length of the ATIM window. Of course, the usable period of the ATIM window shrinks by the length of the delay in beacon transmission.

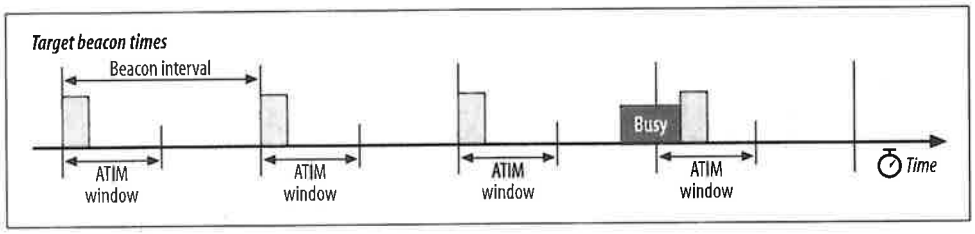

Figure 7-14. ATIM window

To monitor the entire ATIM window, stations must wake up before the target beacon transmission. Four situations are possible: the station has transmitted an ATIM, received an ATIM, neither transmitted nor received, or both transmitted and received. Stations that transmit ATIM frames must not sleep. Transmitting an ATIM indicates an intent to transmit buffered traffic and thus an intent to stay active. Stations to which ATIM frames are addressed must also avoid sleeping so they can receive any frames transmitted by the ATIM's sender. If a station both transmits and receives ATIM frames, it stays up. A station is permitted to sleep only if it neither transmits nor receives an ATIM. When a station stays up due to ATIM traffic, it remains active until the conclusion of the *next* ATIM window, as shown in

Figure 7-15. In the figure, the station goes active for the first ATIM window. If it does not send or receive any ATIM frames, it sleeps at the end of the ATIM window. If it sends or receives an ATIM frame, as in the second ATIM window, the station stays active until the conclusion of the third ATIM window.

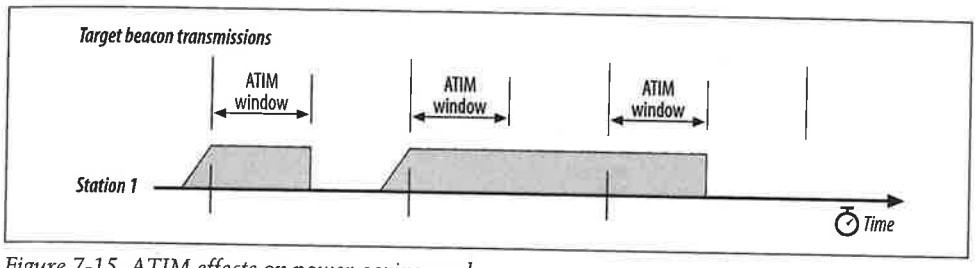

Figure 7-15. ATIM effects on power-saving modes

Only certain control and management frames can be transmitted during the ATIM window: Beacons, RTS, CTS, ACK, and, of course, ATIM frames. Transmission takes place according to the rules of the DCF. ATIM frames may be transmitted only during the ATIM window because stations may be sleeping outside the ATIM window. Sending an ATIM frame is useless if other stations in the IBSS are sleeping. In the same vein, acknowledgments are required for unicast ATIM frames because that is the only guarantee that the ATIM was received and that the frame destination will be active for the remainder of the beacon interval. Acknowledgments are not required for multicast ATIM frames because multicast frames cannot be efficiently acknowledged by a large group of stations. If all potential recipients of an ATIM frame were required to acknowledge it, the mass of acknowledgments could potentially interrupt network service.

Buffered broadcast and multicast frames are transmitted after the conclusion of the ATIM window, subject to DCF constraints. Following the transmission of broadcast and multicast frames, a station may attempt to transmit unicast frames that were announced with an ATIM and for which an acknowledgment was received. Following all transmissions announced with an ATIM, stations may transmit unbuffered frames to other stations that are known to be active. Stations are active if they have transmitted the Beacon, an ATIM, or are not capable of sleeping. If contention is severe enough to prevent a station from sending the buffered frame it announced with an ATIM, the station must reannounce the transmission with an ATIM at the start of the next ATIM window.

Figure 7-16 illustrates several of these rules. In the first beacon interval, the first station transmits a multicast ATIM to stations 2, 3, and 4. Multicast ATIM frames need not be acknowledged, but the transmission of the ATIM means that all stations must remain active for the duration of the first beacon window to receive multicast frames from station 1. When the ATIM window ends, station 1 can transmit its multicast frame to the other three stations. After doing so, station 4 can take advantage of the

remaining time before the beacon to transmit a frame to station 1. It was not cleared with an ATIM, but it is known to be active.

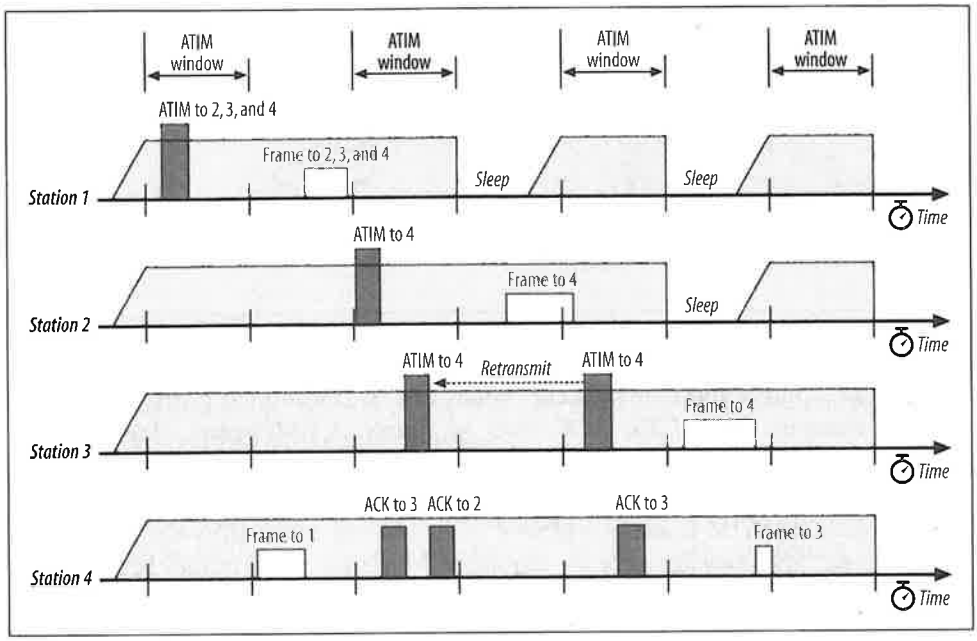

Figure 7-16. Effect of ATIM on power-saving modes in an IBSS network

In the second beacon interval, stations 2 and 3 have both buffered a frame for station 4, so each transmits an ATIM. Station 4 acknowledges both. At the conclusion of the ATIM window, station 1 has neither transmitted nor received an ATIM and can enter a low-power state until the next beacon interval. However, station 2's frame is extremely long and robs station 3 of the opportunity to transmit its frame.

Station 3 still has a buffered frame for station 4 when the third beacon interval opens. It therefore retransmits its ATIM frame to station 4, which is acknowledged. Station 2 is not involved in any ATIM exchanges and can enter a low-power state when the ATIM window ends. At that time, no broadcast or multicast frames have been buffered, and the ATIM-cleared frame from station 3 to station 4 can be transmitted. After the frame from 3 to 4 is transmitted, station 4 can again take advantage of the remaining time before the beacon frame to transmit a frame of its own to station 3, which is known to be active because of the ATIM exchange.

Stations are responsible for maintaining sufficient memory to buffer frames, but the buffer size must be traded off against the use of that memory for other purposes. The standard allows a station in an independent network to discard frames that have been buffered for an "excessive" amount of time, but the algorithm used to make that determination is beyond the scope of the standard. The only requirement placed

on any buffer management function is that it retain frames for at least one beacon period.

Timer Synchronization

Like other wireless network technologies, 802.11 depends a great deal on the distribution of timing information to all the nodes. It is especially important in frequency-hopping networks because all stations on the network must change frequency channels in a coordinated pattern. Timing information is also used by the medium reservation mechanisms.

In addition to local station timing, each station in a basic service area maintains a copy of the *timing synchronization function* (TSF), which is a local timer synchronized with the TSF of every other station in the basic service area. The TSF is based on a 1-MHz clock and "ticks" in microseconds. Beacon frames are used to periodically announce the value of the TSF to other stations in the network. The "now" in a timestamp is when the first bit of the timestamp hits the PHY for transmission.

Infrastructure Timing Synchronization

The ease of power management in an infrastructure network is based on the use of access points as central coordinators for data distribution and power management functions. Timing in infrastructure networks is quite similar. Access points are responsible for maintaining the TSF time, and any stations associated with an access point must simply accept the access point's TSF as valid.

When access points prepare to transmit a Beacon frame, the access point timer is copied into the Beacon's timestamp field. Stations associated with an access point accept the timing value in any received Beacons, but they may add a small offset to the received timing value to account for local processing by the antenna and transceiver. Associated stations maintain local TSF timers so they can miss a Beacon frame and still remain roughly synchronized with the global TSF. The wireless medium is expected to be noisy, and Beacon frames are unacknowledged. Therefore, missing a Beacon here and there is to be expected, and the local TSF timer mitigates against the occasional loss of Beacon frames.

To assist active scanning stations in matching parameters with the BSS, timing values are also distributed in Probe Response frames. When a station finds a network by scanning, it saves the timestamp from the Beacon or Probe Response and the value of the local timer when it was received. To match the local timer to the network timer, a station then takes the timestamp in the received network advertisement and adds the number of microseconds since it was received. Figure 7-17 illustrates this process.

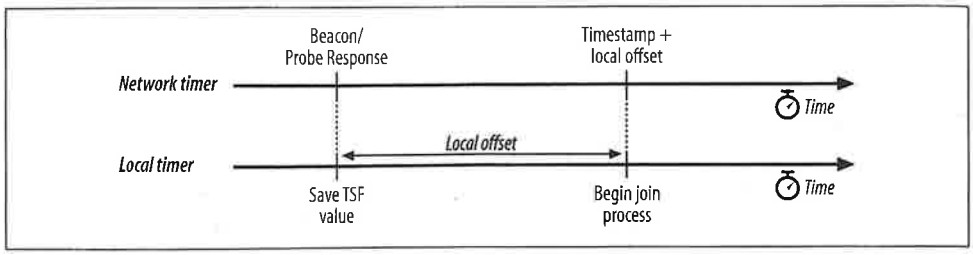

Figure 7-17. Matching the local timer to a network timer

IBSS Timing Synchronization

IBSSs lack a central coordination point, so the Beacon process is distributed. TSF maintenance is a subset of the Beacon generation process. Time is divided into segments equivalent to the interbeacon timing period. Beacon frames are supposed to be transmitted exactly as the beacon interval ends, at the so-called *target Beacon transmission time* (TBTT). Independent networks take the TBTT as a guideline.

All stations in the IBSS prepare to transmit a Beacon frame at the target time. As it approaches, all other traffic is suspended. Timers for the transmission of frames other than Beacon frames or ATIM frames are stopped and held to clear the medium for the important management traffic. All stations in the IBSS generate a *backoff timer* for Beacon transmission; the backoff timer is a random delay between 0 and twice the minimum contention window for the medium. After the target beacon interval, all stations begin to count the Beacon backoff timer down to 0. If a Beacon is received before the station's transmission time, the pending Beacon transmission is canceled.

In Figure 7-18, each station selects a random delay; station 2 has randomly generated the shortest delay. When station 2's timer expires, it transmits a Beacon, which is received by stations 1 and 3. Both stations 1 and 3 cancel their Beacon transmissions as a result. Because timer synchronization ensures that all stations have synchronized timers, multiple Beacon frames do not pose a problem. Receivers simply process multiple Beacon frames and perform multiple updates to the TSF timer.

Beacon generation interacts closely with power management. Beacon frames must be generated during the active period around each Beacon interval so that all stations are available to process the Beacon. Furthermore, the Beacon sender is not allowed to enter a low-power state until the end of the next active period. The latter rule ensures that at least one station is awake and can respond to probes from new stations scanning to discover networks.

Rules for adopting the received timestamp are more complex in an independent network. No centralized timer exists, so the goal of the standard is to synchronize all timers to the timer of the fastest-running clock in the BSS. When a Beacon is received, the timestamp is adjusted for processing delays and compared to the local

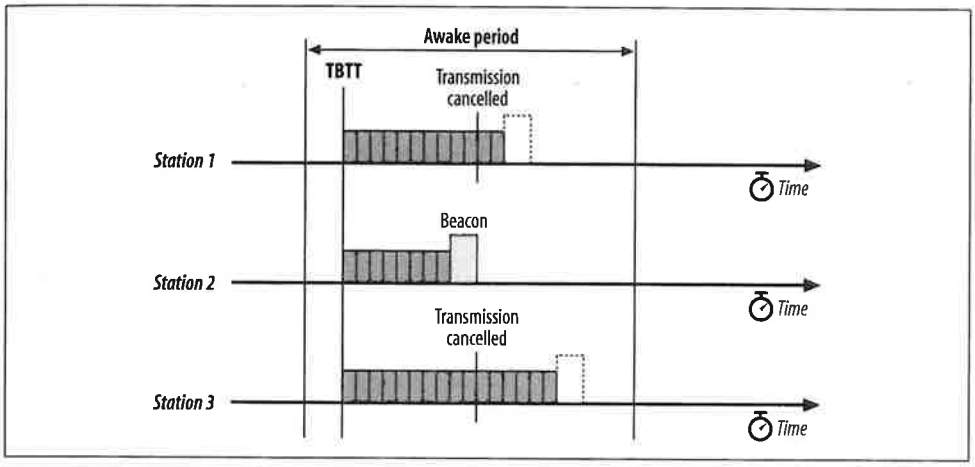

Figure 7-18. Distributed Beacon generation

TSF. The received timestamp updates the local timer only if it is later than the local timer.

CHAPTER 8

Contention-Free Service with the PCF

To support applications that require near real-time service, the 802.11 standard includes a second coordination function to provide a different way of accessing the wireless medium. The point coordination function (PCF) allows an 802.11 network to provide an enforced "fair" access to the medium. In some ways, access to the medium under the PCF resembles token-based medium access control schemes, with the access point holding the token. This chapter describes medium access under the PCF, detailed frame diagrams for the PCF frames, and how power management operations interact with the PCF.

The PCF has not been widely implemented. This chapter is included for two reasons. Readers interested in the standard itself may also be interested in how the PCF works. It is also possible that products based on the PCF may someday hit the market, in which case, network engineers will need to understand the PCF so they can implement it. But most readers can skip this chapter safely.

Contention-Free Access Using the PCF

If contention-free delivery is required, the PCF may be used. The PCF is an optional part of the 802.11 specification; products are not required to implement it. However, the IEEE designed the PCF so stations that implement only the distributed coordination function (DCF) will interoperate with point coordinators.

Contention-free service is not provided full-time. Periods of contention-free service arbitrated by the point coordinator alternate with the standard DCF-based service. The relative size of the contention-free period can be configured. 802.11 describes the contention-free periods as providing "near isochronous" services because the contention-free periods will not always start at the expected time, as described in the section "Contention-Free Period Duration."

Contention-free service uses a centralized access control method. Access to the medium is restricted by the point coordinator, a specialized function implemented in

access points. Associated stations can transmit data only when they are allowed to do so by the point coordinator. In some ways, contention-free access under the PCF resembles token-based networking protocols, with the point coordinator's polling taking the place of a token. Fundamentals of the 802.11 model remain in place, however. Although access is under the control of a central entity, all transmissions must be acknowledged.

PCF Operation

Figure 8-1 shows a transfer using the PCF. When the PCF is used, time on the medium is divided into the contention-free period (CFP) and the contention period. Access to the medium in the former case is controlled by the PCF, while access to the medium in the latter case is controlled by the DCF and the rules from Chapter 7. The contention period must be long enough for the transfer of at least one maximum-size frame and its associated acknowledgment. Alternating periods of contention-free service and contention-based service repeat at regular intervals, which are called the contention-free repetition interval.

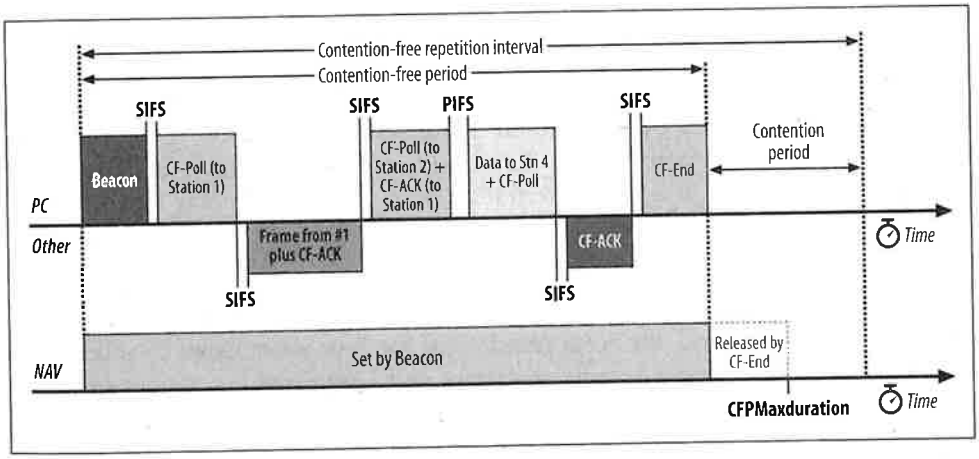

Figure 8-1. Using the PCF

Reserving the medium during the contention-free period

At the beginning of the contention-free period, the access point transmits a Beacon frame. One component of the beacon announcement is the maximum duration of the contention-free period, *CFPMaxDuration*. All stations receiving the Beacon set the NAV to the maximum duration to lock out DCF-based access to the wireless medium.

As an additional safeguard to prevent interference, all contention-free transmissions are separated only by the short interframe space and the PCF interframe space. Both

are shorter than the DCF interframe space, so no DCF-based stations can gain access to the medium using the DCF.

The polling list

After the access point has gained control of the wireless medium, it polls any associated stations on a *polling list* for data transmissions. During the contention-free period, stations may transmit only if the access point solicits the transmission with a polling frame. Contention-free polling frames are often abbreviated CF-Poll. Each CF-Poll is a license to transmit one frame. Multiple frames can be transmitted only if the access point sends multiple poll requests.

The polling list is the list of privileged stations solicited for frames during the contention-free period. Stations get on the polling list when they associate with the access point. The Association Request includes a field that indicates whether the station is capable of responding to polls during the contention-free period.

Transmissions from the Access Point

Generally, all transmissions during the contention-free period are separated by only the short interframe space. To ensure that the point coordinator retains control of the medium, it may send to the next station on its polling list if no response is received after an elapsed PCF interframe space. (Such a situation is illustrated in Figure 8-1.) The access point polled the second station on its list but received no response. After waiting one PCF interframe space, the access point moves to the third station on the list. By using the PCF interframe space, the access point ensures that it retains access to the medium.

The access point may use several different types of frames during the contention-free period. During this period, the point coordinator has four major tasks. In addition to the "normal" tasks of sending buffered frames and acknowledging frames from the stations, the point coordinator can poll stations on the polling list to enable them to send frames; it may also need to transmit management frames.

Time in the contention-free period is precious, so acknowledgments, polling, and data transfer may be combined to improve efficiency. When any subset of these functions are combined into a single frame, the result is a bit strange. A single frame could, for example, acknowledge the receipt of the previous frame, poll a different station for buffered data, and send its own data to the station on the polling list.

Several different frame types can be used in the contention free period:

Data
: The standard vanilla Data frame is used when the access point is sending a frame to a station and does not need to acknowledge a previous transmission. The standard Data frame does not poll the recipient and thus does not allow the recipient

to transmit any data in return. The Data-Only frame used in the contention-free period is identical to the Data frame used in contention-based periods.

CF-Ack

This frame is used by stations to acknowledge the receipt of a frame when no data needs to be transmitted. Contention-free acknowledgments are longer than the standard control frame acknowledgment, so this frame may not be used in actual implementations.

CF-Poll

CF-Poll frames are sent by the access point to a mobile station to give the mobile station the right to transmit a single buffered frame. It is used when the access point does not have any data for the mobile station. When a frame for the mobile station is available, the access point uses the Data+CF-Poll frame type.

Data+CF-Ack

This frame combines data transmission with an acknowledgment. Data is directed to the frame recipient; the acknowledgment is for the previous frame transmitted and usually is not for the recipient of the data.

Data+CF-Poll

This frame is used by access points to transmit data to a mobile station and request one pending frame from the mobile station. The Data+CF-Poll can only be sent by the access point during the contention-free period.

CF-ACK+CF-Poll

This frame acknowledges the last frame from one of the access point's clients and requests a buffered frame from the next station on the polling list. It is directed to the next station on the polling list, though the acknowledgment may be intended for any mobile station associated with the access point.

Data+CF-ACK+CF-Poll

This frame brings together the data transmission, polling feature, and acknowledgment into one frame for maximum efficiency.

CF-End

This frame ends the contention-free period and returns control of the medium to the contention-based mechanisms of the DCF.

CF-End+CF-Ack

This is the same as the CF-End frame but also acknowledges the previously transmitted Data frame.

Any Management

No restriction is placed by the standard on which management frames can be transmitted during the contention-free period. If the rules applying to a particular frame type allow its transmission, the access point may transmit it.

Contention-Free Period Duration

The minimum length of the contention period is the time required to transmit and acknowledge one maximum-size frame. It is possible for contention-based service to overrun the end of the contention period, however. When contention-based service runs past the expected beginning of the contention-free period, the contention-free period is *foreshortened*,as in Figure 8-2.

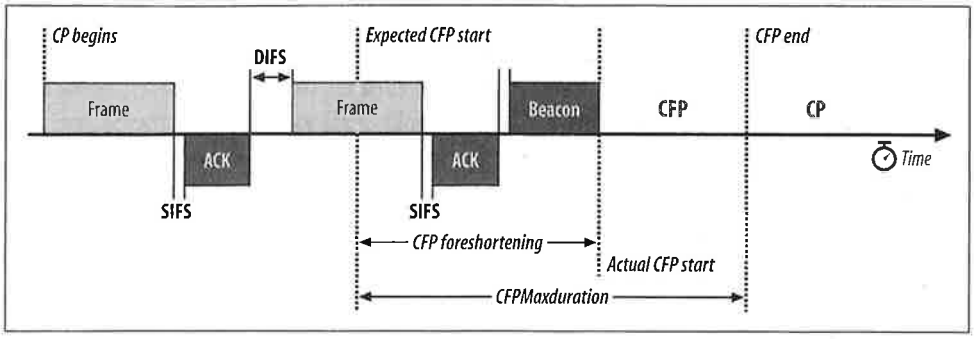

Figure 8-2. Data+CF-Ack and Data+CF-Poll usage

When the contention-free period is foreshortened, the existing frame exchange is allowed to complete before the beacon announcing the start of contention-free operation is transmitted. The contention-free period is shortened by the amount of the delay. Contention-free service ends no later than the maximum duration from the expected beginning point, which is referred to as the Target Beacon Transmission Time (TBTT).

The point coordinator may also terminate the contention-free period prior to its maximum duration by transmitting a CF-End frame. It can base this decision on the size of the polling list, the traffic load, or any other factor that the access point considers important.

Detailed PCF Framing

Several frame types are used exclusively within the contention-free period. They combine, in various states, data transmission, acknowledgment, and polling. This section describes when various frames are used and how the different functions interact during frame exchanges.

Data+CF-Ack

The Data+CF-Ack frame combines two different functions for transmission efficiency. Data is transmitted in the frame payload, and the frame implicitly acknowledges the receipt of data received one short interframe space previously. Generally, the data and the acknowledgment are intended for two separate stations. In Figure 8-3, the contention-free acknowledgment is coupled with the

data for transmission to the access point in the previous frame, but the data may be intended for any station on the 802.11 network.

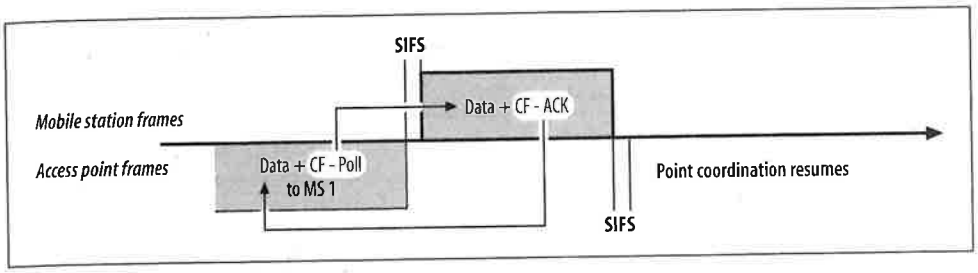

Figure 8-3. Data+CF-Ack usage

This frame is used only in infrastructure networks because it is transmitted during the contention-free period. It may be transmitted by either the access point or a mobile station. During the contention-free period, however, the access point is responsible for polling, and it is unlikely that it would transmit this frame subtype because it does not include a poll.

Data+CF-Poll

The Data+CF-Poll frame is used by access points in infrastructure networks during the contention-free period. When the access point does not need to acknowledge any outstanding frames, it sends a Data+CF-Poll to transmit data to the recipient and allows the recipient to send one buffered frame in response. The data in the frame body must be intended for the recipient of the poll; the two operations cannot be "split" across two different receivers. In Figure 8-3, the access point uses a Data+CF-Poll frame to send one frame to the mobile station and to solicit the response.

Data+CF-Ack+CF-Poll

The Data+CF-Ack+CF-Poll frame is used by access points in infrastructure networks during the contention-free period. When the access point has data to transmit, must acknowledge a frame, and needs to poll a station on the polling list, all the functions can be combined into one frame. Figure 8-4 illustrates the usage of Data+CF-Ack+CF-Poll. As with Data+CF-Ack, the components of the Data+CF-Ack+CF-Poll frame are generally intended for different stations. The data transmission and polling must be intended for the same station, but the acknowledgment is for the previous transmission.

The figure begins with mobile station 1 (MS1) transmitting a Data+CF-Ack frame. The Data must go to the access point, but the CF-Ack is used to acknowledge the previous Data frame transmitted by the access point. (That frame is not shown in the figure.) Moving down the polling list, the access point then polls mobile station 2 (MS2). However, the access point must acknowledge the data from MS1, which it does by transmitting a frame with a CF-Ack component. When the access point also has data to transmit, all three features can be combined into one

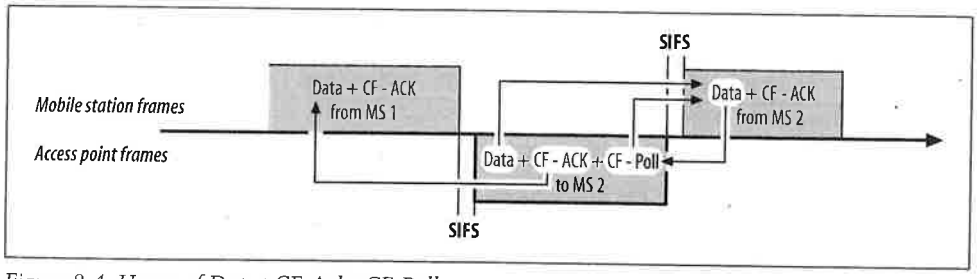

Figure 8-4. Usage of Data+CF-Ack+CF-Poll

omnibus frame. The Data and CF-Poll components are intended for the recipient of the frame, but the CF-Ack is intended for the transmitter of the *previous* frame. MS1 must listen to the access point frames to note the acknowledgment.

CF-Ack (no data)

When only an acknowledgment is required, a header-only frame with just the CF-Ack function can be transmitted. In Figure 8-4, if MS2 had no data to transmit, it would have responded with a CF-Ack frame.

CF-Poll (no data)

CF-Poll can also be transmitted by itself. Naturally, only access points perform this function, so the CF-Poll frame is transmitted only by access points in infrastructure networks during the contention-free period.

"Naked" CF-Polls are transmitted when the access point has no buffered data for the recipient and does not need to acknowledge the receipt of previous frames. One common situation in which no acknowledgment is necessary is when the access point transmits a CF-Poll and the polled station has no data and does not respond. If the access point has no data for the next station on the polling list, it transmits a CF-Poll, as in Figure 8-5.

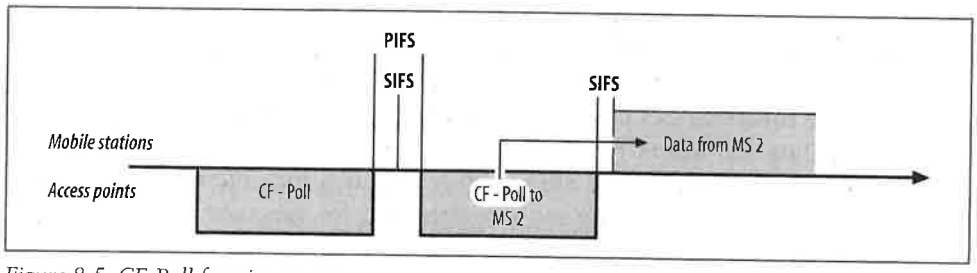

Figure 8-5. CF-Poll framing usage

In Figure 8-5, the access point attempts to transmit data to MS1 but does not receive a response. After the PCF interframe space has elapsed, the access point can proceed down the polling list to MS2. No frame from MS1 needs to be acknowledged, and if the access point has no data for MS2, it can use a CF-Poll to allow MS2 to send data.

CF-Ack+CF-Poll (no data)
: The final subtype of Data frame is the CF-Ack+CF-Poll, which is also transmitted by access points. Like all CF-Poll frames, it is used only during the contention-free period and only by access points. It incorporates the acknowledgment function and the polling function into a frame with no data. Figure 8-6 illustrates its usage.

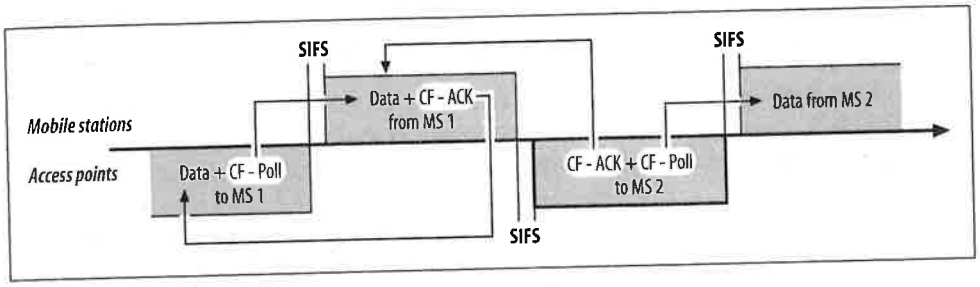

Figure 8-6. CF-Ack+CF-Poll usage

The scenario is a slight variation on the previous setting. Instead of a timeout waiting for MS1 to respond, MS1 returns a frame. When the access point takes control of the medium, it uses a CF-Ack+CF-Poll to acknowledge receipt of the frame from MS1 and notifies MS2 that it is allowed to send a frame.

Contention-Free End (CF-End)

When the contention-free period ends, the access point transmits a CF-End frame to release stations from the PCF access rules and begin contention-based service. The format of the CF-End frame is shown in Figure 8-7. Four fields make up the MAC header of the CF-End frame:

Frame Control
: The frame subtype is set to 1110 to indicate a CF-End frame.

Duration
: CF-End announces the end of the contention-free period and thus does not need to extend the virtual carrier sense. Duration is set to 0. Stations that receive the CF-End frame cut the virtual carrier sense short to resume contention-based access.

Address 1: Receiver Address
: CF-End is relevant to the operation of all mobile stations, so the receiver address is the broadcast address.

Address 2: BSSID
: CF-End is announced by the access point to all the stations associated with its BSS, so the second address field is the BSSID. In infrastructure networks, the BSSID is the address of the wireless interface in the access point, so the BSSID is also the transmitter address.

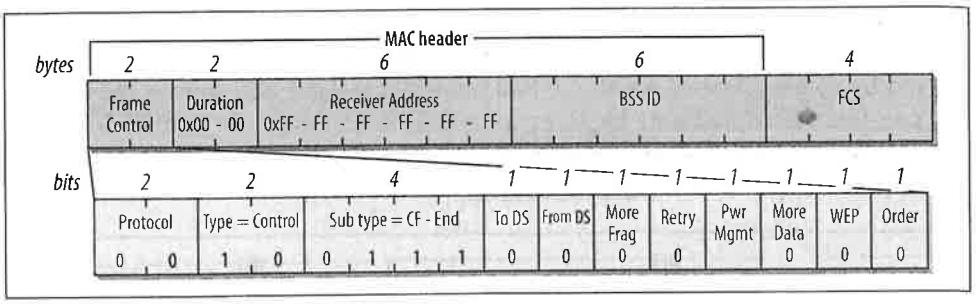

Figure 8-7. CF-End frame

CF-End+CF-Ack

When the contention-free period ends, the access point transmits a CF-End frame to release stations from the PCF access rules and then begins contention-based service using the DCF. If the access point must also acknowledge receipt of data, it may simultaneously end the contention-free period and acknowledge the previous frame by using the CF-End+CF-Ack frame, which combines both functions. The format of the CF-End+CF-Ack frame is shown in Figure 8-8. Four fields make up the MAC header of the CF-End+CF-Ack frame:

Frame Control
: The frame subtype is set to 1111 to indicate a CF-End+CF-Ack frame.

Duration
: CF-End+CF-Ack announces the end of the contention-free period and thus does not need to extend the virtual carrier sense. Duration is set to 0.

Address 1: Receiver Address
: CF-End+CF-Ack is relevant to the operation of all mobile stations, so the receiver address is the broadcast address.

Address 2: BSSID
: CF-End+CF-Ack is announced by the access point to all the stations associated with its BSS, so the second address field is the BSSID. In infrastructure networks, the BSSID is the address of the wireless interface in the access point, so the BSSID is also the transmitter address.

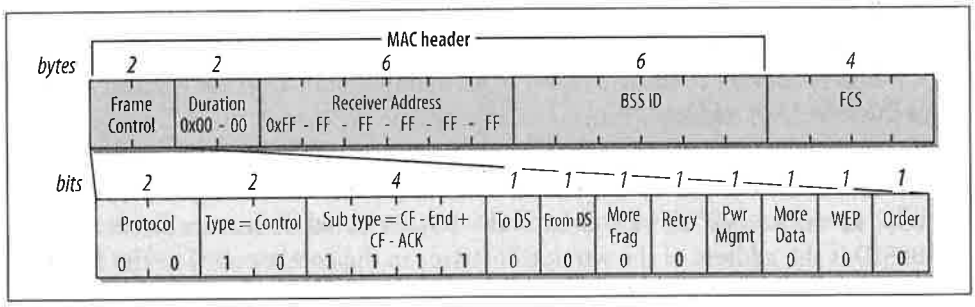

Figure 8-8. CF-End+CF-Ack frame

CF Parameter Set

Access points that support contention-free operation may include the CF Parameter Set information element, which is shown in Figure 8-9. CF Parameter Set elements are included in Beacon frames to keep all mobile stations apprised of contention-free operations. They are also included in Probe Response frames to allow stations to learn about contention-free options supported by a BSS. Four fields make up the CF Parameter Set information element:

CFP Count
: This field, which is one byte in length, tells how many DTIM frames will be transmitted before the start of the next contention-free period. Zero indicates that the current frame is the start of contention-free service.

CFP Period
: This one-byte field indicates the number of DTIM intervals between the start of contention-free periods.

CFP MaxDuration
: This value is the maximum duration of the contention-free period as measured in time units (TUs). Mobile stations use this value to set the NAV to busy for the entire contention-free period.

CFP DurRemaining
: This value is the number of TUs remaining in the current contention-free period. Mobile stations use it to update the NAV throughout the contention-free period. When DCF-based contention-free service is provided, it is set to 0.

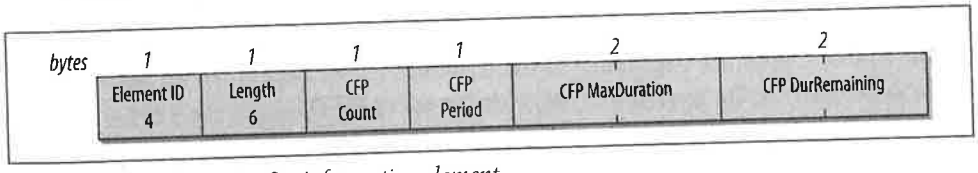

Figure 8-9. CF Parameter Set information element

Power Management and the PCF

Power conservation during the contention-free period is similar to power conservation during the contention-based period, with a few minor exceptions. The basic distinction between the two is that frame delivery must obey the PCF rules, so buffered frames can be delivered only to CF-Pollable stations. Stations that do not support PCF operations must wait until contention-based service resumes before retrieving buffered frames.

Stations on the polling list are not allowed to sleep during the contention-free period. When the access point is performing its point coordination functions, it may poll any station on the polling list at any time. Frames destined for stations on the polling list

do not need to be buffered during the contention-free period because those stations do not sleep.

Frame buffering is identical under contention-free and contention-based service. By maintaining power-saving status for each station, the access point can buffer frames for any station in a low-power mode. Broadcast and multicast frames are buffered whenever an associated station is in a low-power mode.

In addition to the buffer status associated with contention-free service, the access point also sets bits in the TIM for any station it intends to poll. The reason for setting these bits is related to how buffered frames are delivered. Like contention-based service, DTIM frames trigger the transmission of broadcast and multicast frames. If the total time required to transmit multicast and broadcast frames exceeds the Beacon interval, the access point will transmit one Beacon interval's worth of buffered frames and stop. Remaining frames will, however, cause the access point to keep the bit corresponding to AID 0 set.

After transmitting the buffered broadcast and multicast frames, the access point goes through the list of AIDs whose TIM bits are set in increasing order and transmits any pending data. Transmissions are conducted according to the rules of the PCF, so it is not necessary to include a delay before beginning transmission. Stations on the polling list are added to the TIM, so they will be included in this process. Multiple buffered frames can be transmitted, but this is entirely up to the access point implementation—in contention-free service, mobile stations can transmit only when given permission by the access point. A station is not allowed to resume sleeping until all frames have been delivered to it, as indicated by a 0 More Data bit. When a station is cleared to resume sleeping, it sleeps until the next DTIM transmission. DTIM frames signal the beginning of the contention-free period, so all stations that implement the PCF are required to wake up for every DTIM.

If a station switches from a low-power mode to the active mode, any frames buffered for it are transferred to the point coordination function for delivery during the contention-free period. The transfer does not result in immediate delivery, but the access point can place the frames into a queue for transmission as soon as the point coordination function permits.