CHAPTER

10 Wireless PANs/IEEE 802.15x

Upon completion of this chapter, the student should be able to:

- Explain the basic differences between wireless PANs and WLANs.
- Discuss the evolution of the IEEE 802.15 standard from the Bluetooth standard.
- Discuss the various types of wireless PAN networks that may be set up under the IEEE 802.15.1 standard.
- Discuss the details of the WPAN physical and baseband layer.
- Discuss the Bluetooth protocol stack.
- Discuss the various extensions to the IEEE 802.15x standard.
- Explain the basic characteristics of the IEEE 802.15.3 and IEEE 802.15.4 WPAN implementations.

This chapter introduces the IEEE 802.15x standard for wireless personal area networks (WPANs). Beginning with a comparison of the functionality provided by a wireless LAN and a wireless PAN, the reader is next introduced to the origins of IEEE 802.15x—the Bluetooth specifications. Wireless PAN applications and network architectures are discussed in some detail. The basic characteristics of a WPAN are presented and contrasted against the basic operation of a WLAN. The various possible ad hoc network topologies that may be formed under the IEEE 802.15.1 standard are discussed and evaluated.

Once most of the introductory material about WPANs has been presented, the Bluetooth physical layer details for 2.4-GHz operation are presented. The function and operation of the Bluetooth link controller in carrying out baseband protocols and low-level link operations are discussed in the context of circuit-switched data connections and packet-switched connectionless data transfers. The details of timeslots, the different types of physical links, and packet formats are introduced to the reader. Next, the operational states of the Bluetooth link controller are covered in sufficient detail to provide a sense of how everything comes together to yield a functional system. A short coverage of the Bluetooth-specific protocols and host control interface operation is provided to complete the material about the IEEE 802.15.1 Bluetooth standard.

The chapter ends with an overview of the other IEEE 802.15x standards. Emphasis is placed on the new low-rate and high-rate WPAN technologies. Basic physical system characteristics are presented and modes of operation introduced and summarized.

10.1 INTRODUCTION TO IEEE 802.15X TECHNOLOGIES

On April 15, 2002, the Standards Board of the IEEE approved IEEE 802.15.1: the wireless medium access control (MAC) and physical layer (PHY) specifications for wireless personal area networks (WPANs). Before describing the details of the standard, it is necessary to define what is meant by a wireless **personal area network.** Essentially, a wireless PAN is used to transfer information over short distances (approximately ten meters maximum) between private groupings of participant devices. The goal of the standard is to provide wireless connectivity with fixed, portable, and moving devices either within or entering a **personal operating space** (POS). A POS is further defined as the space around an individual or object that typically extends ten meters in all directions and envelops the individual whether that person is stationary or in motion. One might even visualize this POS as a bubble surrounding the individual. The standard has been developed to coexist with all other IEEE 802.11 networks and to eventually allow a level of interoperability that would provide a means by which a WPAN could transfer data between itself and an IEEE 802.11 device. One should certainly also be aware of the differences between a wireless PAN and a wireless LAN. The WLAN's primary function is to extend the reach of a wired LAN with its inherent connectivity to the various data services available on an Ethernet-based computer network. Unlike the WLAN, the WPAN provides a wireless connection between devices that involves little or no physical infrastructure or direct connectivity to the world outside of the link. Due to these facts, the implementation of a WPAN can be achieved through small, extremely power-efficient, battery-operated, low-cost solutions for a wide and diverse range of personal devices.

Since a large portion of the IEEE 802.15.1 standard is adapted from portions of the Bluetooth wireless specification, it is appropriate to give some background information about the **Bluetooth** standard at this time. A Bluetooth Special Interest Group (SIG) was formed during the late 1990s in response to an industry-perceived market need. Composed of many important stakeholders within the wireless, telecommunications, and software industries (i.e., Ericsson, IBM, Intel, Nokia, Toshiba, Microsoft, 3COM, Agere, and Motorola), the group set about the task of developing a short-range wireless technology that could be used to eliminate cables between both stationary and mobile devices. Furthermore, this new technology would facilitate the transfer of both voice and data traffic through the formation of ad hoc networks between the rapidly expanding assortments of personal devices being introduced to the marketplace. The interesting origin of the use of the Bluetooth name for this new wireless technology will be left as an exercise for the reader. Also, the reader may want to visit an interesting Web site devoted to the Bluetooth standard, its proposed applications, and the Bluetooth SIG located at www.bluetooth.org.

The Bluetooth SIG has produced and published specifications for the initial Bluetooth standard (Specification 1.1) and a follow-on extension (Specification 1.2). The IEEE 802.15.1 standard derives most of its text from the initial Bluetooth standard (Spec 1.1). A follow-up revision project to IEEE 802.15.1 has recently been initiated that will incorporate the changes provided by the new Bluetooth specification (Spec 1.2) and add some additional modifications and improvements to both the MAC and PHY layers. The IEEE 802.15.1 standard is a rather formidable document that, as of this writing, stretches some 1169 pages in length. This chapter will attempt to provide an overview of the basic details about the architecture, implementation, and operation of the IEEE 802.15.1 standard and the derivative work that forms the basis for the other IEEE 802.15x standards. In presenting this material, the details that differentiate the WPAN standard from the WLAN standard will be emphasized and similarities will only be pointed out but not discussed to any great length. Particular attention will be given to the follow-on extensions to IEEE 802.15.1. Some of these emerging technologies have the potential to someday transform various segments of the industrialized world through the implementation of ubiquitous wireless sensor networks (WSNs) supporting applications for use in both the commercial/industrial and consumer/home environments.

10.2 WIRELESS PAN APPLICATIONS AND ARCHITECTURE

In recent times, there has been a proliferation of personal electronic devices entering the marketplace. These devices are being designed with ever increasing data capabilities and are becoming more intelligent and interactive in their behavior. The fact that these devices exist is a testament to the microelectronics industry's continuing ability to follow the roadmap predicted by Moore's law, the device designer's continuing improvements to the man-machine interface, and the public's rapid acceptance of these devices. As the semiconductor industry has been able to integrate more and more active devices onto the surface of a piece of silicon, it has also begun to use innovative techniques to lower overall IC chip power consumption through the use of lower operational voltages and nontraditional RF circuit designs using CMOS and SiGe/BiCMOS technologies. Additional circuit efficiencies are also being achieved through the use of novel on-chip RF microelectromechanical systems (MEMS) to replace former off-chip subassemblies. These facts have allowed designers to build devices with a great deal of embedded processing power, compact and efficient RF electronics circuitry, and small form factors that are acceptable to users of these products.

These and other factors have figured heavily into the appearance of low-cost, battery-operated personal digital assistants (PDAs), personal MP3 music players, digital cameras, and multimedia-enhanced mobile phones. These devices in conjunction with the more traditional notebook/laptop and newer tablet computers have driven this personal devices product space. The use of increased memory, processing power, and more ubiquitous high-speed wireless connectivity has provided these devices with the ability to retain, process, and transfer large amounts of digital information. Many of these devices, like the PDAs, maintain personal information management (PIM) databases. These databases are used to maintain personal calendars, address books, and so-called to-do lists. It is highly desirable to have all of one's personal devices in synchronization (i.e., all of the PIM databases in agreement). An obvious solution to this difficulty is to provide wireless interconnectivity (via an ad hoc network) between the various personal devices that are typically used by an individual during the course of his or her daily activities.

The other projected major use of wireless PAN technology is for the elimination of the numerous cables that are presently needed to provide wired connectivity between the aforementioned personal devices and PCs. The idea here is that within one's POS there would be no need for interconnecting and oftentimes propriety cables. This application would provide the user with the ability to transfer data between various devices by just moving into close proximity with the desired device. Actually, an effort to provide this data transfer functionality to laptop computers and several other personal devices was attempted toward the late 1990s through the use of infrared (IR) technology. This technique has never really enjoyed much popularity due to the need for extremely close range between devices and the need to aim or direct the IR signal between the devices. Although the public has embraced this technique for the operation of remote controls for consumer entertainment products, channel surfing is a far different activity than the implementation of hands-free, cableless, user-friendly, transparent data transfers between personal devices.

Before considering the basic characteristics of wireless PANs, it is appropriate to consider how the integration of WPAN functionality into the device should impact its use. First and most importantly, the personal device's primary functionality should not be affected by any type of wireless connectivity built into it. Furthermore, the device should retain its prior form factor, weight, power demands, cost, and basic usability. As stated before, the interconnecting of personal devices is different than the connecting of computing devices with a wired computer network. In the case of the WPAN, the individual is primarily concerned with the personal devices in either his or her possession or in his or her vicinity, wherever that location might be. This emphasis might change in the future but for the time being the personal and private nature of the wireless PAN connection is of major importance.

Basic WPAN Characteristics

A seemingly effective way to explain the basic characteristics of wireless PANs is to contrast them with the characteristics of wireless LANs. Both WLANs and WPANs appear to be very similar in their operation

(i.e., both are able to connect wirelessly to their surrounding environment and exchange data over an unlicensed portion of the frequency spectrum). However, that is where the similarity ends. The WLAN has been designed to support transportable types of computing (i.e., clientlike devices) like that provided by laptops/notebooks or tablet PCs. The WPAN standard has been designed to support more mobile personal devices. With this in mind, the next few sections will discuss the three fundamental ways in which these two technologies differ:

- WPAN power levels and coverage areas
- Media control techniques
- Network life span or duration

WPAN Power Levels and Coverage Areas

A WLAN is typically deployed to ensure as large a coverage area as possible. This translates into power levels of approximately 100 mW, with coverage distances of approximately 100 meters, supplied by radio base stations (access points). Due to the WLAN power requirements, these access points need to be placed in optimized fixed locations, connected to a wall outlet via a power cord, and connected to the wired computer network via a LAN cable. The use of a WLAN enables the deployment of a LAN where the use of cables is either difficult or costly to install. The WLAN must still be deployed and set up in any case and primarily serves to extend the reach of a portable device to connect to an established infrastructure. The WPAN, on the other hand, is designed to interconnect multiple personal devices. Furthermore, these personal devices tend to be mobile versus portable. A definition of a mobile device is warranted here. Personal mobile devices typically are battery powered and experience brief interconnection periods with other devices. Portable devices tend to be moved less frequently, usually experience longer network connection durations, and tend to be powered from standard wall outlets. To reiterate, the portable device is typically using the wireless connection to access LAN-based services.

A WPAN uses low power consumption to enable true mobility. The typical WPAN has a coverage area of approximately ten meters with a transmitting output power of 1 mW. Personal devices are able to achieve low-power modes of operation that allow several devices to share data through the use of WPAN technology. Again, it should be pointed out that a personal device typically does not have the need to access LAN-based services; however, that possibility is not excluded in the IEEE 802.15 standards.

WPAN Media Control Techniques

Since many different types of personal devices may participate in a wireless PAN, there is a need for the standard to support numerous different types of applications. Furthermore, the applications will typically require different levels of QoS (i.e., scheduled or unscheduled bandwidth). With this goal in mind the basic structure and operation supported by the wireless PAN standard consists of the formation of **ad hoc networks** that are controlled by a single member of the PAN known as the **master** (i.e., a mechanism that controls other similar mechanisms). The other member or members of the ad hoc PAN function as **slaves** meaning that they are mechanisms that are controlled by the personal device that has taken on the role of the master. With this type of structure in place, and through the use of a time-multiplexed slotted system, the master is able to poll the slave members of a wireless PAN and thus determine the required bandwidth needs. The device serving as the master is then able to regulate the bandwidth assigned to the various slave personal devices based upon the required QoS requested. Through use of a system that employs short timeslots high-quality traffic may be supported.

Furthermore, the ad hoc nature of the wireless PAN necessitates that a personal device must be able to act as either the master or the slave within a newly formed network. This fact drives the design of WPAN technology, since the personal device, regardless of the role it assumes, must still conform to a low-cost, low-power implementation. Personal devices that are able to provide WPAN functionality are primarily purchased for their personal appeal and the services that they can provide. They are typically not meant to be members of an established networking infrastructure. A WLAN device is required to maintain a

management information database (MIB) to facilitate end-to-end network operations of a larger infrastructure. The WPAN device presently does not need to maintain a network-observable and network-controllable state to provide this type of WLAN functionality. This does not mean that end-to-end solutions cannot be implemented by WPAN technology, just that presently they will need to be overlaid onto it unless extensions to the standard evolve to implement these functions.

WPAN Network Lifespan

Once a WLAN has been deployed it is placed into existence. An access point may or may not have any wireless LAN stations associated with it for the WLAN to exist. A WPAN does not conform to this model. In all cases, for communications to occur over a WPAN a master must exist. If the WPAN master does not participate in an ad hoc network, the network no longer exists. For a WPAN, a device can create a connection that lasts only as long as needed and therefore the network has a finite life span. If a digital picture is to be transferred from a camera to a PC, the network might exist only as long as needed to transfer the picture. Since the connections created in a WPAN are ad hoc and temporary in nature, the personal devices that are connected at one moment may bear little or no resemblance to what was previously connected by the network or what will be connected by the network in the future. As an example, a person might return home from work and allow both PDA and cell phone to connect to his or her home PC and to each other. Afterwards, the person might download a movie from a digital camera into the PC, and so on and so forth. In all cases, the WPAN allows for the rapid formation of ad hoc networks that provide wireless connectivity without any predeployment activity necessary.

Bluetooth WPAN Overview

The Bluetooth wireless specification provides for communications over a relatively short-range radio link that has been optimized for battery-operated, compact, personal devices (see Figure 10–1 for a typical Bluetooth device). The Bluetooth WPAN provides support for both asynchronous communications channels for data transfer and synchronous communications channels for telephony-grade voice communications. Using Bluetooth wireless technology, a user could simultaneously be provided hands-free cellular telephone operation via a Bluetooth-enabled wireless headset and at the same time be transferring packet data from the cellular mobile phone to a laptop/notebook PC.

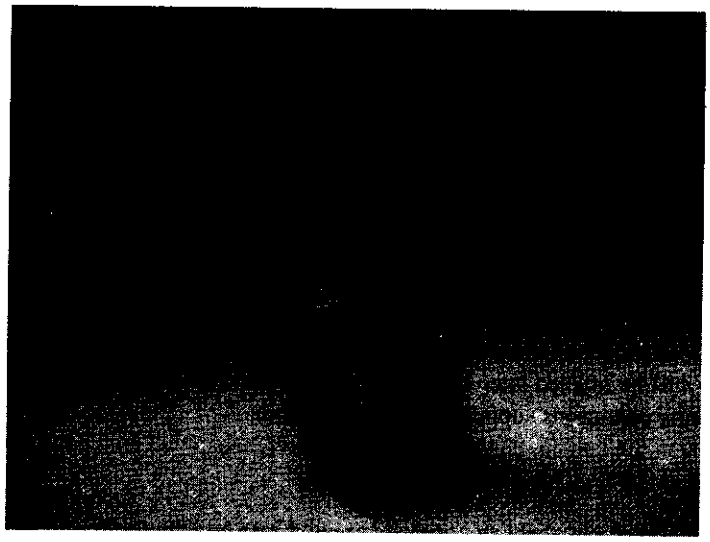

Figure 10–1 Bluetooth to USB port adaptor.

The Bluetooth specification calls for the use of the 2.4-GHz unlicensed ISM band. A fast frequency hopping scheme is employed to prevent interference and signal fading. The baseband data is preshaped with a Gaussian filter and then modulated by using binary frequency shift keying (BFSK) at a symbol rate of 1 msps. The use of frequency hopping at a rate of 1600 hops/s or 625 μs/hop and binary FSK modulation yields a fairly simple transceiver that can typically be implemented as a system on a chip (SOC) integrated circuit (IC). A slotted channel format is used with a slot (or hop) duration of 625 μs. This allows for full-duplex operation using a fast time division duplex (TDD) scheme. Over the radio link, information is transferred in packets. Because of the frequency hopping scheme, each packet is transmitted on a different frequency. A packet normally is only a single slot in length but can be extended up to three or five slots. Data traffic can have a maximum asymmetric rate of 723.2 kbps between two devices. Bidirectional, synchronous 64-kbps channels are able to support voice traffic between two devices. Various combinations of asynchronous and synchronous traffic are allowed. Figure 10–2 shows the format of an over-the-air, single-slot Bluetooth packet. The figure indicates that each packet consists of an access code, a header, and a payload. More detail will be offered later about the functions of the access code and header portions of a Bluetooth packet.

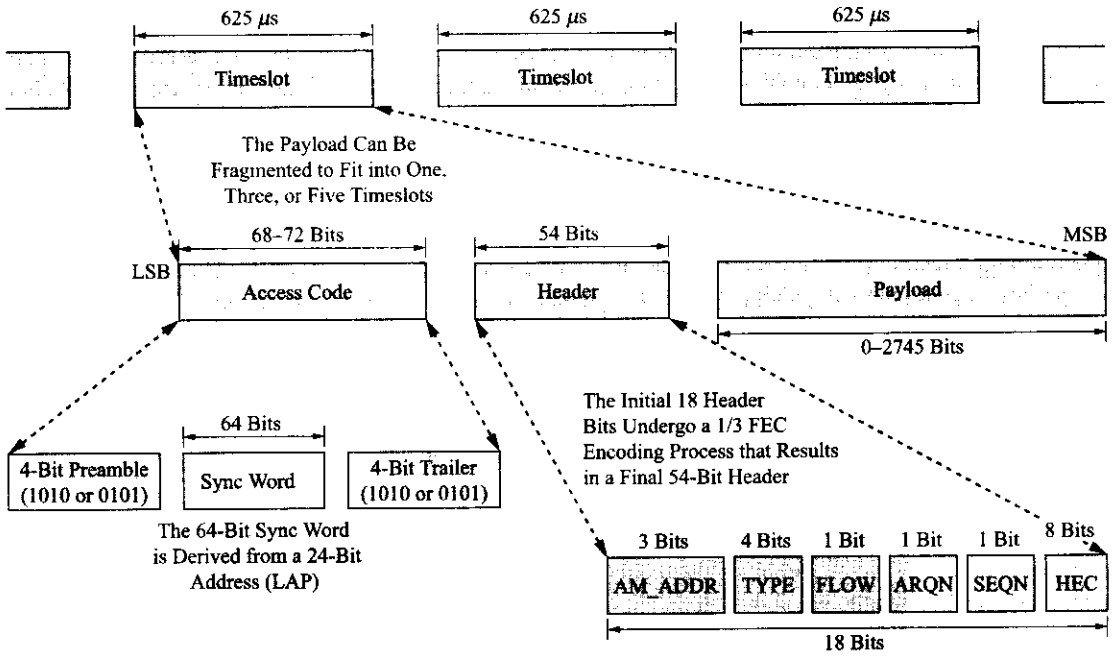

Figure 10–2 Format of a single-slot Bluetooth packet (Courtesy of IEEE).

Before presenting more about the architectural details of the actual ad hoc networks allowed by the Bluetooth specification, a few words about the mapping of the IEEE 802.15.1 WPAN standard to the OSI model are appropriate. Figure 10–3 depicts Bluetooth wireless technology and the OSI protocol stack. As the reader can see, the standard maps to the physical and the MAC layer as was the case for IEEE 802.11 WLANs. More detail will be offered later about the functions of the various MAC sublayer, baseband protocols, and physical layer operations supported by the standard.

Bluetooth WPAN Ad Hoc Network Topologies

The two basic types of ad hoc networks that may be entered into by Bluetooth-enabled devices are piconets or scatternets. A **piconet** (see Figure 10–4) is formed by a Bluetooth device serving as a master and at least

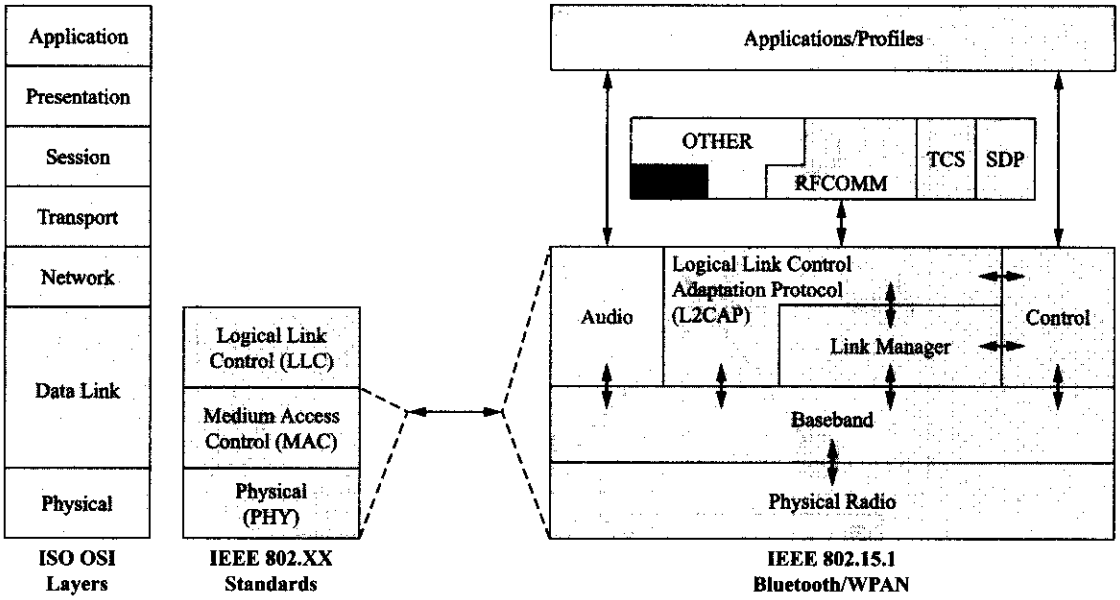

Figure 10–3 Comparison of Bluetooth technology and the OSI model (Courtesy of IEEE).

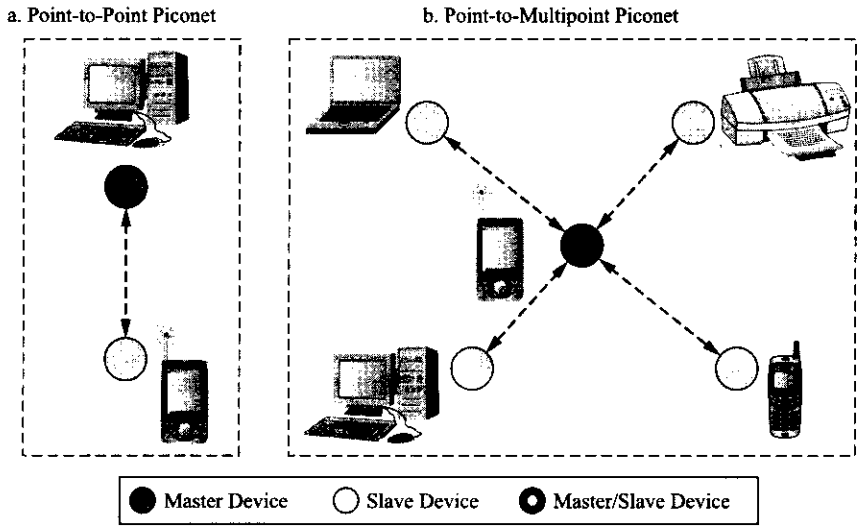

Figure 10–4 Bluetooth piconet architectures (Courtesy of IEEE).

one or more (up to a maximum of seven) Bluetooth devices acting as slaves. The piconet is defined by the frequency hopping scheme of the master. All devices that are taking part in a piconet are synchronized to the clock of the master of the piconet and hence to the same frequency hopping sequence. The piconet slaves only communicate with the piconet master in a point-to-point fashion and under the direct control of the master. However, the piconet master may communicate in either a point-to-point or point-to-multipoint fashion. Various usage scenarios might tend to define a certain device's role within a piconet as always being either a master or a slave; however, the standard does not define permanent masters or slaves. A device that has served as a slave for one application could just as easily be the master in another situation.

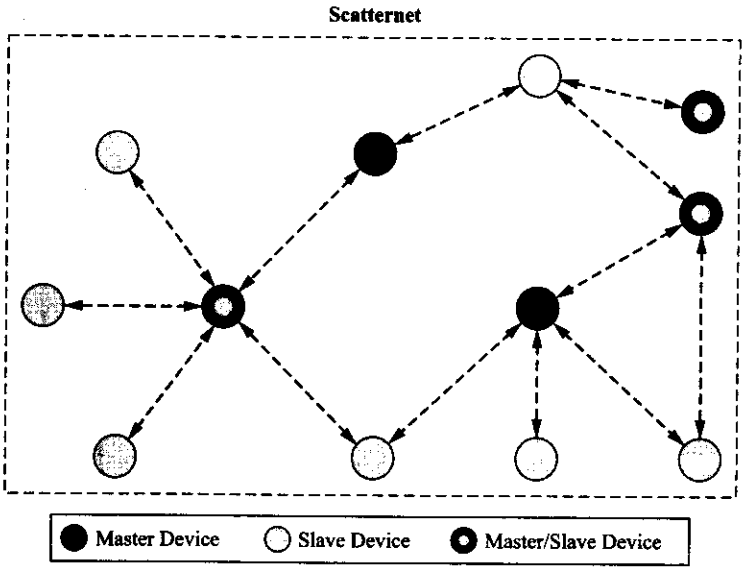

Figure 10–5 Typical Bluetooth scatternet structure.

The scatternet is the other Bluetooth ad hoc network supported by the IEEE 802.15.1 standard. The **scatternet** is a collection of functioning piconets overlapping in both time and space (see Figure 10–5). Through the scatternet structure, a Bluetooth device may participate in several piconets at the same time. A device in a scatternet may be a slave in several piconets but can only be a master of a single piconet. Furthermore, a device may serve as both a master and a slave within the scatternet. An interesting result of the scatternet structure is that information may flow beyond the coverage area of a single piconet.

Although what is to be presented next is not another form of Bluetooth ad hoc network, the integration of a Bluetooth WPAN with other LANs is possible and needs to be discussed. Figure 10–6 shows this situation.

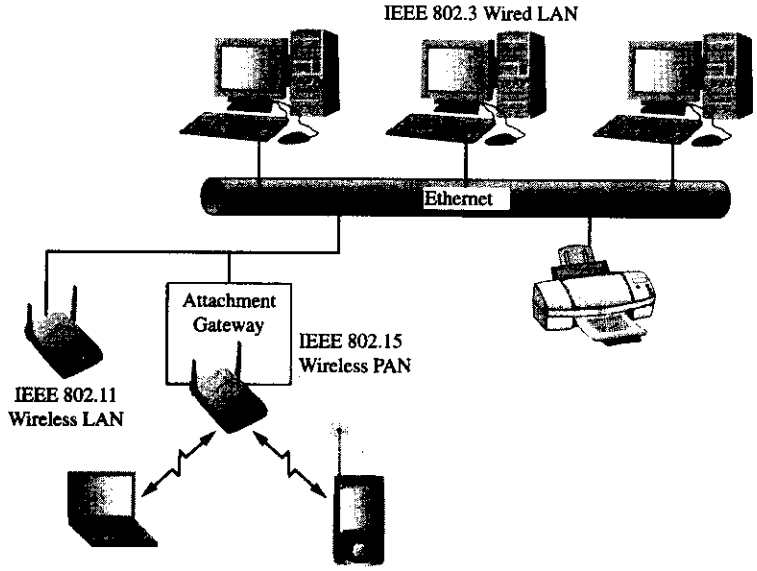

Figure 10–6 Integration of IEEE 802.15 and IEEE 802.11 networks.

Through the use of an IEEE 802 LAN attachment gateway (AG) a Bluetooth WPAN may connect to and participate in the transfer of data with other LANs in the IEEE 802 family. The LAN attachment gateway allows for the transfer of MAC service data units (MSDUs) from or to other LANs via the wireless connectivity afforded by the Bluetooth WPAN.

Components of the Bluetooth Architecture

As is the case with all communications protocols, the ultimate goal of the protocol is to enable applications running in different devices with the ability to exchange data with each other. To facilitate this operation requires that compatible protocol stacks are running in these devices and that the applications that reside on top of these protocol stacks are also equivalent. Therefore, to provide the desired WPAN functionality, the Bluetooth standard calls for a set of communications protocols and a set of interoperable applications that are used to support the usages addressed in the specifications. To give the reader a sense of how this relates to the physical and MAC layers within the IEEE 802.15.1 standard an overview of this concept will be presented here. Figure 10–7 shows the Bluetooth protocol stack. Recall from Figure 10–3 that the functions within the lower box shown in the figure are equivalent to the physical and MAC layers of the IEEE 802 standards.

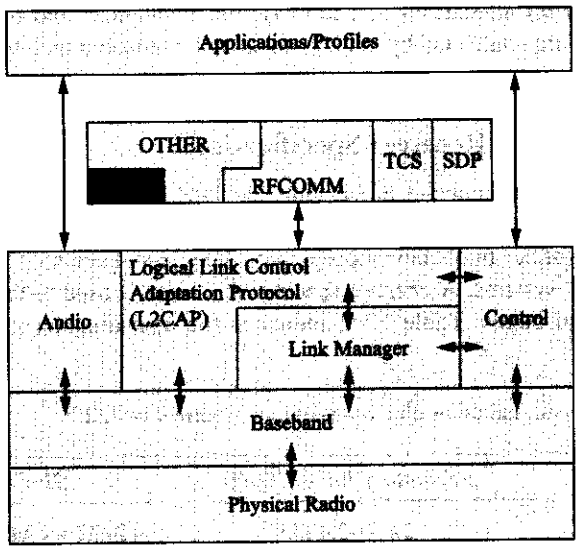

Figure 10–7 The Bluetooth protocol stack (Courtesy of IEEE).

Figure 10–7 shows both Bluetooth-specific protocols and other non-Bluetooth-specific protocols. The link manager protocol (LMP) and the logical link control and adaptation layer protocol (referred to as L2CAP) are Bluetooth specific whereas the protocols within the "Other" box are not. Some of these other protocols would probably include the point-to-point protocol (PPP), the wireless application protocol (WAP), and other similar or to-be-implemented future protocols. The design of the Bluetooth protocol stack was based on the reuse of existing protocols for support of higher-layer functions. The reuse of protocols also allows for the easy adaptation of existing applications to work with Bluetooth wireless PAN technology and the development of new applications that can take advantage of the Bluetooth technology. The figure also shows that the logical link control (LLC) protocol is not part of the Bluetooth specification but is grouped within the "Other" box. LLC traffic would therefore have to be encapsulated by the Bluetooth network encapsulation protocol (BNET) before being passed down to the MAC sublayer and physical

layer. The details of the other protocols in the Figure 10–7 are as follows: RFCOMM is a serial cable emulation protocol based on an ETSI standard, TCS is a telephony control and signaling protocol, and SDP is a service discovery protocol used to allow Bluetooth devices to determine what services other Bluetooth devices may provide.

10.3 IEEE 802.15.1 PHYSICAL LAYER DETAILS

The physical layer of the Bluetooth protocol stack is shown in Figure 10–7. At the sending device, the function of this layer is to receive a bit stream from the MAC sublayer and transmit the bit stream to another Bluetooth-enabled device over a radio link existing between the two devices. In a complementary fashion, the function of this layer at the receiving device is to receive the radio link signal from the sending device, convert the radio signal into a bit stream, and pass the demodulated bit stream to the MAC sublayer of the receiving device. The Bluetooth physical layer provides for the radio transmission and reception functions but provides no interpretation functions. The Bluetooth specification calls for the transceiver to operate in the 2.4-GHz ISM band. This being the case, the Bluetooth device must satisfy the regulatory requirements for the geographic area that it is operated in. Presently, the standard calls for compliance with established regulations for Europe, Japan, and North America. The next section will present more detailed information about the channel allocations and the required transmitter and receiver specifications. Frequency hopping sequences are controlled by the Bluetooth link controller and will be discussed later in this chapter.

Channels, Transmitter, and Receiver Specifications

IEEE 802.15.1-compliant systems operate in the 2.4-GHz ISM band. Some portions of these frequencies (2.400 to 2.4835 GHz) are available in a majority of countries around the world. In the countries with limited spectrum availability, special frequency hopping sequences have been specified similar to what is done with IEEE 802.11x WLAN systems. A channel spacing of 1 MHz is used with a guard band employed at both lower and upper band edges. Table 10–1 indicates the operating frequency bands for Bluetooth devices.

Table 10–1 Operating frequency bands for Bluetooth devices (Courtesy of IEEE).

Region of Use	*Regulatory Range (GHz)*	*RF Channels*
United States, Europe and most other countries	2.400–2.4835	$f = 2402 + k$ MHz, $k = 0, \ldots, 78$
France	2.4465–2.4835	$f = 2454 + k$ MHz, $k = 0, \ldots, 22$

The transmitters used for Bluetooth operation must conform to the specifications shown by Table 10–2. As indicated by the table, transmitters can fall into three power classes. Class 1 devices must provide for a power control mechanism that is used to limit output power over 1 mW or 0 dBm. Furthermore, a Class 1 device with a maximum power output of 100 mW or +20 dBm must be able to control its transmitted output power down to +4 dBm or less. The transmitter output power should be adjustable in equal-step sizes ranging from 8 dB to 2 dB per step. A Class 1 device will optimize its output power in a radio link through the use of link management protocol (LMP) commands. This operation is accomplished by the measuring of received signal strength (RSS) by the receiving device, the use of the RSS indication (RSSI) to determine if the transmitted power should be increased or decreased, and the return transmission indicating a

Table 10–2 Bluetooth transmitter power output specifications (Courtesy of IEEE).

Power Class	*Maximum output power (P_{max})*	*Nominal output power*	*Minimum output power*[1]	*Power control*
1	100 mW (+20 dBm)	N/A	1 mW (0 dBm)	P_{min} < +4 dBm to P_{max} Optional: P_{min}[2] to P_{max}
2	2.5 mW (+4 dBm)	1 mW (0 dBm)	0.25 mW (–6 dBm)	Optional: P_{min}[2] to P_{max}
3	1 mW (0 dBm)	N/A	N/A	Optional: P_{min}[2] to P_{max}

[1]Minimum output power at maximum power setting

[2]The lower power limit P_{min} < –30 dBm is suggested but not mandatory

request to alter the transmitting device's output power. A Class 1 device may not send packets to a device that does not support RSSI measurements. If the receiving device cannot support power control operation by making the required measurements, a Class 1 device must comply with the rules of operation for a Class 2 or Class 3 transmitter.

The transmitter modulation details are essentially identical to those for the original IEEE 802.11 FHSS physical layer for a data transmission rate of 1 mbps. Further detailed specifications for both in-band and out-of-band spurious emissions and RF tolerances are given by the standard but will not be discussed here.

Bluetooth receiver specifications call for a sensitivity of –70 dBm or better to achieve a raw bit error rate (BER) of 0.1% or less. Although many other receiver specifications are listed by the standard, at this point, only one of these will be considered. It is the received signal strength indicator (RSSI) specification. If a receiver is to be able to participate in a power controlled link, it must be able to provide an RSSI measurement that can be compared to two threshold power levels that define the ideal receiver power range. Figure 10–8 depicts these two levels. The lower threshold value is between –56 dBm and 6 dB above the actual receiver sensitivity and the upper threshold is 20 dB above the lower threshold level ±6 dB.

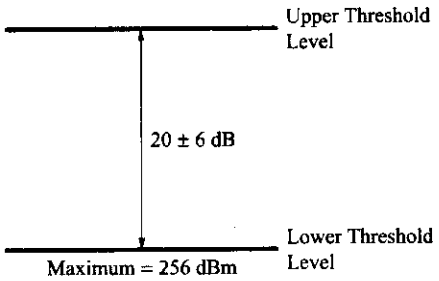

Figure 10–8 Upper and lower threshold powers for ideal receiver operation (Courtesy of IEEE).

10.4 BLUETOOTH LINK CONTROLLER BASICS

The Bluetooth system consists of a 2.4-GHz radio transceiver unit, a link control unit, and a link manager (see Figure 10–9). The basic radio (physical layer) functions have already been described in the last section, the link controller unit that carries out the baseband protocols and other low-level link operations will be described in this section, and the link manager that provides link setup, security, and control will be described in a later section.

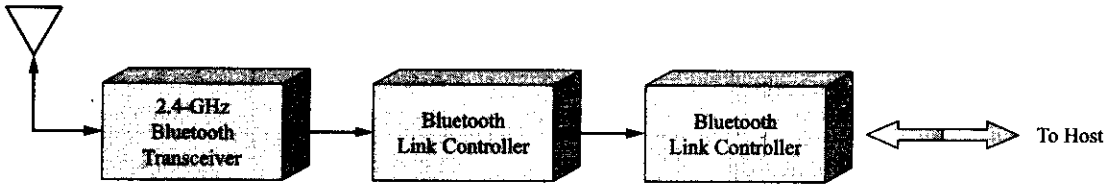

Figure 10–9 Bluetooth system components.

As previously described, the Bluetooth protocol is able to support both circuit-switched data and packet-switched data transfers. This type of functionality is made possible through the use of a slotted format during which data is transferred. Slots may be reserved for synchronous packets as well as asynchronous traffic. The Bluetooth standard is able to support the following different combinations of traffic: up to three simultaneous synchronous voice channels (64 kbps in each direction), a channel that simultaneously supports asynchronous data traffic and synchronous voice traffic, an asynchronous data channel with a maximum rate of 723.2 kbps in one direction and a 57.6 kbps data rate in the return direction, and a symmetric mode of operation with a data rate of 433.9 kbps in both directions.

As mentioned before, the Bluetooth standard calls for either point-to-point or point-to-multipoint connections. For a point-to-point connection as few as two Bluetooth-enabled devices may be involved (i.e., forming a minimal-size piconet of two). For a point-to-multipoint connection the same channel is shared among several Bluetooth-enabled devices. In any piconet one of the devices acts as a master and the other devices act as slaves. Within the piconet seven slaves may be active at any one time but there may be many more slave devices associated with the piconet existing in what is known as the parked state. These devices are not active on the channel but they remain synchronized to the master device's clock and hence to the piconet frequency hopping scheme. In all cases, access to the channel for both active and parked devices is controlled by the master device. The other type of ad hoc network possible under the Bluetooth standard is the scatternet, a combination of various piconets. In this configuration, there can be multiple masters (only one per piconet however) and multiple slaves. A slave is able to belong to more than one piconet through time division multiplexing and each piconet employs its own hopping channel or sequence/phase. The details of how all this occurs will be presented shortly.

Bluetooth Timeslot Format

As previously described, the Bluetooth channel is divided into 625 μs timeslots. The timeslots are numbered according to the clock of the piconet master. Each Bluetooth-enabled device has a clock count that ranges from 0 to $2^{27} - 1$ and then starts over. Figure 10–10 shows the time division duplex (TDD) scheme used by the Bluetooth system. The master and a particular slave take turns transmitting in alternate timeslots. The

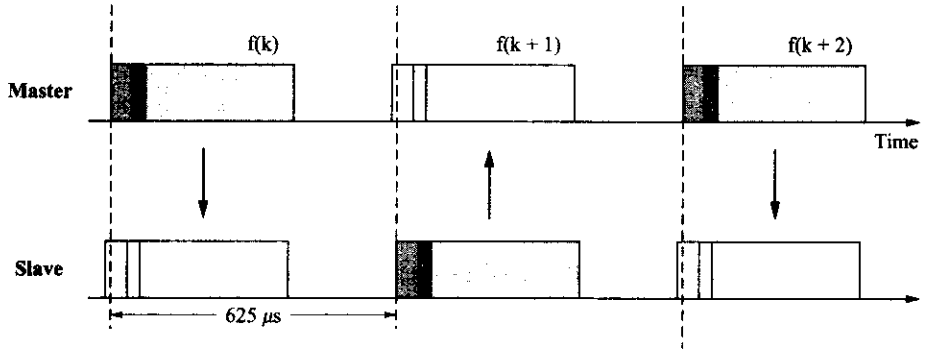

Figure 10–10 Bluetooth time division duplex transmission scheme (Courtesy of IEEE).

master starts transmitting only during even timeslots and the slave will start its transmissions only in odd-numbered timeslots. The start of the data packet coincides with the start of the timeslot. An alternate mode of operation allows the packets transmitted by either the master or slave to extend over a period of up to five timeslots.

As mentioned before, for single-packet transmission, the RF hop frequency remains fixed for the duration of the packet and is derived from the current master clock value. To support the transmission of multiple timeslot packets a slightly different procedure is followed. For a multislot data packet, the RF hop frequency that is used for the entire packet duration is derived from the master clock value at the start of the first timeslot of the packet. The RF hop frequency to be used for the timeslot that is to be used for transmission of a new packet is derived from the current master clock value for that timeslot. Figure 10–11 depicts this system operation in detail.

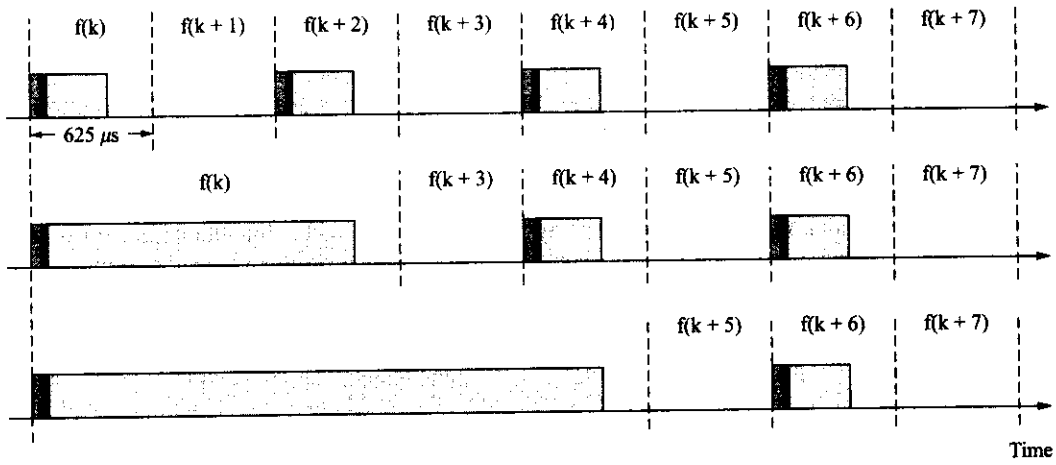

Figure 10–11 Bluetooth multislot packet transmission (Courtesy of IEEE).

Types of Physical Links

To support both synchronous and asynchronous modes of operation between master and slave(s) different types of links can be established. The two link types that have been defined by the Bluetooth standard are:

- Synchronous connection-oriented (SCO) link
- Asynchronous connectionless (ACL) link

The synchronous link is a symmetric, point-to-point link between the piconet master and a single specific slave. The master implements the SCO link by using reserved slots at regular periodic intervals. The asynchronous link is for point-to-multipoint links that exist between the piconet master and all the slave devices actively associated with the piconet. In slots not reserved for SCO operation the master can establish an ACL link on a slot-by-slot basis to any slave including one that is also participating in a SCO link.

The SCO link appears to be a circuit-switched connection between the master and the slave since the link reserves timeslots and therefore provides a known QoS to the connection. This type of physical link is typically used to support voice traffic. The Bluetooth specification allows the piconet master to support combinations of up to three SCO links to the same slave or to other slaves. Conversely, a slave may support up to three SCO links from the same master or up to two SCO links from different masters. On an SCO link, packets are never retransmitted. The SCO link is created by the piconet master by sending an SCO link setup message using the link management protocol (LMP). This message will indicate the link timing parameters such as the SCO interval (given as a number of slots) and a timing offset (again, given in slots) to identify the start of the link. The SCO slave is always allowed to respond with an SCO packet during the next slave-to-master timeslot following a master-to-slave SCO packet transmission.

The ACL link appears as a packet-switched connection between the master and all the active piconet slaves. In this mode of operation, slots not reserved for SCO links may be used by the master to exchange data with any slave on a slot-by-slot basis. However, only one ACL link at a time may exist between a master and a slave for this configuration. The slave is allowed to return an ACL packet in the next slave-to-master slot only if it has been addressed in the preceding master-to-slave slot. An ACL packet that does not specify a particular slave is considered a broadcast packet and is read by every active piconet slave. For ACL links packet retransmission is allowed. Since there are many types of ACL and SCO packet types, a wide range of data transfer rates may be supported by the Bluetooth standard. The next section will provide more details about this topic.

Packet Formats

Earlier in this chapter Figure 10–2 was used to introduce the reader to the typical over-the-air Bluetooth packet. At this time, a more in-depth look at this packet structure will be undertaken. The reader may want to refer back to Figure 10–2 or the figures included in this section while reading about this topic. As has been already indicated, data is transferred over the Bluetooth channel in packets. The generic packet format consists of an access code, a header, and the payload. The access code and the header information provide the Bluetooth ad hoc network with operational flexibility while performing various functions and procedures. Access codes are basically used for system synchronization and identification. All packets transferred within the same piconet are preceded by the same access code; thus the access code serves to identify the piconet. The access code is also used in both Bluetooth paging and inquiry activities. In this case, the access code is transmitted as a signaling message without any header or payload information. The header consists of six subfields and contains link control information like the address of an active piconet participant and the type of packet being used over the link.

Access Code Types

The Bluetooth specification provides for three different types of access codes. The format used for the access code is shown by Figure 10–12. The three different access codes are used by a Bluetooth-enabled device in different operating modes. The channel access code (CAC) is used to identify and synchronize the piconet. This code is transmitted with every packet data transfer within a piconet. The device access code (DAC) is used for signaling procedures like paging and response to paging. The inquiry access code (IAC) has two options. The general inquiry access code (GIAC) can be used to discover what other Bluetooth devices are within range of the inquirer. The dedicated inquiry access code (DIAC) can be used to discover Bluetooth devices that share a common characteristic or trait and are also within range of the inquirer. As shown in either Figure 10–2 or 10–12 the CAC format consists of a preamble, a sync word used for system timing, and a trailer (72 bits total). For a DAC or IAC signaling message, the trailer bits are not included and the message length is 68 bits. The preamble and trailer are fixed 4-bit binary combinations whose values depend upon the value of the first and last bits in the sync word. The sync word itself is derived from the unique Bluetooth address of the piconet master device for a CAC, from the slave's Bluetooth address for a DAC, and from certain dedicated binary words for an IAC.

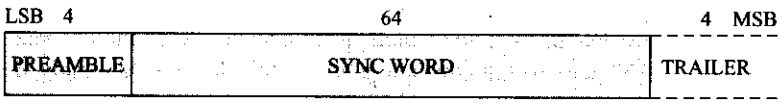

Figure 10–12 Bluetooth access code format (Courtesy of IEEE).

Packet Header Details

The packet header field contains various Bluetooth link control information in six different subfields (see Figure 10–13):

Figure 10–13 Bluetooth packet header field details (Courtesy of IEEE).

- AM_ADDR: 3-bit active member address
- TYPE: 4-bit type code
- FLOW: 1-bit flow control
- ARQN: 1-bit acknowledgement indication
- SEQN: 1-bit sequence number
- HEC: 8-bit header error check

Each active piconet slave is assigned a temporary 3-bit address to identify it when it is active. The value contained in the AM_ADDR field represents a particular active member of a piconet and thus is used to address that piconet device. All packets exchanged between the master and a particular slave carry the AM_ADDR of the slave device. The AM_ADDR address consisting of all 0s is designated as the piconet broadcast address. For this case, the master talks to all the active slaves within a piconet. If a slave disconnects or goes into a parked state, it gives up its AM_ADDR. Upon entering or reentering a piconet, a new AM_ADDR must be assigned to the Bluetooth device.

The 4-bit TYPE field allows the system to distinguish between the sixteen possible packet types that can be supported by the standard. How the TYPE code is interpreted also depends upon the type of link that it has been sent over. Once it has been determined what type of link (i.e., SCO or ACL) is in use, the type of packet may be deciphered. The TYPE code also indicates the total slot length of the packet. This information is useful for the devices that have not been addressed by a message since it allows them to know when the transmission will end. The other three single-bit fields are used for flow control, acknowledgement indications, and sequence control of retransmitted packets. The 8-bit header error check code is used to decide whether the packet is kept or discarded due to possible bit errors encountered during transmission. The total header consists of 18 bits that undergo a rate 1/3 forward error correction (FEC) encoding process that results in a total of 54 bits.

Packet Types

As just mentioned, the types of packets transferred over the piconet are related to the type of physical link they are used over (i.e., either SCO or ACL). For each of these link types several different packet types have been defined that provide unique operational modes over that particular link. Furthermore, there are also four control packet types and an ID packet that are common to both link types. They are NULL, POLL, FHS, DM1, and ID packets. The ID packet has already been described as the device access code (DAC) or the inquiry access code (IAC). The NULL packet is used for returning information to the source device about the success or failure of the previous transmission. The NULL packet consists of only the channel access code and the header and does not need to be acknowledged. The POLL packet is similar to the NULL packet but it does need to be acknowledged. The master may use this packet to poll the slaves and they must acknowledge its receipt. The frequency hop synchronization (FHS) packet is used to provide information about the sending device's address, current clock value, scan mode, power class, and the AM_ADDR the recipient should use. The FHS packet is used for page master response, inquiry response, and during master-slave switching procedures. The DM1 packet carries data information only. It is used to support control messages for each link type. When used over the SCO link it can interrupt the synchronous information flow, to provide control information.

SCO Packets There are four different types of SCO packets that are used over a synchronous SCO link: HV1, HV2, HV3, and DV packets. These packets do not include a CRC code and they are never retransmitted.

Bluetooth SCO packets are routed to the voice or synchronous I/O port. These four SCO packet forms consist of three similar but slightly different synchronous packet types and one packet type that carries both synchronous and asynchronous data. The HV1, HV2, and HV3 packets all carry high-quality voice (HV) at 64 kbps. The voice information may be encoded as either log PCM format (A-law or μ-law) or 64 kbps continuous variable slope delta modulation (CVSD). The difference between the three HVx packets is the amount of voice information carried per packet. HV1 packets consist of 10 information bytes. These bytes undergo a rate 1/3 FEC. The resulting payload length is 240 bits and there is no header present. Therefore, each HV1 packet carries 1.25 ms of speech at 64 kbps. HV2 packets consist of 20 information bytes. These bytes undergo a rate 2/3 FEC. Again, the resulting payload consists of 240 bits or 2.5 ms of speech at 64 kbps. HV3 packets consist of 30 information bytes. There is no FEC employed, the 240-bit payload is equal to 3.75 ms of speech at 64 kbps. To provide continuous voice, an HV1 packet must be sent every two timeslots, an HV2 packet every four timeslots, and an HV3 packet every six timeslots. The DV (data/voice) packet contains 10 voice information bytes and up to 150 bits of data. The DV packet is sent at regular intervals (SCO operation); however, the voice and data delivered by the packet are handled differently. The voice information is never retransmitted (it is always new for every packet) but the asynchronous data is checked for errors and the destination device may ask for it to be resent by the source.

ACL Packets The ACL link supports seven different types of packets for transmission within the piconet. The packet information may be either user data or control data. Six of the packets contain a CRC code and will be retransmitted if no acknowledgement occurs. The AUX1 packet does not employ a CRC code and is not allowed to be retransmitted. The seven ACL packet types are DM1, DM3, DM5, DH1, DH3, DH5, and AUX1 packets. The DMx packets carry varying amount of data (18, 123, and 256 bytes, respectively), are FEC encoded, and have durations of one, three, and five timeslots. The DHx packets are similar to the DMx packets except that they do not employ FEC encoding and therefore can carry 28, 185, and 341 information bytes over one, three, or five timeslots. The AUX1 packet is similar to the DH1 packet except that there is no CRC code so it can carry 30 bytes of information during one timeslot. Tables 10–3 and 10–4 provide a summary of the various ACL and SCO packet types with the resulting data transfer rates. As can be seen from the two tables, the Bluetooth standard provides a fairly good selection of symmetric and asymmetric asynchronous data rates and also provides for support of several channels of high-quality (64 kbps) voice traffic.

Table 10–3 Bluetooth asynchronous communications link types (Courtesy of IEEE).

Payload Type	*Payload Header (bytes)*	*User Payload (bytes)*	*FEC*	*CRC*	*Symmetric Maximum Rate (kbps)*	*Asymmetric Maximum Rate (kbps)*	
						Forward	*Reverse*
DM1	1	0–17	2/3	yes	108.8	108.8	108.8
DH1	1	0–27	no	yes	172.8	172.8	172.8
DM3	2	0–121	2/3	yes	258.1	387.2	54.4
DH3	2	0–138	no	yes	390.4	585.6	86.4
DM5	2	0–224	2/3	yes	286.7	477.8	36.3
DH5	2	0–339	no	yes	433.9	723.2	57.6
AUX1	1	0–29	no	yes	185.6	185.6	185.6

Table 10–4 Bluetooth synchronous communications link types (Courtesy of IEEE).

Payload Type	*Payload Header (bytes)*	*User Payload (bytes)*	*FEC*	*CRC*	*Symmetric Maximum Rate (kbps)*
HV1	N/A	10	1/3	no	64.0
HV2	N/A	20	2/3	no	64.0
HV3	N/A	30	no	no	64.0
DV[1]	1 D	10 + (0–9) D	2/3 D	yes D	64.0 + 57.6 D

[1]Items followed by "D" relate to data field only

In conjunction with the various combinations of data rates, the Bluetooth standard provides for several different levels of error correction. Both FEC and automatic request for retransmission (ARQ) schemes are implemented for header and payload data but will not be discussed any further here.

Transmitter/Receiver Timing

An important aspect of Bluetooth operation is the ability of the master and slave(s) to become time synchronized. Since data is transferred via time division duplex operation, successful system operation can only occur when all members of a piconet are in time synchronism. The piconet relies on the system clock of the master to provide this timing. The master never adjusts its clock. Instead, the piconet slaves adapt their clocks by providing a timing offset that causes their clocks to match the master clock. Each time the slave receives a transmitted packet during the master-to-slave timeslot, the received channel access code provides timing information that can be used to correct any timing misalignments and update the required offset time. A timing uncertainty window of ±10 μs is allowed for received packets for both the master and slave devices. Even though a timeslot is 625 μs in length, during the connection state, the length of a single-slot packet is limited to 366 μs (naturally, multislot packets have a longer duration). Both master and slave devices expect to receive packets that are synchronized to the start of timeslots and at intervals that are multiples of 1250 μs. Since the slave derives its timing from the piconet master, its transmission is scheduled to occur at $N \times 625$ μs (where N is an odd integer) after the start of a master transmission timeslot.

Depending upon the relative state of the Bluetooth device, timing behavior may deviate somewhat from what has just been described. While in the connection state, a slave Bluetooth device may be put into a hold mode during which it neither transmits nor receives information. Upon returning to normal operation from a hold mode, the slave must listen for the master before it may transfer any data. For this case, due to possible clock drift, the uncertainty window may be extended quite dramatically to allow for the resynchronization of the slave's clock offset. Other slave operational modes are the park and sniff modes. These modes are similar to the hold mode and require special operation on the part of the piconet slave. While in these modes, the slave device periodically wakes up to listen for transmissions from the master so it may resynchronize its clock offset. As was the case for the hold state, the slave may significantly increase its search window as shown in Figure 10–14.

In the page state, the piconet master transmits the device access code (an ID packet) for the slave that is to be connected to the ad hoc network. Due to the fact that the ID packet is very short and for other reasons to be revealed shortly, this operation is performed at a frequency hopping rate of 3200 hops/s or every 312.5 μs (twice as fast as normal). Therefore, during each timeslot the ID packet is transmitted on two different hop frequencies. During the receiving timeslot, the master device listens on the corresponding hop frequencies during the beginning of the timeslot and 312.5 μs later during the middle of the timeslot.

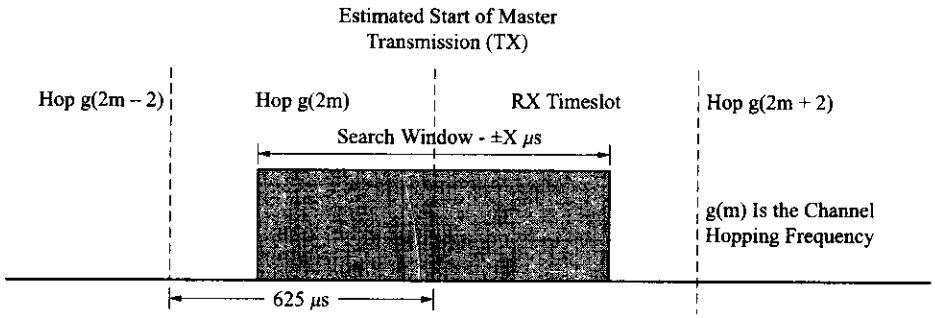

Figure 10–14 Extension of the Bluetooth clock resynchronization search window (Courtesy of IEEE).

At the time of a piconet connection setup and during a master-slave switch, an FHS packet is transferred from the master to the slave device. The packet is used to establish both timing and frequency synchronization. After the slave has received a page message, the slave unit returns a response message 625 μs later that consists of the ID packet. Since frequency hopping is taking place every 312.5 μs during the paging operation, it is possible that the paging response from the slave device will occur during the middle of the timeslot or only 312.5 μs before the beginning of the next timeslot. For this case, the master then sends the FHS packet 312.5 μs later in the next transmit slot after the receive slot during which the slave device responded and on the correct next hop frequency. The slave will then adjust its timing according to the received FHS packet and acknowledge its receipt 625 μs later.

Bluetooth Channel Control

Now that a great deal of detail about Bluetooth operation has been presented, it is time to wrap up some loose ends about system operations. An overview of the creation of a piconet will be provided here in the context of the states of operation of the Bluetooth devices that facilitate these functions. As stated previously, the channel in the piconet is defined totally by the piconet master. The Bluetooth device address (BD_ADDR) of the master device is used to determine the frequency hopping sequence and the channel access code. Furthermore, the master device's system clock determines the phase of the frequency hopping sequence and sets the system timing. Finally, the master device controls the piconet traffic through a polling scheme. It is again pointed out that all Bluetooth devices are identical and any unit can become either a master or a slave. The master is by definition the device that initiates the connection that forms the piconet. Once a piconet has been formed, a master and slave may even switch roles.

The Bluetooth Clock

An internal system clock is used by each Bluetooth device to determine the timing and frequency hopping of its transceiver subsystem. This clock (typically an accurate, low-drift, crystal oscillator) is free running and is never adjusted or turned off. Its value has no relationship to the time of day and therefore however it might randomly initialize itself upon power-up is of no consequence to system operation. Figure 10–15 depicts the Bluetooth clock implemented as a 28-bit binary counter that resets back to all zeros after counting up to $2^{28} - 1$. As shown by the figure, the master oscillator driving the clock has a frequency of 3200 Hz or a period of 312.5 μs (i.e., half of a timeslot cycle). The clock cycle itself takes about one day to complete. However, the important time periods generated by the clock and used by the system to trigger various operations are indicated in the figure. These periods are 312.5 μs, 625 μs, 1250 μs, and 1.28s. Also, during low-power states of operation (i.e., standby, hold, park, or sniff), a low-power oscillator with somewhat less timing accuracy may be used to replace the crystal oscillator.

The timing and frequency hopping on the channel of a piconet is determined by the master device's Bluetooth clock. During the creation of a piconet, the value of the piconet master clock is communicated to

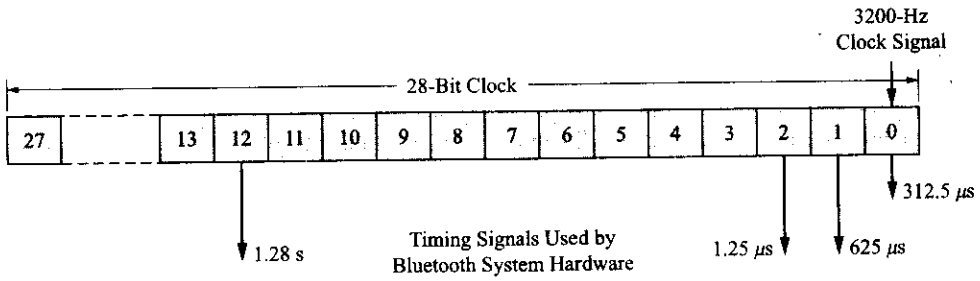

Figure 10–15 Bluetooth system clock (Courtesy of IEEE).

the slaves via an FHS packet transfer. Each slave will add an offset to its own clock to become synchronized with the master clock. Because the clocks are free running the offsets must be updated regularly. The use of an offset by a slave(s) provides for temporary Bluetooth clocks that are synchronized to the master clock.

10.5 BLUETOOTH LINK CONTROLLER OPERATIONAL STATES

The operational states of the Bluetooth link controller provide a systematic sequencing of the necessary operations that allow a Bluetooth device to enter into a piconet connection. Figure 10–16 depicts the different states used by the Bluetooth link controller. For a Bluetooth device there are two possible major states: standby and connection. Furthermore, there are seven substates: page, page scan, inquiry, inquiry scan, inquiry response, slave response, and master response. These substates are temporary states that are used to add new slaves to a piconet. A Bluetooth device will move from one state to another, either in response to commands from the Bluetooth link manager or in response to an internal trigger signal generated by the link controller.

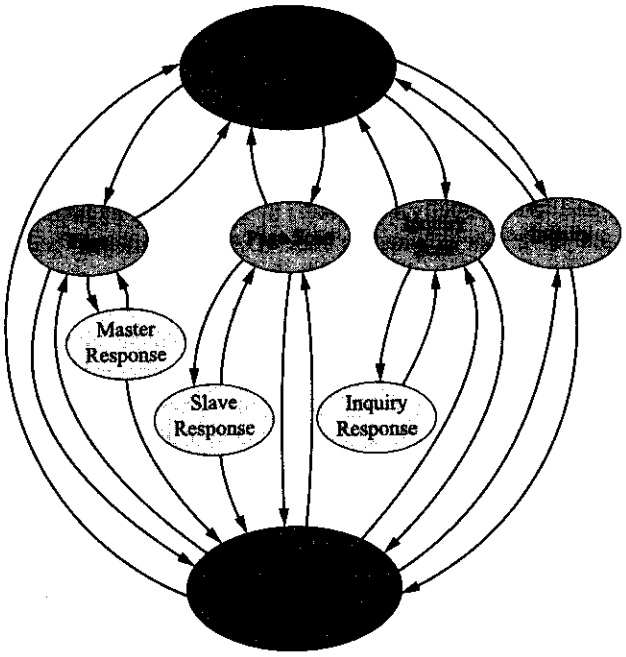

Figure 10–16 Bluetooth link controller operational states (Courtesy of IEEE).

No attempt will be made here to provide all the details of all the possible changes in state that can occur for a Bluetooth device. Instead, an overview of the most important aspects of how a device becomes either a master or a slave of an ad hoc piconet will be the focus of this section. When a Bluetooth device is first powered up, it goes into the default standby state and assumes a low-power mode while in this state. The Bluetooth link controller may leave the standby state to enter a page or inquiry mode or a page or inquiry scan mode. If the device enters the inquiry or page mode, it is taking the first steps in the process of becoming a piconet master while in the connection state. Conversely, if the device enters the page or inquiry scan modes it can end up as a piconet slave in the connection state. A device may return to any one of these four substates from the connection state as shown by the figure.

Bluetooth Access Procedures

If a Bluetooth device desires to form a piconet, the inquiry and paging procedures provided within these substates are used. The use of the inquiry operation allows a device to determine the presence of other Bluetooth-enabled devices. During the inquiry procedure, the inquiring unit is able to collect all the addresses and clock states of any devices within range that respond to the inquiry message. This inquiry message will contain no information about the source device but can provide an indication of which class of devices should respond to the inquiry. Recall that there is a general inquiry access code (GIAC) and a number of dedicated inquiry access codes (DIACs) (a total of sixty-three) that only ask certain types of devices to respond (printers, fax machines, etc.).

Once the inquiring device has obtained information about other units within its immediate environment, if desired, it can attempt to make a connection to any one of them by entering the paging substate. The inquiry operation itself consists of the continuous transmission of the inquiry message at different hop frequencies at double the normal rate. The frequency hopping sequence is derived from the GIAC value. A device that will allow itself to be discovered will regularly enter the inquiry scan substate and scan for the inquiry access code using the same frequency hopping sequence as derived from the GIAC value. Since the phases of the frequency hopping sequences for the devices are most likely different (and occurring at different rates), it may take some time before they finally match and the discovered device is able to transfer any data (in the form of an FHS packet) to the inquirer while in the inquiry response substate. Provisions are included in the standard for the resolution of possible contention problems that might occur between discovered Bluetooth units during this procedure.

With the paging procedure an actual connection (the creation of an ad hoc network) can occur. A default paging scheme is used when Bluetooth devices meet for the first time or in the case of device paging that occurs directly after an inquiry procedure. The paging substate is used by a master device to trigger the activation of a slave unit and to eventually connect to it. During the paging operation, the master device will repeatedly transmit the slave's device access code (DAC) at different hopping frequencies using the hopping sequence called for by the slave's device address until it receives a response from the slave device. Again, since the frequency hopping phase of the master and slave devices is most likely different, the master will use any clock information it might have about the slave's clock from either an inquiry procedure or a past connection to the slave device to offset its own clock in an attempt to synchronize its hopping with that of the slave. In any case, to make sure that the two devices finally communicate, the frequency hopping rate of the master is doubled. Moreover, when a slave receives a page and returns a response to the master during the response routine their clocks are "frozen" at the same hop frequency to facilitate the operation. This operation is illustrated by Figure 10–17.

In the page scan substate a Bluetooth device listens for its own device access code (DAC). Its frequency hopping sequence is determined by its own Bluetooth device address (BD_ADDR) and every 1.28s a new frequency (phase) is selected. When the device is triggered by the reception of its own DAC, it will enter the slave response substate. In this substate the slave will transmit a respond message that simply consists of its device access code. The slave then awaits the return of an FHS packet from the master device during the

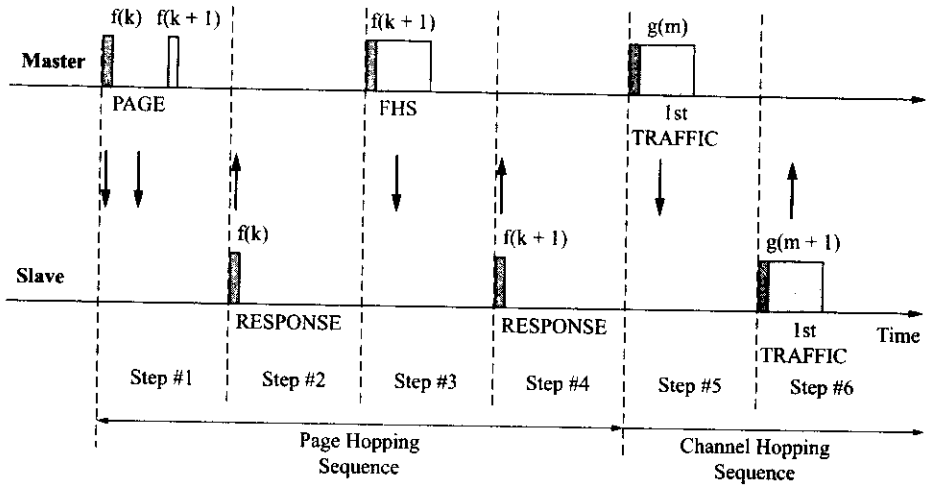

Figure 10–17 Bluetooth page and channel hopping sequences (Courtesy of IEEE).

next master-to-slave transmit timeslot. This FHS packet contains information that will allow the slave device to enter the connection state. During the connection state the piconet's channel access code and channel hopping scheme are derived from the master device's Bluetooth device address (BD_ADDR) and the piconet timing is determined by the master clock. As described previously, an offset is added to the slave's clock to temporarily synchronize it with the master device's internal clock. If no master device FHS packet is received by the slave before a timer expires, the slave returns to the page scan mode.

The master device performs several actions in response to a successful page. When a master receives a response message from a slave, it enters the master response substate. It transmits an FHS packet to the slave that contains the master device's clock value, 48-bit BD_ADDR, and its class of service. After the transmission of the FHS packet the master waits for a second response from the slave device that acknowledges the receipt of the FHS packet. If no second response is forthcoming, the master device repeats the transmission of the FHS until either a second slave response is received or a timer expires (at which point it returns to the page substate). If a second response is received from the slave, the master changes to its own parameters for the channel access code and the master clock. It then enters the connection state and uses its own BD_ADDR to derive the new channel hopping sequence that will be used by the piconet (step 5 in Figure 10–17). At this point, the master device transfers its first traffic packet, which happens to be a poll packet. If the slave does not receive the packet or return a packet within a certain time, both the slave and master return to the page scan and page modes, respectively.

Connection State

Once the connection state has been achieved, packets may be transferred back and forth between the master and slave(s). Again, the piconet uses the master's internal clock and BD_ADDR to define the piconet channel. After the poll packet has been sent by the master device, the first packets to be transferred contain control messages that are used to characterize the link and give more details about the Bluetooth units that are connected by the piconet. These messages are exchanged between the Bluetooth link managers (refer back to Figure 10–7) and typically define the type of data links (SCO or ACL) to be used and other details like sniff parameters. The connection state is left through the receipt of a detach or reset command message. The detach command is not as absolute as the reset command since all the configuration data in the Bluetooth link controller remains valid after a detach operation is performed. While in the connection state, a Bluetooth device is able to take on several different modes of operation. These are the active, sniff, hold, and park modes.

Active, Sniff, and Hold Mode

In the active state a device actively participates on the piconet channel. The master schedules traffic based on the demands of the piconet slaves. Active slaves listen to the channel for master-to-slave packets meant for them. An active slave that is not addressed may go into a low-power sleep mode and wake when the next transmission is scheduled. The slaves may use any packet transmitted from the master device to synchronize with the master.

Sniff mode may be invoked by either the master or the slave through the LM protocol. This mode is similar to SCO link operation except that it is performed with ACL links. In the sniff mode the slave does not have to listen to every ACL timeslot. It can be instructed to listen to the timeslots on a reduced periodic basis and therefore save battery power during this mode of operation.

During the connection state, the ACL link to a particular slave may be put on hold. This allows the slave to perform other operations like scanning, paging, inquiring, or participating in another piconet during this period. While in the hold mode the slave may go into a low-power sleep mode but it retains its active piconet address (AM_ADDR). Before going into the hold mode, the master and slave agree on the duration of the hold and set a timer that will be used to bring the slave out of the hold mode when it expires (i.e., the slave wakes up).

Park Mode

If a slave does not need to participate in a piconet but still wants to remain synchronized for possible future participation, it can enter the parked mode and go into a low-power mode. In this mode the slave gives up its active piconet member address (AM_ADDR) but receives two new addresses to be used while in the parked mode. One of these addresses is used for a master-initiated unpark and the other for a slave-initiated unpark. Parked slaves will wake up periodically to resynchronize and to listen for broadcast messages from the master device. The use of the parking mode essentially allows for an unlimited number of piconet slaves. Although only seven slave devices may be active at a time, through swapping techniques, the piconet can take on virtually any size. To facilitate the parking process the master provides a **beacon channel** or slot. The beacon channel is transmitted periodically and serves four purposes: the resynchronization of parked slaves, carrying messages about the beacon to parked slaves, carrying broadcast messages, and unparking parked slaves. The parking and unparking of a slave by the master is carried out through the use of LM protocol messages. A slave may request to be unparked through the transmission of an access request message to the master during a special access window time interval that typically occurs during the beacon slots. If the reader is interested in obtaining more detail about any of the operations discussed in this section, he or she is encouraged to obtain the latest version of the IEEE 802.15.1 specifications.

Scatternet Operation

It is possible for multiple piconets to exist in the same area. Because a different master exists for each piconet, the piconets will have their own channel access codes, channel hopping sequences, and phase determined by the particular master device of the piconet. However, as the number of piconets within an area grows, a likely consequence of this fact is a decrease in overall system performance. This is not unexpected since this is typical for frequency hopping schemes such as those employed by the Bluetooth standard.

The Bluetooth standard allows for the interconnection of piconets to form scatternets (refer back to Figure 10–5). In this mode of operation a master or a slave in one piconet may become a slave in another piconet by being paged by the master of the other piconet. Also, a slave in one piconet may page the master or the slave of another piconet. By default, this slave would become the master of the newly formed piconet. Therefore it is possible for a Bluetooth device to take on two modes of operation within a scatternet. Alternating between operation as a master and a slave on a time division multiplexing (TDM) basis, a device is actually able to participate in two or more piconets within a scatternet. To perform these types of

operations, the Bluetooth standard has incorporated a great deal of functionality within the Bluetooth devices that has just been touched on in this treatment of the specifications. Suffice to say, modes of device operation to facilitate the scatternet option exist, interpiconet communications are supported on a TDM basis, and appropriate master-slave/slave-master switching procedures exist. Since the participation in any piconet requires a device to use the associated master device address and proper clock offset, switching back and forth between several piconets requires setup time and therefore restricts the use of this technique to certain types of links (an SCO link using HV3 type packets, an ACL link, etc.). This mode of connection typically restricts the total possible Bluetooth device data transfer rate and the various possible combinations of data traffic.

Hopping Sequence Selection

The last topic that needs to be discussed is the Bluetooth system selection of the frequency hopping sequence. There are ten types of hopping sequences defined (five for 79-hop systems and five for 23-hop systems) within the standard. This section will confine itself to the five 79-hop sequences used in the United States and most of Europe. These sequences are:

- A page hopping sequence that consists of 32 unique wake-up frequencies that are equally distributed over the 79 channels with a period length of 32;
- A complementary page response sequence covering 32 unique frequencies that have a one-to-one correspondence with the page hopping sequence. The master and slave use different rules to obtain the same sequence;
- An inquiry sequence with 32 unique wake-up frequencies also distributed equally over the 79 channels with a period length of 32;
- A complementary inquiry response sequence of 32 frequencies;
- A channel hopping sequence with a very long period length that does not display repetitive patterns over a short time interval.

The general scheme used to select the hopping sequence is shown in Figure 10–18. The inputs to the selection unit are the device system clock and the current address. In the connection state the slave's current clock value is offset by an amount necessary to equal the 27 most significant bits (MSBs) of the master's clock. In the page and inquiry substates, all 28 bits of the clock are used as an input to the selection unit. For a device in the page substate, the master clock will be offset to a value equal to the master's best guess of the paged device's clock value. The address input to the selection unit is derived from the Bluetooth device address. In the connection state the master's BD_ADDR is used, in the page substate the BD_ADDR of the paged device is used, and in the inquiry state the address corresponding to the GIAC is employed. The output from the selection unit is a pseudorandom sequence covering the seventy-nine possible hopping channels.

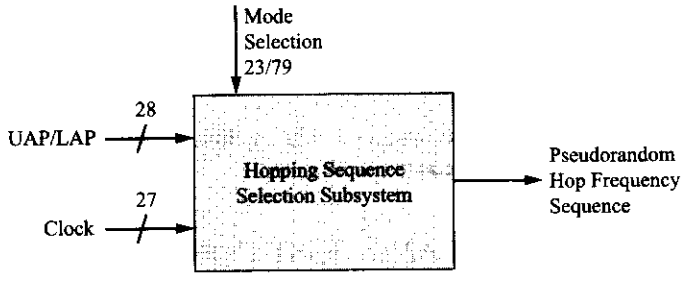

Figure 10–18 Selection of Bluetooth hopping sequence (Courtesy of IEEE).

The selection unit, through a complex algorithm, chooses a 32-hop frequency segment that spans about 64 MHz and visits these hops once in random order. Then another 32-hop segment is chosen and another and so on. In the case of the page, page scan, or page response substate, the same 32-hop segment is used over and over with each page being a unique 32-hop sequence since its value depends upon the unique paged device's address (BD_ADDR). In the connection state, the output of the selection unit consists of a pseudo-random sequence produced from 32-channel (64 MHz) segments of the 79-channel register that is organized as shown in Figure 10–19. Each successive 32-channel segment is offset from the prior 32-channel segment by 16 channels or 32 MHz.

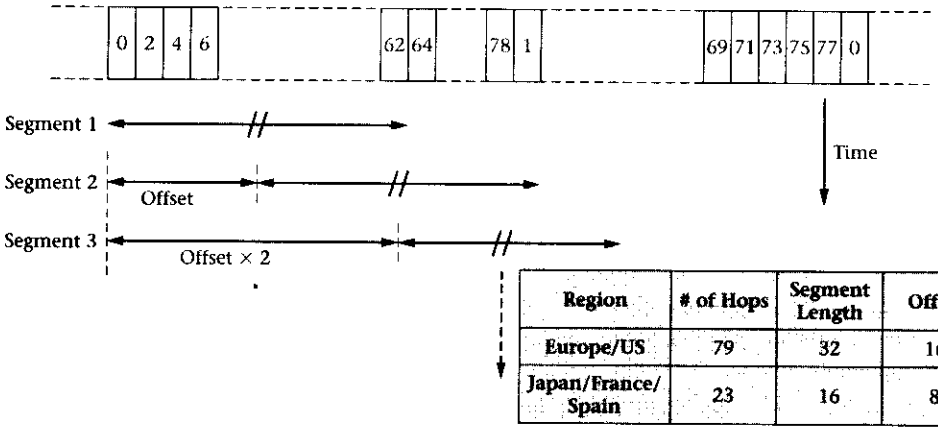

Region	# of Hops	Segment Length	Offset
Europe/US	79	32	16
Japan/France/ Spain	23	16	8

Figure 10–19 Further details of the operation of the hopping sequence selection subsystem (Courtesy of IEEE).

Bluetooth Addresses and Encryption

Each Bluetooth device is allocated a unique 48-bit Bluetooth address known as BD_ADDR. Part of the address is assigned by the manufacturer and part of it consists of the manufacturer's ID. This address is segmented by the Bluetooth standard into three parts that are used at various times by the Bluetooth system during various system operations. Figure 10–20 shows the address format with the lower address part (LAP), the upper address part (UAP), and the nonsignificant address part (NAP) labeled in the figure.

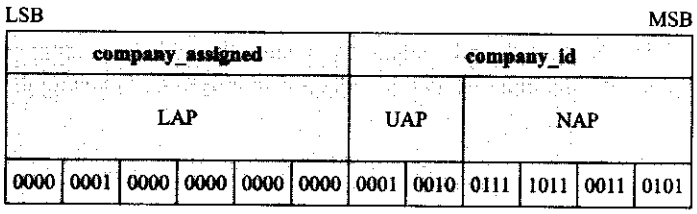

Figure 10–20 Bluetooth device address format (Courtesy of IEEE).

The Bluetooth standard provides for various levels of security through the use of complex authentication and encryption procedures that will not be discussed further here.

10.6 IEEE 802.15.1 PROTOCOLS AND HOST CONTROL INTERFACE

A great deal of detail about the Bluetooth RF physical layer, the actual physical data links, and the operations and possible states of the Bluetooth data link controller have been presented to the reader of this chapter. It is

hoped that this detail has given the reader a clear sense of the basic operation of Bluetooth-enabled devices, their physical limitations, and an insight to their vastly untapped potential future uses. To round out this coverage of IEEE 802.15.1, some discussion of the Bluetooth-specific protocols and the host control interface entity that reside within the protocol stack is appropriate. Earlier in this chapter, Figure 10–7 was introduced. This figure depicts the Bluetooth protocol stack and the relationships of the two Bluetooth-specific protocols: link manager protocol and logical link control and adaptation protocol with both the physical layer and the (host) control interface unit. At various times during the discussion of the physical layer and the data link controller, references have been made to operations that have involved these protocols and the control interface portion of the protocol stack. This section will give a short overview of the operation of these two protocols and the host interface control unit and their relationship to the overall operation of the Bluetooth system.

Link Manager Protocol

Within the Bluetooth specification, the link manager protocol (LMP) is used for link setup, security, and control. The various LMP messages are sent from either the piconet master to the slave or from the slave to the piconet master. In each case, they are interpreted and filtered out by the link manager on the receiving side and therefore are not transferred to higher layers of the protocol stack. LMP messages are transferred as protocol data units (PDUs) via the physical link in the payload section of a packet instead of L2CAP data. LMP messages have higher priority than user data and are always sent as single-slot packets. Each LM PDU consists of a 7-bit opcode that distinguishes the PDU and a 1-bit transaction identity bit that indicates the source (i.e., master or slave) of the transaction. The many possible PDU transaction parameter values are carried in additional bytes. Typically, either a DM1 or DV packet will be used to carry the PDU.

There are many different types of transactions supported by LMP. A listing of all the various transaction types will be given here and then several examples of typical procedures will be presented to give the reader a sense of how these procedures are used. The various transaction types are authentication, pairing, change link key, change the current link key, encryption, clock offset request, slot offset information, timing accuracy information request, LMP version, supported features, switch of master-slave role, name request, detach, hold mode, sniff mode, park mode, power control, channel quality-driven change between DM and DH, quality of service (QoS), SCO links, control of multislot packets, paging scheme, link supervision, and connection establishment. There are also some additional LMP PDUs used to support a Bluetooth device test mode. In each case, there are usually several subcategorizes to the primary LM transaction type. These various LM procedures are used during the establishment of a piconet, during the operation of a piconet, during testing of a piconet device, and during the dissolution of a piconet. Two representative examples of LMP transactions are given next.

Example 10–1: Encryption Mode

For encryption to be used, both the master and slave must agree upon both its use and the extent of its use (i.e., only for point-to-point packets or for both point-to-point and broadcast packets). If both the master and slave agree on the encryption mode, then the master will provide more detail about the encryption parameters to the slave. This transaction takes place in the following fashion as shown by Figure 10–21. The link manager that initiates the transaction finishes the transmission of the current ACL packet (that contains user data) and sends the LMP_encryption_mode_req message. Depending upon whether the change is accepted, the other Bluetooth device finishes the transmission of the current ACL packet and responds with either the LMP_accepted or LMP_not_accepted message.

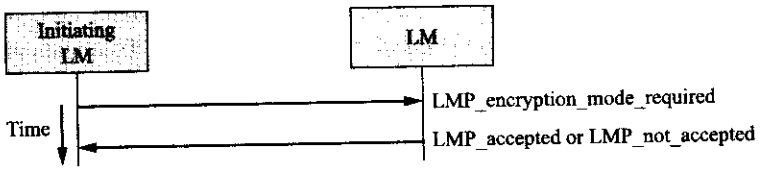

Figure 10–21 Messages exchanged during a Bluetooth encryption mode transaction using LMP (Courtesy of IEEE).

Example 10–2: Power Control

If the RSSI value differs too much from the preferred value for a Bluetooth device and if the other device supports it, the Bluetooth device can request a power adjustment. It does this by transmitting either the LMP_incr_power_req or the LMP_decr_power_req messages as shown in Figure 10–22. Depending upon the status of the other device, several message may be returned in response to this request. In any case, in this release of the standard, the device will change its power only one step per power change request.

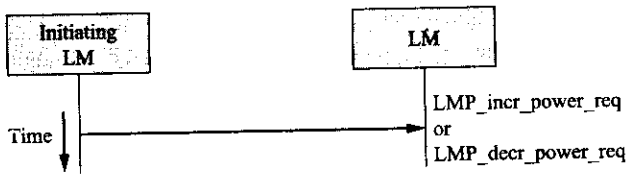

Figure 10–22 Messages exchanged during a Bluetooth power control transaction using LMP (Courtesy of IEEE).

Logical Link Control and Adaptation Protocol

The logical link control and adaptation protocol (L2CAP) is used to provide support for higher-level protocol multiplexing, packet segmentation and reassembly, the transfer of QoS information, and group abstractions. The L2CAP is layered over the baseband protocol and provides both connectionless and connection-oriented data services over ACL links to the layers above it. Furthermore, the L2CAP layer supports the transmission and reception of L2CAP data packets that are up to 64 kilobytes in length. Figure 10–23 illustrates how L2CAP interfaces with these other protocols. This section will provide a short overview of L2CAP features and its operation.

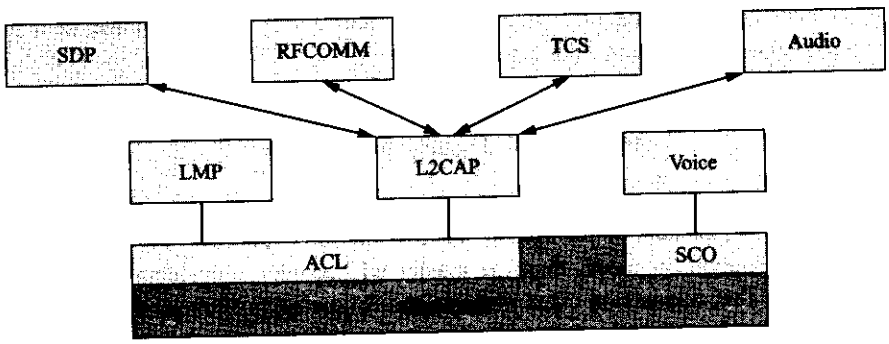

Figure 10–23 Interface of Bluetooth L2CAP protocol with other higher-level protocols (Courtesy of IEEE).

To carry out the various tasks just listed, L2CAP must support the following operations:

- Protocol Multiplexing—Since the baseband protocol does not support a type field that identifies higher-layer protocols, the L2CAP layer performs this operation.

- Segmentation and Reassembly—Since baseband data packets are limited in size (341 bytes for a DH5 packet), the L2CAP layer provides this service to upper-layer protocols that permit larger packet sizes.
- Quality of Service—The L2CAP connection establishment process allows for the exchange of QoS information between two Bluetooth devices. The L2CAP layers at either end of the link monitor the QoS during the link activation.
- Groups—Since many protocols support the concept of group addresses and the baseband protocol supports the piconet concept, L2CAP group abstraction is used to map higher-layer protocol groups onto the devices of a piconet.

The operation of L2CAP is based on a logical channel concept. Each one of the end points of an L2CAP channel has its own channel identifier (CID). Figure 10–24 shows how CIDs are used between peer L2CAP entities in separate Bluetooth devices. For connection-oriented data links (that represent a connection between two devices), each end point has a CID. For connectionless links, data flow is restricted to a single direction. These types of connections are used to support a group of devices and in these cases a CID at a source may represent one or more other remote devices. There are several CID values reserved for special purposes. A signaling channel is an example of a reserved channel. A signaling channel is used to provide the establishment of a connection-oriented data channel and to provide the ability to change the characteristics of this channel.

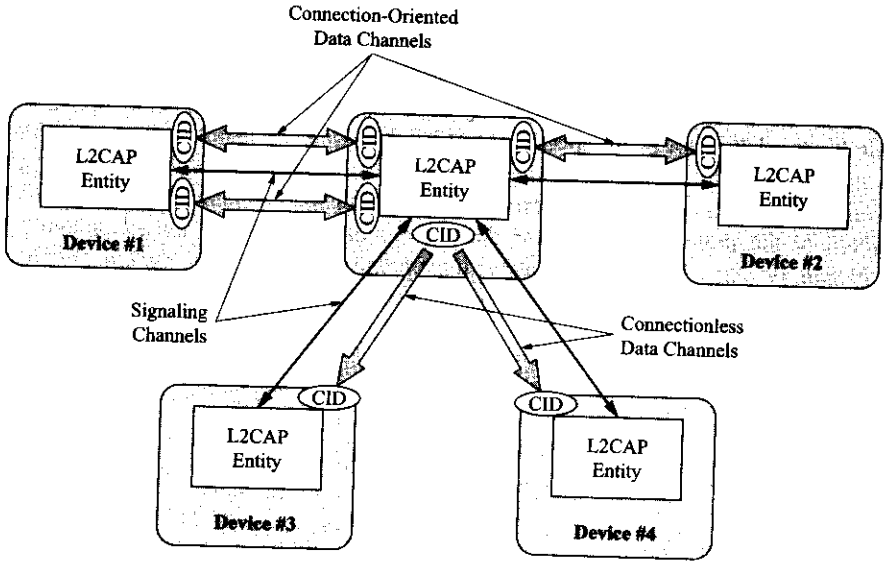

Figure 10–24 Bluetooth channel identifiers (Courtesy of IEEE).

The L2CAP layer should be able to transfer data between higher-layer protocols and lower-layer protocols, provide various services to these other layers, provide signaling commands between peer L2CAP implementations in other devices, and react to events from lower layers by passing an indication of these event occurrences on to higher layers. The detailed operation of how the L2CAP layer accomplishes these tasks will not be considered here but the reader is certainly encouraged to explore L2CAP operation in more detail by consulting the latest IEEE 802.15.1 standard.

Host Control Interface

The last topic that needs to be discussed is the host control interface (HCI) or control interface as labeled in Figure 10–7. The Bluetooth HCI provides a command-level interface to the baseband controller and link

manager. Furthermore, the HCI has access to Bluetooth hardware control and status registers. The IEEE 802.15.1 control interface specifications are based totally upon those outlined in the Bluetooth specifications for the HCI function. The HCI interface provides a consistent method of integrating the functionality of the Bluetooth baseband capabilities into a host digital device. The HCI unit provides two basic functions within the Bluetooth specification:

- Describes how a Bluetooth module may be physically (electrically) interfaced to a device
- Describes the necessary control functions for Bluetooth implementations

This last function is described by the IEEE 802.15.1 standard. A short description of this function will conclude the coverage of this topic for this chapter.

The HCI offers the host device the ability to control the link-layer connections to other Bluetooth-enabled devices through LMP commands exchanged with remote devices. HCI policy commands are used to affect the operation of both local and remote link managers and how they manage a particular piconet. The host device has access to various host controller registers through an assortment of HCI commands. Additionally, the host device receives immediate real-time notifications of HCI management events and therefore has knowledge of when various events have occurred or have not occurred. The user of the host device may react accordingly.

10.7 EVOLUTION OF IEEE 802.15 STANDARDS

When the IEEE 802.15 standard was first being considered, there were certain applications for wireless PANs that could not be addressed under one comprehensive standard. Therefore, in an effort to adequately address other application areas and in acknowledgement of the probable rapid technologic advances that would provide higher data throughput speeds, the IEEE 802.15 standards have been subdivided into four separate but related standards. Three of these standards deal specifically with different application areas whereas one (IEEE 802.15.2) deals with interoperability issues with other wireless applications (i.e., IEEE 802.11/WLANs). At the same time, follow-on derivative work has already started on some of these standards areas. Figure 10–25 graphically illustrates the operating space of the various WPAN and WLAN standards. This section will attempt to provide an overview of the increasing interest and activity in the wireless PANs space.

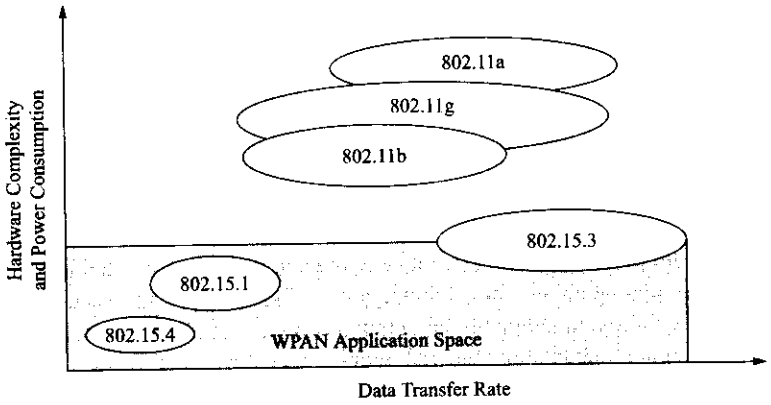

Figure 10–25 Application space of the various WLAN and WPAN standards (Courtesy of IEEE).

IEEE 802.15.1

The IEEE 802.15.1 standard has already been described in great detail. It provides for short-range wireless connectivity for personal devices at moderate data rates and supports high-quality voice connections. As

mentioned earlier, the Bluetooth SIG has already developed a follow-on Bluetooth Specification 1.2. In response, the IEEE 802.15 working group has put into place a revision project that will result in IEEE 802.15.1a. This revision will provide complete backward compatibility with the present standard and at the same time incorporate the functional changes provided in the new Bluetooth specifications. Additionally, new features and improvements to both the physical and MAC layer will be included in the revision. Another important goal of this working group is to move toward a level of device interoperability that would allow a WPAN device and a WLAN to exchange data.

IEEE 802.15.2

IEEE 802.15.2-2003 is a revised project that has as its ultimate goal the facilitation of the coexistence of IEEE 802.15x devices and other devices that both use the same unlicensed frequency spectrum. Specifically included in these other devices are IEEE 802.11 WLANs. The working group involved with this standard has been looking at two basic types of technologies to achieve the goals of coexistence. They are known as non-collaborative and collaborative mechanisms. The first method, as can be deduced from its name, does not depend on any cooperation or interoperability between the interfering units. At this time, the favored approach is to invoke a form of adaptive frequency hopping that would be employed by the Bluetooth device. This technique would assign a classification of either good or bad to hop channels. These ratings might be achieved by any number of different techniques but once rated as a bad channel (high BER or FER due to EM interference) it would be dropped out of the device's frequency hopping sequence. LMP messages between a master and slave would be used to invoke this mode of operation within a piconet and to exchange data about channel status.

A collaborative mechanism involves the exchange of information between an IEEE 802.11 device and an IEEE 802.15 device. Several different proposals to implement this type of coexistence mechanism have been suggested. In all cases the two systems have knowledge of each other's operation. One scenario provides for a tightly coordinated queuing and scheduling algorithm that would dynamically adapt to the traffic type and manage the operation of a single WLAN station and a single piconet. Another scenario is to use a form of time division multiple access (TDMA) that provides dedicated times within a WPAN timeslot for either WPAN traffic or WLAN traffic. At present, the general consensus is to combine all of these approaches. The IEEE 802.15.2 working group is in the draft stage of the formulation of the standard. What will finally be adopted as the new standard remains to be seen at this time.

IEEE 802.15.3 and IEEE 802.15.3a

The IEEE 802.15.3 standard was adopted during late 2003. It provides for low cost and complexity, low power consumption, and high-data-rate (20 mbps or more) wireless connectivity of devices within or entering a POS. The original IEEE 802.15 standard did not provide the necessary higher data transfer rates to satisfy a number of multimedia industry needs for WPAN communications. IEEE 802.15.3 provides for WPAN-HR (WPAN high rate) data transfers with QoS support capabilities. The IEEE 802.15.3a working group is looking at a follow-on standard that will raise the data transfer rate to 110 mbps or more through the use of either ultra-wideband technology (refer back to Chapter 8) or some other new transmission technology. Through the use of new transmission techniques, it is felt that extremely high data transfer rates may be achieved that would open the door to many new and novel multimedia applications within this WPAN space. At this time a short overview of the IEEE 802.15.3 standard will be offered with emphasis on the differences between it and IEEE 802.15.1.

IEEE 802.15.3 Physical Layer

The key changes embodied in IEEE 802.15.3 are the higher data transfer rates supported by the standard and the different modulation schemes used to achieve these rates. The single-carrier data transfer rates enabled by the standard include 11, 22, 33, 44, and 55 mbps (see Table 10–5). These rates are achieved

Table 10–5 IEEE 802.15.3 modulation schemes and data rates (Courtesy of IEEE).

Modulation Type	*Coding*	*Data Rate*
QPSK	8-state TCM	11 mbps
DQPSK	none	22 mbps
16-QAM	8-state TCM	33 mbps
32-QAM	8-state TCM	44 mbps
64-QAM	8-state TCM	55 mbps

Table 10–6 IEEE 802.15.3 carrier frequency allocations (Courtesy of IEEE).

CHNL_ID	*Center Frequency*	*High-Density*	*802.11b Coexistence*
1	2.412 GHz	X	X
2	2.428 GHz	X	
3	2.437 GHz		X
4	2.445 GHz	X	
5	2.462 GHz	X	X

through the use of complex digital modulation schemes and the use of Trellis coding techniques. In all cases, the symbol transmission rate is 11 msps, with the modulation schemes affording higher data transfer rates. Presently, a total of five channels in the 2.4-GHz unlicensed spectrum are assigned for operation (see Table 10–6). Three of these channels are designated as 802.11b coexistence channels and four are from a high-density application set. Signal bandwidth is limited to 15 MHz and output transmitter output levels are restricted to the milliwatt range.

The IEEE 802.15.3 frame format for the 22, 33, 44, and 55 mbps transmission rates is shown in Figure 10–26. The entire transmitted frame consists of a preamble, physical layer header, MAC header, a header check sequence (HCS), the frame check sequence (FCS) and frame payload, stuff bits, and tail bits as shown in the figure. The header portion of the frame is transmitted at 22 mbps (using both QPSK and DQPSK modulation) and then for the payload portion of the frame the necessary modulation scheme is implemented to achieve the desired data rate (i.e., 22, 33, 44, or 55 mbps). The frame format for the 11-mbps transmission mode is slightly different but will not be discussed here.

IEEE 802.15.3 Piconets

An IEEE 802.15.3 piconet consists of several devices (DEVs). One device assumes the role of the **piconet coordinator** (PNC) and as such provides the basic piconet timing through the use of a beacon frame. The PNC also is responsible for managing the QoS requirements of the piconet, power saving modes of the DEVs, and control of access to the piconet by other devices. The IEEE 802.15.3 standard allows for a DEV to request the formation of a **dependent piconet.** When there are other dependent piconets, the original piconet is referred to as the **parent piconet** and the dependent piconets are referred to as either **child** or **neighbor piconets** depending upon how the DEV that formed them associated with the parent piconet.

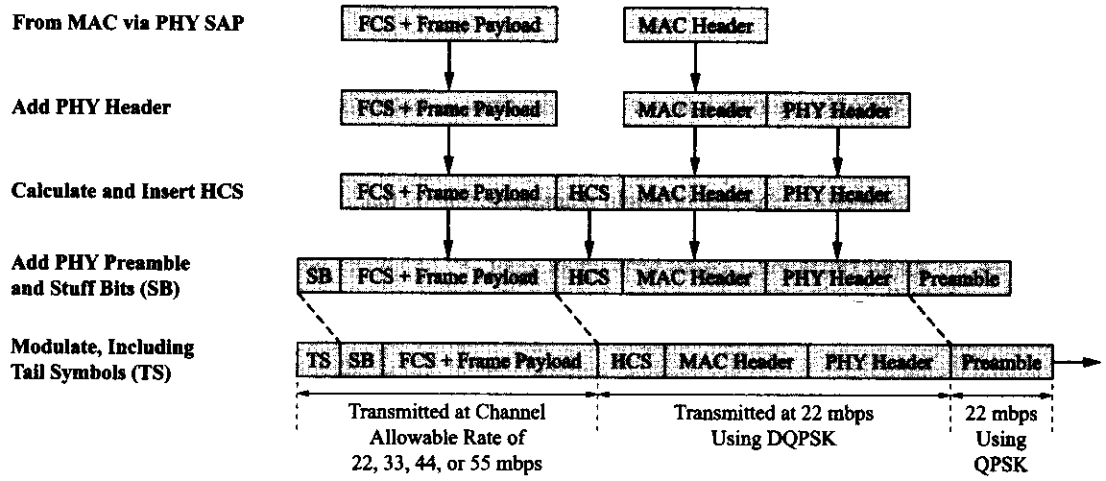

Figure 10–26 IEEE 802.15.3 frame format (Courtesy of IEEE).

A DEV that wants to start a piconet must first scan the available channels to see if there are beacon frames from any existing PNCs. The scanning DEV collects statistics about each channel including information about any parent, child, or neighbor piconets that were detected and then rates the channels for their suitability for the start of a new piconet. Then the DEV listens to the best candidate channel for a certain length of time and if the channel is still clear it now becomes a PNC by commencing to broadcast a beacon once every superframe time period. A **superframe** consists of the beacon frame, a contention access period (CAP), and the channel time allocation period (CTAP). Figure 10–27 shows the basic superframe format. The beacon is used to set timing allocations and to communicate management information about the piconet. The CAP time period is used to communicate commands and data by either DEVs or the PNC. Access to this time period is through the use of a CSMA/CA scheme with backoff algorithms. During the CTAP portion of the superframe, the PNC controls channel access by assigning channel time allocations (CTAs) to an individual DEV or a group of DEVs. Using TDMA, all the CTAs have a fixed start time and duration. The PNC determines the allocation of the CTAs to the DEVs within a piconet. For child or neighbor piconets, a portion (or possibly a longer period) of the parent superframe is dedicated to use by the child or the neighbor piconet if the parent piconet so allows it.

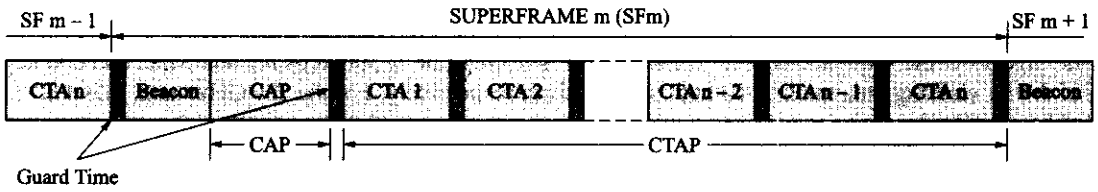

Figure 10–27 IEEE 802.15.3 superframe format (Courtesy of IEEE).

As the reader may surmise, there are detailed procedures spelled out in the IEEE 802.15.3 standard (over 300 pages worth) for starting, joining, leaving, and stopping a piconet, channel access and time management, synchronization, power management, security, encryption, and so forth. The goal of this short section has been to present an overview of IEEE 802.15.3, which has been accomplished. For the interested reader, the latest version of the IEEE 802.15.3 standard should be consulted for additional operational details.

IEEE 802.15.4

The IEEE 802.15.4 standard addresses the low-rate WPAN (LR-WPAN) application space. The principal characteristics of LR-WPANs are data transfer rates less than or equal to 250 kbps, ultralow power consumption, a small form factor, and low cost and complexity (refer back to Figure 10–25). There are many proposed applications for these types of WPANs including those involving wireless sensor networks (WSNs). An interesting Web site the reader might want to explore is that of the ZigBEE Alliance (an association of manufactures with an interest in IEEE 802.15.4) at www.zigbee.org. The goal of this section will be to point out the major operational characteristics of IEEE 802.15.4 but not a discussion of its applications. Chapter 13 dealing with emerging technologies will discuss the topic of WSNs in the context of IEEE 802.15.4 in more detail.

IEEE 802.15.4 Physical Layer

The IEEE 802.15.4 standard calls for operation in three different unlicensed frequency bands: the 868–868.6 MHz band using DSSS to provide 20-kbps data transfer rates, the 902–928 MHz band using DSSS to provide 40-kbps data rates, and the 2.4-GHz band using DSSS to provide 250-kbps data rates. The standard specifies only one channel in the 868-MHz band, ten channels in the 915-MHz band, and sixteen channels in the 2.4-GHz band. Furthermore, the use of these bands is not universal. The 868-MHz band is limited to Europe, the 915-MHz band to the Americas, and the 2.4-GHz band is used by most of the world's countries. Table 10–7 indicates various data transfer rates and modulation format details of the operation of the IEEE 802.15.4 standard in these various bands.

Table 10–7 IEEE 802.15.4 frequency bands and data transfer rates (Courtesy of IEEE).

Band (MHz)	*Frequency Band*	*Bit Rate (kbps)*	*Symbol Rate (kbps)*	*DSSS Spreading Parameters*	
				Modulation Technique	*Chip Rate*
868	868–868.6 MHz	20	20	BPSK	300 kcps
915	902–928 MHz	40	40	BPSK	600 kcps
2400	2400–2483.5 MHz	250	62.5	O-QPSK	2 mcps

The IEEE 802.15.4 standard calls for a DSSS scheme employing differential BPSK in the 868- and 915-MHz bands. A single 15-chip pseudo-random sequence is transmitted in a symbol period to represent a 1 and the inverse of the sequence is used to encode a 0. The chip rate is 300 kcps for 868 MHz and 600 kcps for 915 MHz, yielding data rates of 20 kbps and 40 kbps, respectively. For the 2.4-GHz band, offset QPSK (OQPSK) modulation with a quasi-orthogonal spreading scheme is used. Every 4 bits of data to be transmitted is spread/encoded by a 32-bit chip code. The 32-bit code is split into two 16-bit streams that are applied to the I (even bits) and Q (odd bits) channels of the OQPSK modulator with a one-half chip delay in the Q channel to provide the offset for the OQPSK. The OQPSK modulator encodes 2 bits every 1 μs; thus the entire process takes 16 μs per symbol providing a rate of 62.5 ksps. With 4 bits encoded per symbol this yields a data transfer rate of 250 kbps. An IEEE 802.15.4 transceiver must be capable of an output power of at least –3 dBm, with higher output powers limited by the regulatory body in charge of the particular geographic location. The receiver portion of the transceiver must be capable of a sensitivity of –85 dBm in the 2.4-GHz band and –92 dBm in the lower bands.

IEEE 802.15.4 Piconets

The MAC layer of IEEE 802.15.4 provides for the support of two wireless network topologies: a star and a peer-to-peer topology. Figures 10–28 to 10–30 illustrate these various LR-WPAN topologies. Home

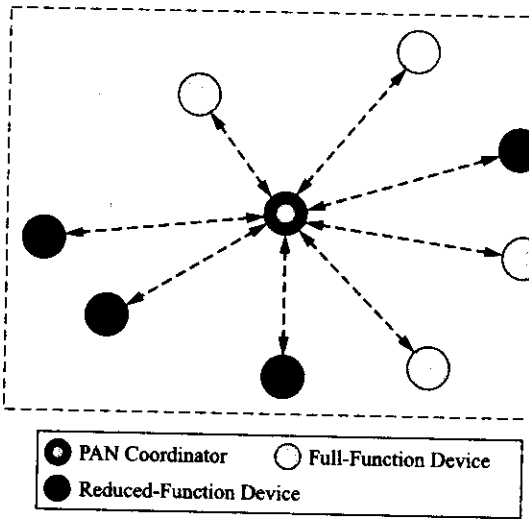

Figure 10–28 IEEE 802.15.4 LR-WPAN star topology (Courtesy of IEEE).

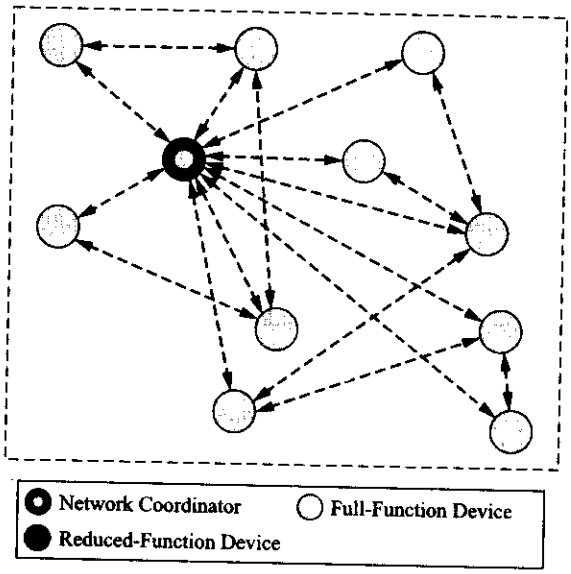

Figure 10–29 IEEE 802.15.4 peer-to-peer network topology (Courtesy of IEEE).

applications would typically employ the star structure whereas industrial and commercial applications have driven the strategy of the peer-to-peer structure. In the star structure, all communications within the network are controlled by a unique **PAN coordinator.** This PAN coordinator acts as the network master, transmits beacon frames for device synchronization, and maintains the association management status of the other devices within the network. Only a full-function device (FFD) capable of transmitting a beacon frame may become a PAN coordinator; however, reduced-function devices (RFDs) may participate in star networks. Furthermore, a star network operates independently of any other IEEE 802.15.4 networks. An FFD may establish a star network after performing a channel scan. If the FFD does not detect any transmitted beacon frames, it may begin operating as a PAN coordinator by sending beacon frames that contain a unique network identifier or ID. All devices participating in LR-WPANs use their unique IEEE 64-bit

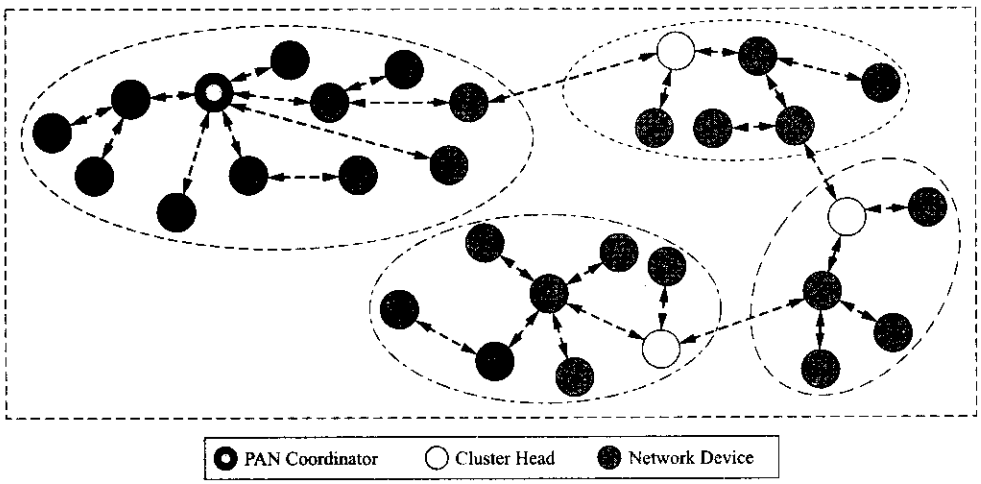

Figure 10–30 IEEE 802.15.4 cluster-tree network topology (Courtesy of IEEE).

addresses. Once the device has started to send beacon frames, other devices may ask to associate with it (i.e., send an association request message) thus forming an ad hoc network.

The peer-to-peer network organization (see Figure 10–29) allows any FFD to communicate with any other FFD within its range and to also have messages relayed to FFDs outside its range. This type of topology enables more complex ad hoc wireless networks with added coverage areas due to multihop and mesh network configurations that allow the functionality of message relaying. RFDs may participate in peer-to-peer networks but they are unable to act as relays.

A type of peer-to-peer network known as a cluster-tree network is also possible (refer to Figure 10–30). In this network organization, a number of network devices can take on the role of "cluster heads." This structure provides path redundancy between various network devices and the ability to have the network span a greater area. In any case, the peer-to-peer network must have one device that acts like the PAN coordinator. During the formation of a peer-to-peer network, a discover phase of operation allows the network devices to determine the features and services that are supported and available from each other.

The IEEE 802.15.4 standard supports a superframe structure that is managed by a PAN coordinator. The superframe format is shown by Figure 10–31. The superframe starts with the transmission of a beacon

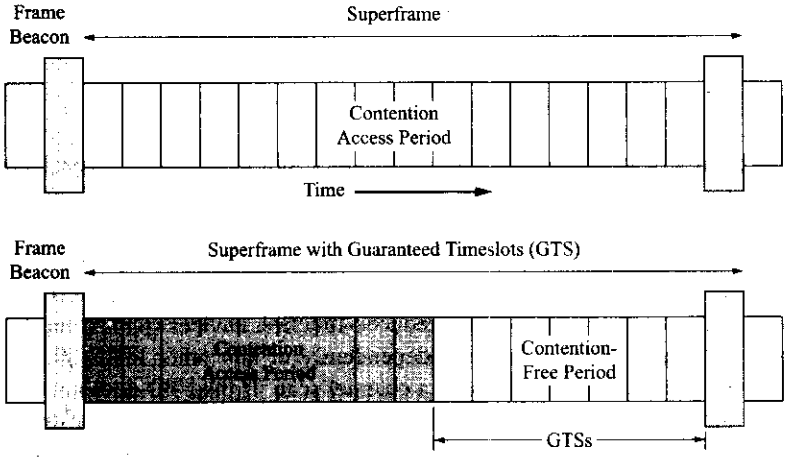

Figure 10–31 IEEE 802.15.4 superframe structure (Courtesy of IEEE).

frame that is used by devices to synchronize to the network, provides the network ID, and information about the superframe structure. The superframe is divided into sixteen timeslots that provide a contention access period (CAP). Using a CSMA/CA scheme, network devices attempt to communicate with the PAN coordinator during this time period. The network coordinator can assign dedicated portions of the superframe to a requesting network device. Known as guaranteed timeslots (GTS) these segments of time allow for particular bandwidth requirements or QoS requirements. The GTS slots are placed at the end of the superframe and form the contention-free period.

The physical protocol data unit that is passed over the physical interface is shown in Figure 10–32. It consists of a synchronization header, physical layer header, and the MAC PDU. The MAC protocol data unit consists of a MAC header, MAC payload, and MAC footer. The MAC header contains a frame control field and an addressing field. The frame control field provides information about the type of frame, the format and content of the address field, security information, and whether an acknowledgement is required. The MAC header addressing field contains the source or destination address. The MAC payload contains information germane to the type of MAC transaction being performed and it can also be further subdivided if necessary. The MAC footer consists of a 16-bit frame check sum (FCS). There are four types of MAC frames that can be transferred: beacon, command, data, and acknowledgement.

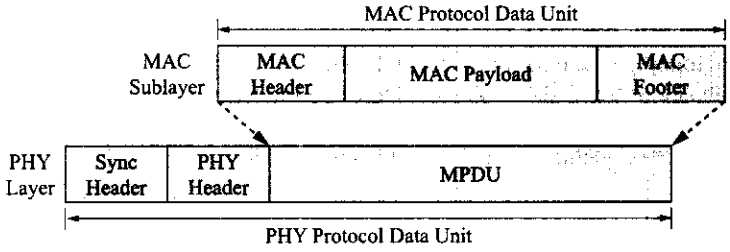

Figure 10–32 IEEE 802.15.4 physical protocol data unit (Courtesy of IEEE).

As was the case for IEEE 802.15.3, the IEEE 802.15.4 standard adopted during mid-2003 contains a great deal of detail about the operation and various procedures used to create, maintain, secure, and dissolve a LR-WPAN. As before, the goal of this section was to provide an overview of the standard. For the interested reader, the latest version of the IEEE 802.15.4 standard should be consulted to obtain additional details about system operation.

QUESTIONS AND PROBLEMS

1. Define a wireless personal area network.
2. Define a personal operating space in the context of a wireless PAN.
3. Describe the basic difference between a wireless LAN and a wireless PAN.
4. What is the Bluetooth standard?
5. Describe a basic application of wireless PANs.
6. Contrast the transmitting power and range for a wireless LAN versus a wireless PAN.
7. Define the basic functions of the master and slave elements of a wireless PAN.
8. Contrast the life spans of wireless LANs and wireless PANs.
9. Describe a wireless PAN piconet.
10. Describe a wireless PAN scatternet.
11. How is full-duplex operation supported by the Bluetooth standard?
12. What is a wireless PAN Class 3 device?
13. What is the present maximum power output of a wireless PAN device?
14. Describe wireless PAN power control.

15. How does the Bluetooth system support both circuit-switched and packet-switched data?
16. Describe the Bluetooth SCO link.
17. Describe the Bluetooth ACL link.
18. What information is encoded by the "TYPE" field of a packet header?
19. What is the maximum symmetric data rate allowed by the Bluetooth standard?
20. What are the maximum and minimum asymmetric data rates allowed in the forward direction by the Bluetooth standard?
21. How is timing synchronization achieved in a Bluetooth system?
22. Describe the basic steps needed to setup a piconet.
23. In what state must a wireless PAN device be in before packets may be exchanged?
24. What is the function/purpose of the Bluetooth "sniff" mode?
25. For a wireless PAN device, describe the "hold" mode.
26. For a wireless PAN device, describe the "park" mode.
27. How is a Bluetooth device uniquely identified?
28. What is the function/purpose of the link manager protocol?
29. What function/purpose does the host control interface provide?
30. What application area(s) does IEEE 802.15.3 address?
31. What is a "piconet coordinator" as defined in IEEE 802.15.3?
32. What application area does IEEE 802.15.4 address?
33. What network topologies are supported by IEEE 802.15.4?
34. Go to the Web site of the ZigBEE Alliance. Describe the type of companies involved with this wireless PAN technology.
35. How is QoS supported within the IEEE 802.15.4 standard?

CHAPTER

11

Broadband Wireless MANs/IEEE 802.16x

Upon completion of this chapter, the student should be able to:

- Describe the relatively short history of the IEEE 802.16 standard.
- Explain the basic differences between wireless MANs, WLANs, and WPANs.
- Describe the changes in the regulatory environment that have provided for increased interest in broadband wireless access technology.
- Describe the typical wireless MAN deployment scenario.
- Explain the basic operation of the IEEE 802.16 MAC layer.
- Describe the basic operation of the IEEE 802.16/802.16a physical layers including their frame structure.
- Explain the difference between WMAN point-to-multipoint and mesh operation.
- Describe the basic WMAN operations of initialization, uplink scheduling, bandwidth requests, and radio link control functions.

This chapter introduces the IEEE 802.16x standard for wireless metropolitan area networks (WMANs). A brief overview of the current state of the standard is given and related to the use of different frequency allocations and classes of licensed and license-exempt frequency bands. The typical use and deployment of a WMAN is discussed next and contrasted to the operation and use of WLANs and WPANs.

A fairly rigorous coverage of the wireless MAN's MAC and physical layers is provided next. Details about the transport of ATM cells and packet data are provided in the context of the convergence sublayers and the MAC common part sublayer. The connection-oriented operation of the MAC is explained and related to typical MAC management messages and operations. The physical layer IEEE 802.16 specifications for the 10–66 GHz range and the IEEE 802.16a specifications for the 2–11 GHz range are presented with sufficient detail to allow the reader to obtain a good grasp of their operation as well as their differences. Duplexing techniques, framing, power control, modulation, diversity schemes, frame structure, and mesh network operation are all discussed for the different air interface access and modulation schemes employed by the standard.

Finally, common WMAN system operations like initialization, ranging, bandwidth allocation, and radio link control are discussed.

11.1 INTRODUCTION TO WMAN/IEEE 802.16X TECHNOLOGIES

The IEEE 802.16x standards provide the details of the physical and MAC layers for fixed point-to-multipoint broadband wireless access (BWA) systems that can be used to provide multiple types of data services to

system subscribers. The MAC sublayer is structured to provide support for multiple physical layer implementations over a broad range of frequencies in the **microwave** and **millimeter wave** regions. The original IEEE 802.16 project was started in 1999 in an effort to promote the use of innovative and cost-effective broadband wireless products on a worldwide basis. The first specifications provide for the transport of data, video, and voice services at frequencies in the range of 10 to 66 GHz.

It should be pointed out that several manufacturers already provided equipment for local multipoint distribution service (LMDS) in the 10–40 GHz range during the late 1990s and sold this equipment on a worldwide basis. However, most of their sales occurred in countries other than the United States. The United States, with its substantial installed base of telecommunications infrastructure (i.e., telephone and cable systems), did not prove to be very receptive to this relatively expensive type of broadband access technology except in a few isolated, scattered areas. Many other countries of the world with inferior or almost nonexistent telecommunications infrastructure proved to be much more receptive to a wireless technology that could be installed in a relatively rapid fashion and expanded on an "as needed" basis. During the worldwide telecommunications economic downturn of late 1999 and the early 2000s, many of the manufacturers in this telecommunications space changed product lines or ceased operations entirely as the market for these systems quickly contracted and basically disappeared. In retrospect, one might conclude that without a standard in place that could provide for multivendor product interoperability and with the relatively high cost of an, as yet, immature technology that these early products were ahead of their time and the economics were not quite correct for their widespread acceptance and adoption.

Since those early days, a great deal of technologic change has occurred and the regulatory environment has dramatically changed to embrace license-free operation in newly created unlicensed frequency bands (e.g., 2.4-GHz and U-NII bands) that now exist on an almost universal basis. The effect of Moore's law has been mentioned elsewhere in this text in the context of its effect on other types of wireless technologies. For wireless MANs, the ability of the semiconductor industry to mass-produce specialized RF and millimeter wave ICs has only recently become a reality but it has had a dramatic effect on the cost point of this type of equipment. What was once relatively cost prohibitive now becomes affordable or at least competitive with other more established broadband high-speed data transfer technologies (i.e., high-speed cable modems and xDSL service over existing telephone lines).

Technologic advances aside, the transformation of the regulatory environment in the United States might have even more of an impact on the eventual adoption of this form of wireless technology. The original IEEE 802.16 standard called for operation in licensed bands in the 10- to 66-GHz frequency range where line-of-sight (LOS) is required for satisfactory operation. Previously, the installation of new systems had to be coordinated with other preexisting systems on a case-by-case basis to prevent interference and to allow for system coexistence. The original standard has been amended to include operation in the 2- to 11-GHz frequency range in both licensed and unlicensed bands. In the United States this includes operation in the 2.4-GHz and recently expanded (see Chapter 9) 5-GHz U-NII bands where non-line-of-sight (NLOS) operation is possible. In many other parts of the world unlicensed operation is also allowed in the 3.5-GHz band. Furthermore, the recent (late 2003) decision by the FCC to provide additional spectrum in the 71–76 GHz, 81–86 GHz, and 92–95 GHz bands (except 94.0–94.1 GHz) for broadband millimeter wave local area networks and broadband Internet service is certainly going to impact this new technology in the U.S. marketplace. Other FCC rulings meant to allow some forms of unlicensed operation in the 50–70 GHz bands and to align the United States' frequency allocations with those of the international community are also sure to have their impact.

Since these regulatory changes (i.e., unlicensed operations are permitted) have occurred, many new low-cost products have been introduced into this marketspace and other major technology players including semiconductor manufacturer Intel have started to take an interest in this area. Intel is reportedly planning to design and produce an IEEE 802.16-compatible chip set. Furthermore, the term **Wi-Max** (similar to Wi-Fi) has been adopted to describe this technology space and several Web sites exist that are devoted to news and events about this reinvigorated broadband technology (e.g., www.wimaxforum.org). Present

telecommunications industry predictions are for annual sales of Wi-Max products to increase from approximately $250 million today (2004) to over $2 billion by the year 2008.

IEEE 802.16 and 802.16a Standards

The IEEE 802.16-2001 standard was adopted by the IEEE Standards Board late in 2001 and as an ANSI standard during 2002. As already stated, this standard only covered physical layer implementations for the 10–66 GHz frequency range. The MAC layer only supports LOS operation over fairly large channels (i.e., 25 to 28 MHz wide) that can support raw data rates in excess of 120 mbps. At the time of the standard formulation, the perceived application area for this form of wireless technology was broadband Internet access for the small office/home office (SOHO) through medium-sized to large office complexes. Several other IEEE 802.16x projects were also initiated during the same period and in some cases were either superceded or rolled into existing projects. For instance, IEEE 802.16.1 was incorporated into IEEE 802.16 and IEEE 802.16.3 became 802.16a. IEEE 802.16.1b was an amendment project that sought to extend the physical layer implementations to license-exempt bands designated for public network access. It was to specifically focus on the 5–6 GHz range but was to also cover all frequencies between 2–11 GHz. Furthermore, it was expressly intended to address issues involving coexistence with other unlicensed applications. Specifically, it was to propose strategies for coexistence with IEEE 802.11 and 802.15 wireless technologies. IEEE 802.16.1b has since been changed to IEEE 802.16b and then withdrawn as a project since this area has been addressed and covered by IEEE 802.16a-2003, another amendment to the IEEE 802.16 standard. IEEE 802.16a-2003 adds support for operation in license-exempt bands and an optional mesh topology (for NLOS propagation) at these lower frequencies. A further revision to IEEE 802.16 is presently in the formulation stages and is meant to consolidate IEEE 802.16, 802.16a, and 802.16c (another amendment) into one unified and updated 802.16 wireless standard.

Another new amendment project, IEEE 802.16e, provides enhancements to the standard to support operation of subscriber stations moving at vehicular speeds. This would expand 802.16 operations to include both fixed and mobile broadband wireless access for both the Enterprise and consumer market. This project has some similarities to the IEEE 802.20 project but focuses at the higher frequencies already addressed by IEEE 802.16.

IEEE 802.16.2-2004

IEEE 802.16.2-2004 is a newly adopted standard that revises IEEE 802.16.2-2001 (limited to 10–66 GHz) and provides guidelines for the coexistence of fixed broadband wireless access systems. It addresses two issues: the coexistence between multipoint systems and point-to-point systems in the 10–66 GHz frequency range and between fixed licensed systems in the 2–11 GHz bands. The standard's stated purpose is to facilitate the deployment and operation of fixed broadband wireless access systems. A primary goal of this standard is to provide a means by which these systems may be deployed without having to go through a time-consuming case-by-case coordination process. The standard defines a set of consistent design and deployment recommendations that can be used to facilitate the coexistence of fixed BWA systems. The standard addresses equipment design parameters including such specifications as radiated power, modulation spectral masks, antenna radiation patterns, and limits on both in-band and out-of-band fixed BWA system emissions. The standard further provides a systematic approach for the deployment and coordination plans for fixed BWA systems. Much of the standard is based on newer propagation models and simulation techniques that are used to analyze various wireless system coexistence scenarios.

IEEE 802.16/Conformance Standards

There are also IEEE 802.16 standards for conformance that specify the tests that are to be used to check the conformance of both base and subscriber stations to the specifications of the physical layer for the

WirelessMAN-SC (single carrier) air interface. These standards deal with protocol implementation, interoperability, and radio conformance tests. The reader may question the need for this particular type of standard. It should be pointed out that microwave and millimeter wave equipment has different forms of measurement associated with it than equipment that works at lower RF frequencies. Components and systems at these frequencies are characterized through specialized test and measurement equipment (network and spectrum analyzers, etc.) that tends to be peculiar to the microwave/millimeter wave industry.

Since a great deal of regulatory and technologic change is expected to occur in the wireless broadband access field in the next few years, it is reasonable to expect that there will be more amendments added to the IEEE 802.16 standard to address these changes as they occur. This chapter will focus its discussions on the differences between IEEE 802.16 and the other IEEE 802 wireless technologies (i.e., wireless LANs and PANs).

11.2 IEEE 802.16 WIRELESS MANS

Wireless metropolitan area networks (MANs) provide network access to buildings (see Figure 11–1) through exterior antennas communicating with a central radio **base station** over a point-to-multipoint radio link. Therefore, the wireless MAN offers an alternative to "wireline" type access networks. Due to the relatively low cost of deployment of wireless MAN (WMAN) technology, it certainly should prove to be cost-effective compared to the installation of fiber-optic links for instance. Furthermore, for business applications most broadband cable networks do not provide cable drops for small-business enterprises nor do they provide the required bandwidth capacity since they are shared systems designed for high-speed Internet access for the home user. In many instances, wireless MAN technology could provide network access where DSL technology would fail due to distance limitations or severe copper pair signal impairments. In developing countries without an extensive installed infrastructure base, WMANs might be the first choice for network access.

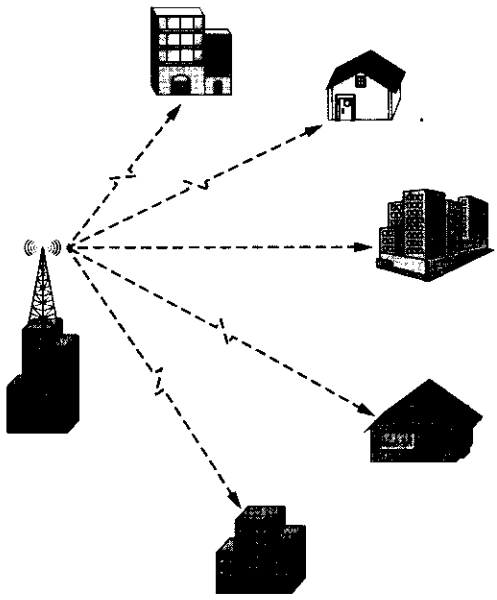

Figure 11–1 Typical wireless metropolitan area network.

At present, the use of wireless MANs that conform to the IEEE 802.16 standard brings the network to the building. Users inside the building connect to the network with conventional in-building network technologies like Ethernet or possibly wireless LANs (IEEE 802.11). Amendments to the standard (recall IEEE 802.16e) may eventually allow an extension of the standard to provide network connectivity to an individual's laptop/notebook/tablet computer or PDA while outside, in one's home or apartment, or while in a moving vehicle.

A wireless MAN effectively serves as a bridge to an existing network infrastructure. This bridge function may extend the network wirelessly to multiple new fixed locations where the network deployment might use standard network infrastructure (wired and wireless LAN technologies) to provide network connectivity to the end users. This functionality is similar to some degree with the primary purpose of IEEE 802.11, to extend the reach of the network. However, it is different in scale since IEEE 802.11 has a maximum reach in the order of 100 meters for indoor access points whereas wireless MANs may span several thousands of meters in an outdoor environment. The wireless LAN also may provide an outdoor bridge function over many kilometers of distance but it is usually only for specific point-to-point applications. The wireless PAN does not appear to have any commonality with the wireless MAN except for the functionality exhibited by the wireless PAN piconet and the wireless MAN mesh network operation. In both cases, the network structures extend the reach of a single node beyond what it is physically capable of. Again, the scale of this effect is what is different here. for WPANs, tens of meters (piconets) versus kilometers for the wireless high-speed MAN.

Typical Deployment

The typical deployment of an IEEE 802.16 system is depicted by Figure 11–1. A wireless MAN **base station** is typically located on a tall building to provide an unobstructed or line-of-sight path between the **subscriber stations** and the base station antennas. Although the new IEEE 802.16a physical layer standard provides for NLOS operation at frequencies between 2–11 GHz, the type of installation depicted in the figure (i.e., a substantial base station antenna height) is still desired because the best system operation, with the highest possible data transfer rates, is still dependent upon base station to subscriber station radio channel characteristics. In all cases, a direct LOS path will provide the best channel transmission characteristics. Even with mesh network operation (to be discussed in a later section), the greater the number of mesh stations with LOS views of the mesh base station, the better the system operation. For LOS operation, a typical cell radius for a wireless MAN system with the base station antenna at a height of 30 meters and the SS antenna at 6.5 meters is approximately 3.5 km. For an 80-meter base station antenna height the cell radius increases to about 7 km. System bit rates are dependent upon the system bandwidth and the coding/modulation formats used. Typical operational values range from 5 to 10s of mbps in the 2–11 GHz range and higher values for systems deployed in the 10–66 GHz range. The subscriber station antenna is typically mounted on an outside building wall, base station facing window, or on a pole aimed at the base station antenna.

To increase system capacity, a wireless MAN base station usually supports numerous antenna sectors. Figure 11–2 illustrates the use of a rather complex, high-capacity, four by four-sector system that provides four-frequency, four-sector frequency reuse. As shown in the figure, four different frequency channels are used within every sector. There are four, 90-degree sectors. Therefore a total of sixteen separate sectors (of 22.5 degrees each) can be supported, with numerous subscriber stations per sector. For a configuration like this, the base station would consist of sixteen radio transceivers and sixteen individual sector antennas that would have narrow fan-beam/pencil-beam type radiation patterns. For this example, each one of the sixteen sectors could support the same total data rate that a single omnidirectional base station could. It would be likely that the service provider would need to employ some form of fiber- optic transport/connection to the network to support the total aggregated system bandwidth to and from the base station. Other types of (N by M) frequency reuse schemes can be implemented for a single base station and the use of different

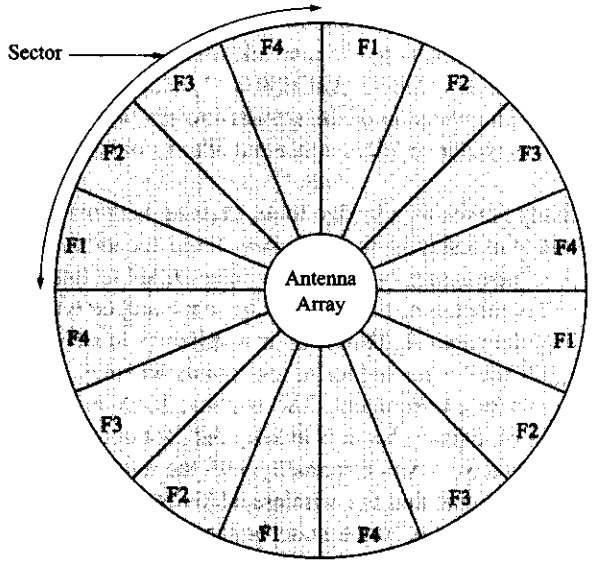

Figure 11–2 Wireless MAN 4 × 4 antenna sectoring scheme.

antenna polarization orientations (i.e., vertical and horizontal) can be used to increase cell capacity or to implement other types of reuse plans. Frequency reuse schemes similar to those used by the cellular industry (refer back to Chapter 4) can be used when it is desired to provide blanket coverage over a given area with wireless MAN cells.

11.3 IEEE 802.16 MAC LAYER DETAILS

Figure 11–3 details the OSI-based reference model for the IEEE 802.16 standard. As indicated in the figure, the MAC layer consists of three sublayers. External network data is received (or transmitted) through the convergence sublayer (CS) service access point (SAP). This external data is transformed or mapped

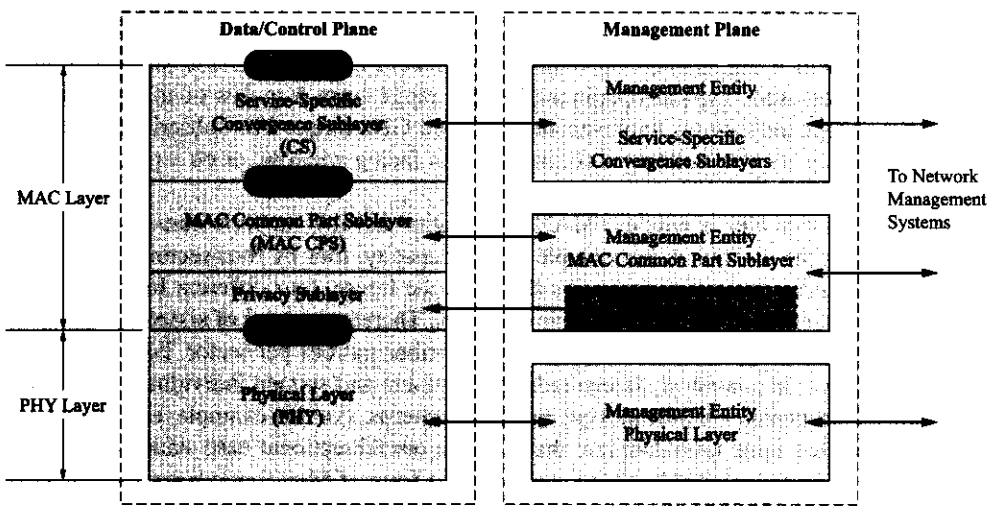

Figure 11–3 OSI reference model for the IEEE 802.16 standard (Courtesy of IEEE).

into MAC service data units (SDUs) by the service-specific convergence sublayer and delivered to the MAC common part sublayer (CPS) through the MAC SAP. The functions provided by the CS include the classification of external network SDUs and their association with the correct MAC service flow and connection identifier (CID). The CS also provides payload header suppression if necessary. There are several CS specifications within this sublayer that provide support for higher-layer protocols like asynchronous transfer mode (ATM), IEEE 802.3 (Ethernet), point-to-point protocol (PPP), and IPv4 and IPv6.

The MAC CPS provides the primary MAC functionality of wireless MAN system access, connection establishment, connection maintenance, and bandwidth allocation between the subscriber units and the base station. This sublayer receives data that has been classified and assigned to particular MAC connections from the CSs. A MAC privacy sublayer exists between the MAC CPS and the physical layer (PHY) that provides the system functions of authentication, secure key exchange, and the subsequent encryption of transferred data. The MAC layer also provides for the assurance of a specific QoS for the connection through proper scheduling and transmission of data over the PHY. The transfer of data, physical layer control, and system statistics (RSSI measurements, etc.) occurs between the MAC CPS and the PHY through the PHY SAP. This wireless MAN standard supports several implementations of the PHY, each devoted to a particular frequency range and air interface environment. The next few sections will provide some additional detail about these operations but what follows is not meant to be a complete coverage by any means. It is hoped that this brief overview will provide the reader with a general sense of how the system operates.

MAC Service-Specific Convergence Sublayer

The MAC service-specific convergence sublayer performs the following tasks on external network data that is to be delivered across the wireless MAN air interface. It receives protocol data units (PDUs) from these higher layers and performs a classification function and any necessary processing before delivering the CS PDUs to the correct MAC SAP. It also receives CS PDUs from its peer entity via the physical layer. At present, the standard provides for the support of the transport of ATM cells and packet data. Therefore, an ATM CS and a packet CS are specified for the system. Other CSs can be added to the standard in the future.

Without going into great detail about ATM operation, this brief discussion will attempt to provide an overview of how the CS works. The ATM CS accepts ATM cells from an ATM network, performs classification, and if so provisioned, payload header suppression, and then delivers the CS PDUs to the appropriate MAC SAP. An ATM cell consists of a 5-byte header and a 48-byte payload. The ATM CS PDU will consist of an ATM CS PDU header and the original ATM cell payload. The original ATM header consists of information about the ATM connection. These connections may be uniquely identified using virtual path and channel identifiers (VPI and VCI) and bits that indicate whether the connections are switched or permanent. The ATM CS PDU header can be either the original ATM header or a modified shortened header that retains sufficient information about the switched ATM connection to allow for correct system operation. In the payload header suppression mode, for virtual path-switched mode, the VPI value is mapped to a 2-byte CID value representing the MAC connection over which it will be transported and the ATM CS PDU header is shortened to 3 bytes of which 2 bytes consist of the VCI value. For virtual circuit-switched mode, the VPI/VCI combination is mapped to a single 16-bit CID value and the header is reduced to a single byte. In both cases, the receiving entity restores the ATM header and the format modifications performed by the MAC layer protocol become transparent to the final higher-level network destination.

The packet convergence sublayer performs the same basic operations as just discussed but for packet-based protocols. The sending CS is tasked with providing a MAC SDU to the MAC SAP. The MAC layer is tasked with the delivery of the MAC SDU to the peer MAC SAP in accordance with the particular connection's service flow characteristics (QoS, fragmentation procedures, etc.). Finally, the receiving CS is tasked with accepting the MAC SDU from the peer MAC SAP and delivering it to a higher-layer entity. For packet data the classification function is performed by applying some form of matching criteria to each

packet that enters the packet CS. The matching criteria is protocol specific and could be something simple, like the destination IP address for instance. If a packet matches the criteria, it is delivered to the MAC SAP for delivery over the MAC connection defined by a particular CID. The service flow characteristics of the particular connection provide the QoS for the transport of the packet. Payload header suppression techniques are also used by the packet CS, but this fairly complex process will not be discussed here.

MAC Common Part Sublayer

The use of a shared wireless medium (the air interface) requires the use of procedures that allow for the efficient allotment and use of this limited resource. Wireless MANs with their two-way point-to-multipoint and mesh topology networks are examples of systems that require many complex media access control (MAC) procedures to achieve efficient operational status. For wireless MANs this system functionality resides in the MAC common part sublayer. A quick overview of IEEE 802.16 system operation will make this more apparent.

The downlink from the base station to the subscriber operates on a point-to-multipoint basis and therefore its transmissions do not have to be coordinated since all users receive the same transmissions. The subscriber station checks the address in the received message and only retains those messages that have been specifically addressed to it. On the uplink, users share the radio link on a demand basis. Typically, the SS must initially request a certain QoS that provides it with scheduled periodic transmission rights or it might have to request transmission time from the base station on a needs basis. The users of the system adhere to transmission protocols that control contention issues and allow for uplink scheduling mechanisms that optimize system performance.

To provide this functionality, the IEEE 802.16 MAC is connection oriented. When an SS is first attached to the system, the initialization process includes the provisioning of service flows. Shortly, after SS registration, MAC connections are associated with the previously provisioned service flows. Let us examine this process more closely. Each SS has a 48-bit universal MAC address. This address uniquely identifies the SS and is used during the registration process (and also during the authentication procedure) to establish the correct connections for an SS. A closer look at the SS initialization process reveals that three different duplex connections are established between the SS and the BS. These MAC connections are identified by 16-bit CIDs that are assigned during the exchange of the ranging request and registration request MAC management messages. These three connections provide different levels of QoS for the MAC management traffic between the BS and the SS. Basic, primary, and secondary connections are used for the exchange of MAC messages that are short and time sensitive, longer and less time sensitive, and delay tolerant, respectively. To support system operation, additional transport connections are allocated to SSs for the users' contracted services. These transport connections are typically assigned in pairs; however, they are usually unidirectional in nature to facilitate different data transfer rates and QoS requirements on the downlink and uplink paths.

The concept of a service flow on a particular SS-to-BS connection is fundamental to the operation of the MAC protocol. The use of service flows provides a built-in method for downlink and uplink QoS management. When an SS requests uplink bandwidth on a per connection basis, the SS is implicitly identifying the QoS needed for the connection. The BS will grant the requested bandwidth in response to the type of SS request. On the downlink side of the system, the BS sets up downlink connections based on the provisioning information provided to it about the existing SS-to-BS connections. Once MAC connections have been established, there are connection maintenance issues that may arise. Dynamic modifications of a connection may be necessary due to the type of traffic that is being handled (i.e., IP traffic is typically very bursty in nature compared to other forms of traffic) or either the SS or the BS may initiate a connection modification due to a change in traffic bandwidth requirements. Finally, connection termination is possible. This usually only occurs if the user cancels his or her subscription to the service (or does not pay the bill!). All of these previously mentioned management functions are handled through the use of MAC management messages

designed to implement static configurations between the SS and the BS that, once in existence, may undergo dynamic addition, modification, and termination operations.

To help accomplish all of the required operations that provide the IEEE 802.16 system functionality, the MAC common part sublayer provides service definitions for the interface (MAC SAP) between the MAC CPS and the CS. These service primitives belong to one of four groups: MAC_CREATE_CONNECTION, MAC_CHANGE_CONNECTION, MAC_TERMINATE_CONNECTION, and MAC_DATA. The standard completely describes the function, semantics, conditions under which they are generated, and the effect of the receipt of each of the primitives. As a typical example, Figure 11–4 illustrates the sequence of logical events that occurs during the creation of a connection that is requested by the convergence sublayer.

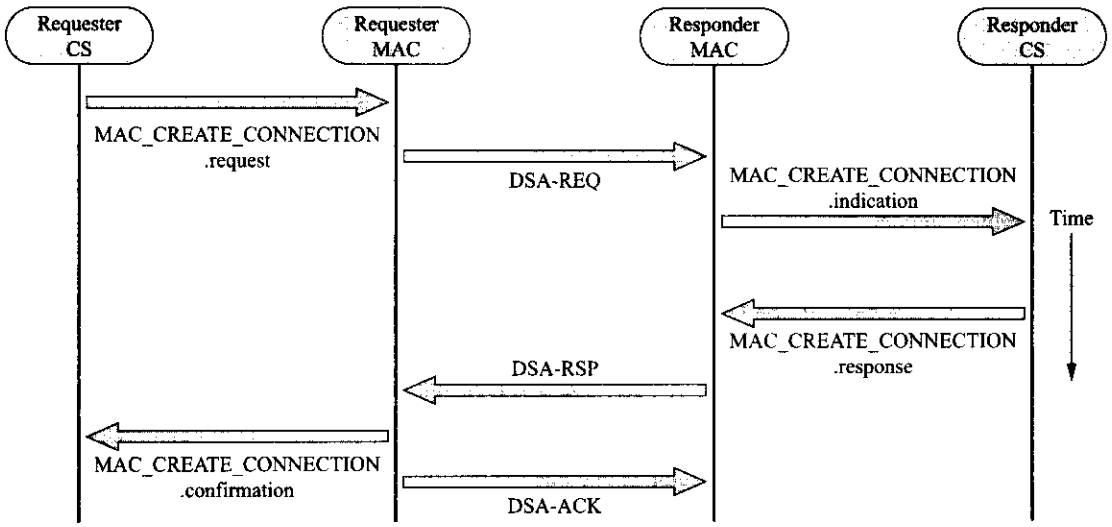

Figure 11–4 Operations occurring during the creation of a IEEE 802.16 connection (Courtesy of IEEE).

The standard provides detailed rules about the construction of a MAC PDU that consists of a MAC header and the MAC payload (a CRC field is optional). The MAC header can take on two distinctly different formats: a generic header that contains MAC management messages or a CS data and a bandwidth request header that is used to request additional system bandwidth. Three types of MAC subheaders are used by the system: fragmentation, grant management, and packing. These subheaders facilitate fragmentation control operations, allow the SS to convey bandwidth needs to the BS, and facilitate the packing of multiple SDUs into a single MAC PDU. Data encryption operations on the MAC PDU are only performed on the payload portion of the MAC PDU.

A set of MAC management messages is defined and they are carried in the payload of a MAC PDU. The message format consists of a management message type field and additional fields used to provide additional message descriptors. It will be helpful to introduce some of these MAC management messages here since they will be referenced during our coverage of the various physical layers supported by IEEE 802.16. **Downlink channel descriptor** (DCD) and **uplink channel descriptor** (UCD) messages are transmitted periodically by the BS to provide details of the characteristics of the downlink and uplink channels (i.e., BS transmit power, physical layer type, FDD/TDD frame durations in the DCD message and uplink preamble length, minislot size, and contention parameters for the UCD message). Note that a separate UCD message must be transmitted for each active uplink channel associated with the downlink channel. DL-MAP and UL-MAP messages are generated by the BS. The **DL-MAP** message defines the access to the downlink information. The DL-MAP message contains information about the following parameters: physical layer synchronization (physical layer dependent), DCD count, base station ID (a 48-bit-long field), and the num-

ber of information elements (IEs) that follow. The **UL-MAP** message allocates access to the uplink channel. The UL-MAP message contains information about the following parameters: uplink channel ID, UCD count, number of information elements, allocation start time (the effective start time of the uplink allocation in units of minislots), and map information elements. The last parameter, MAP IEs, consists of information about the CID, the uplink interval usage code (UIUC), and offset. The IEs define the uplink bandwidth allocations. The CID represents an assignment of an IE to a particular type of address (i.e., unicast, multicast, or broadcast). When used to indicate a particular bandwidth grant, the CID will represent either the basic CID or one of the transport connection CIDs of the SS. An UIUC is used to define the type of uplink access and the uplink burst profile needed for that access. Note that an Uplink_Burst_Profile is included in the UCD for each UIUC to be used in the UL-MAP.

The ranging request (RNG-REQ) message and the ranging response (RNG-RSP) messages are used during the initial attachment of the SS to the system. Other MAC management messages are listed here without explanation. They are privacy key management messages; security association add; authorization request, reply, invalid, and reject; key request, reply, and reject; authentication information messages; a number of dynamic service addition, change, and deletion messages; polling assignment; downlink burst profile change messages; and so forth. For further information and details about these MAC messages, the reader should consult the most recent version of the IEEE 802.16 standard.

This concludes our discussion of the MAC CPS for the moment. After a discussion of the various physical layers supported by the IEEE 802.16 standard, some of the common system operations supported by the MAC layer will be addressed in more detail in a later section of this chapter.

11.4 IEEE 802.16 PHYSICAL LAYER DETAILS

Presently, the IEEE 802.16 standards call for several different physical layer implementations. These different implementations use different randomization, data encoding, symbol mapping, modulation techniques, and channel access methods. In an effort to present this material in a comprehensive yet understandable fashion, the various IEEE 802.16 physical layer specifications will be offered for the 10–66 GHz frequencies first and then for the 2–11 GHz frequencies. According to the standard, the 10–66 GHz physical layer specification (now known as WirelessMAN-SC) is designed with flexibility in mind. This flexibility is meant to allow service providers the ability to optimize system design and deployment. In theory, this flexibility should allow for the most efficient use of the available frequency spectrum in the 10–66 GHz range.

The physical layer specifications for the 2–11 GHz frequency spectrum consist of three different schemes that are all optimized for NLOS operation. These options are designated for use in the licensed bands and with an additional modification they can all be used in the 2–11 GHz license-exempt bands. For use in the licensed bands, the WirelessMAN-SCa physical layer is based on the use of a single carrier, the WirelessMAN-OFDM physical layer is based on the use of OFDM modulation (i.e., multiple carriers), and the WirelessMAN-OFDMA physical layer makes use of a form of wireless channel access using OFDM modulation (known as OFDM access or OFDMA). The use of this last technique is at this time unique to this wireless standard. For use in the license-exempt bands (primarily in the 5–6 GHz range), the WirelessHUMAN physical layer is implemented. This specification adds the use of dynamic frequency selection (DFS) to the three physical layers already listed. As stated earlier, in the interest of providing comprehensive coverage of this material, most emphasis will be placed on technology implementations that have not been discussed elsewhere within this text. Therefore, at times during this section, the reader will be referred to previously covered topics in this or other chapters instead of a rehash of the same explanations.

The reader may also note that details of the physical layer service specifications are not covered here. The physical layer service is provided to the MAC entity at both the BS and SS through the physical layer SAP as shown in Figure 11–3. The physical layer SAP service primitives are related to physical layer management activities, data transfers, and sublayer-to-sublayer interactions. These operations are similar to

those detailed in other chapters of this text for other wireless technologies and therefore will not be chronicled here. The reader is referred to Section 8.1 of the latest edition of the IEEE 802.16 standard for the details of these operations.

IEEE 802.16 Physical Layer for 10–66 GHz (WirelessMAN-SC)

In the IEEE 802.16 standard, the physical layer supports the usage of both time division duplex (TDD) and frequency division duplex (FDD) operation. For both cases, a burst transmission format is used in conjunction with a framing structure that supports the use of adaptive burst profiling. This means that the transmission parameters (i.e., coding and modulation schemes) are able to be adjusted individually for each subscriber station on a frame-by-frame basis. The use of FDD allows subscriber stations (SSs) the ability to operate in both full-duplex and half-duplex modes.

The uplink transmission is based on both time division multiple access (TDMA) and demand assigned multiple access. This is achieved through the use of multiple timeslots on the uplink channel. The MAC layer in the base station dynamically controls the number of timeslots assigned for various system functions (i.e., registration, contention, guard time, or user traffic) for optimal system performance. The downlink channel uses time division multiplexing (TDM). The data for each subscriber station is multiplexed on to a single downlink stream of data that is received by all the subscriber stations serviced by this radio link. Half-duplex subscriber operation is also supported by the use of a TDMA portion of the downlink frame time. The downlink physical layer includes a transmission convergence sublayer function that will insert a pointer byte at the beginning of a MAC protocol data unit (PDU). This pointer is used by the receiver to identify the MAC PDU.

Both the downlink and uplink transmission schemes call for the randomization of the data to be transmitted (thus assuring sufficient bit transitions), FEC encoding, and a mapping of the coded bits to either QPSK, 16-QAM, or optional 64-QAM constellations depending upon various system parameters such as baud rate and channel characteristics.

Duplexing Techniques

The IEEE 802.16 physical layer supports both FDD and TDD operation. For FDD the downlink and uplink channels operate at different frequencies. The fact that the downlink allows for burst transmission facilitates the use of different modulation schemes as well as the simultaneous support for both full-duplex and half-duplex subscriber stations. Full-duplex operation is made possible through the use of two separate transmit and receive channels. Half-duplex operation is supported through the use of different time periods (a form of bandwidth allocation) during the downlink and uplink frames for the transmission and reception of data by the half-duplex station. Figure 11–5 depicts this type of operation in more detail.

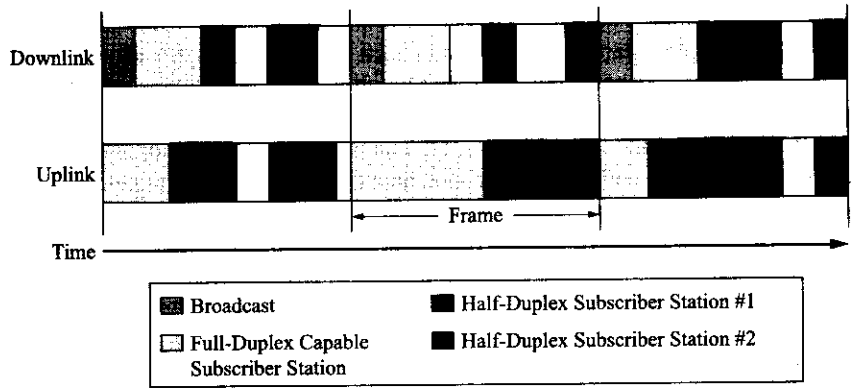

Figure 11–5 IEEE 802.16 uplink and downlink frame structure (Courtesy of IEEE).

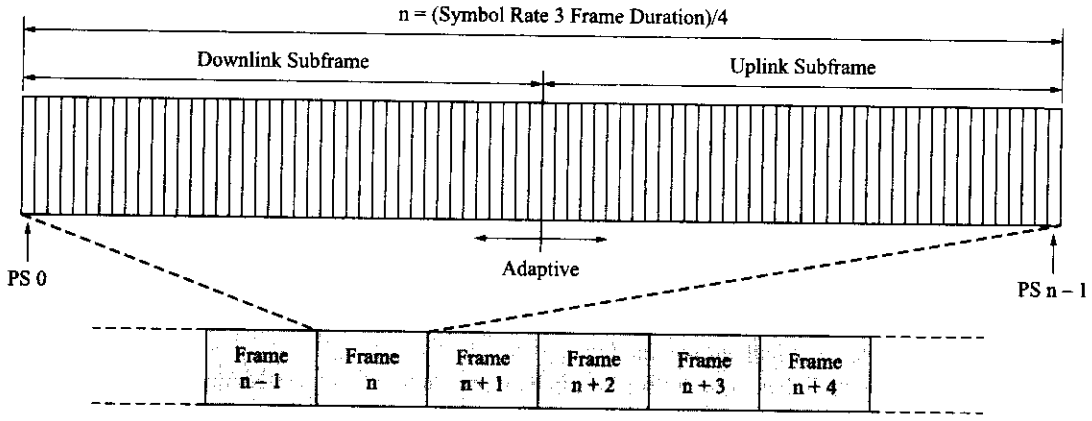

Figure 11–6 IEEE 802.16 time division duplex operation (Courtesy of IEEE).

For time division duplex operation, both the downlink and uplink transmissions share the same single-carrier frequency channel; however, they do so at different times as shown in Figure 11–6. In the TDD mode of operation, each frame is of a fixed duration and contains a downlink and uplink subframe. The entire frame consists of a certain number of equal-size physical slots (PSs). Within the TDD frame structure, the duration and hence the capacity of the downlink and uplink subframes may be varied (i.e., the number of PSs assigned to each operation). Again, this system functionality provides the ability to dynamically allocate bandwidth as needed.

For both the FDD (half-duplex operation) and TDD modes of operation, there are transmit-to-receive and receive-to-transmit time gaps that are used to allow both the subscriber and base station receiver and transmitter time to ready themselves for operation and the antenna transmit/receive (Tx/Rx) switch to change state. The length of these timing gaps is always set at an integer multiple of the PS duration.

Downlink Operation

The IEEE 802.16 downlink specifications for 10–66 GHz provide for different amounts of bandwidth (i.e., the data transfer rate in bps) that is dependent upon and defined by the available bandwidth for one physical slot. The number of PSs within each frame is a function of the symbol rate. The symbol rate is selected in such a fashion as to yield an integral number of PSs within each frame. Frames may be either 0.5, 1.0, or 2.0 ms in length. For example, for a 40-Mbaud symbol rate, the number of PSs for a 2-ms duration frame is 20,000 (see the equation in Figure 11–6). The available uplink bandwidth is defined in terms of uplink minislots where a minislot length may range from 1 to 128 PSs in powers of 2.

The downlink subframe structure for TDD operation is shown in Figure 11–7. As shown by the figure, the subframe consists of a frame start preamble that is used by the receiver for synchronization and equalization. Next is a frame control section that contains downlink and uplink maps (DL-MAP and UL-MAP) that state the physical slots where bursts begin. The next section carries the actual data that is organized into bursts that employ different modulation schemes. The bursts are transmitted according to their robustness (i.e., low-BER formats first) with QPSK-encoded data first, 16-QAM next, and 64-QAM encoded data transmitted last. As shown in the figure, a timing gap separates the downlink subframe from the subsequent uplink subframe. On the receiving end of this downlink subframe, each subscriber station receives and decodes the control portion of the subframe and then looks for MAC headers that indicate data for the individual SSs in the rest of the downlink subframe.

In the case of FDD operation, the downlink subframe is shown by Figure 11–8. The FDD subframe is very similar to the TDD subframe except that it contains both a TDM section and a TDMA section. The TDM section is used to transmit data to any full-duplex SSs, half-duplex SSs that are scheduled to transmit

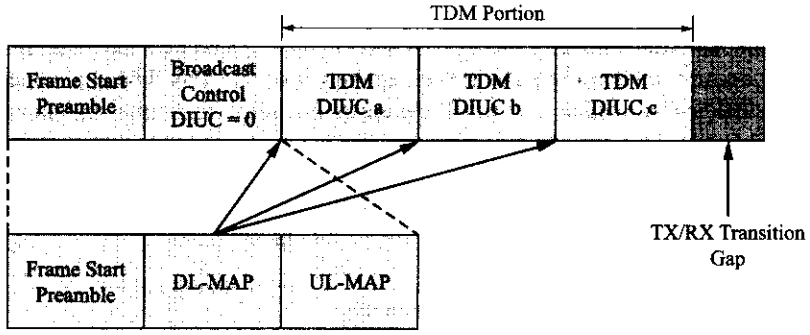

Figure 11–7 TDD downlink subframe structure (Courtesy of IEEE).

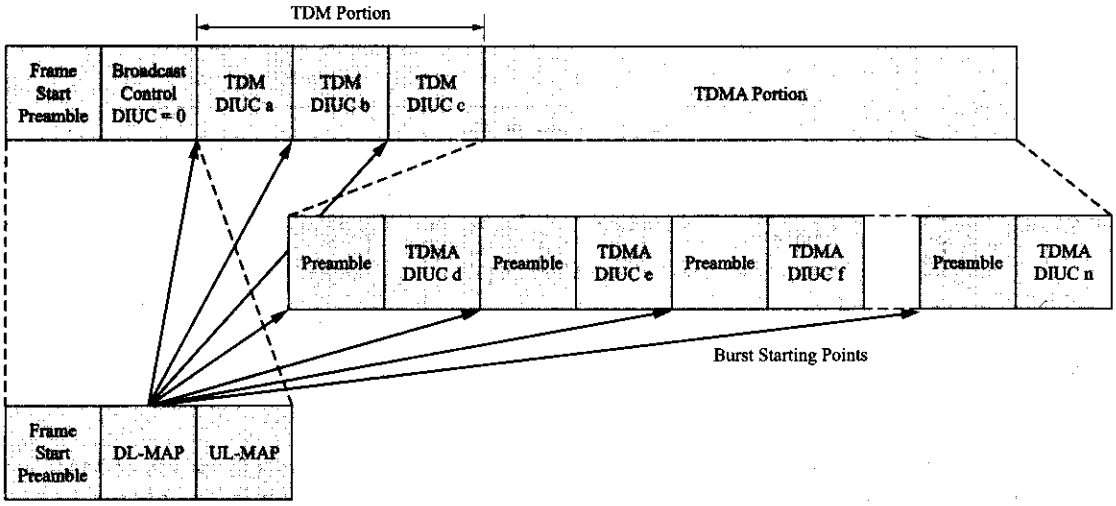

Figure 11–8 FDD downlink subframe structure (Courtesy of IEEE).

later in the frame than they receive, and half-duplex SSs that are not scheduled to transmit during this frame. The TDMA section of the subframe is used to transmit data to any half-duplex SSs scheduled to transmit earlier in the frame than they receive. Using this structure, an individual SS does not need to decode the entire downlink subframe, only the portion that is directed to it. Within the TDMA section, each burst begins with a downlink TDMA burst preamble (using a different format than the frame start preamble) that is used for phase resynchronization for half-duplex SSs. As shown in the figure, the FDD frame control section includes a map of both the TDM and TDMA section bursts. The data bursts within the TDMA section are not required to be transmitted in any particular order, and for the situation where no half-duplex SSs are scheduled to transmit before they receive, the TDD and FDD subframes are identical in structure.

The frame control section is used to transmit broadcast information to all the SSs within a service area. The information transmitted in this section uses the downlink interval usage code (DIUC) of 0, which indicates a specific downlink burst profile. The frame control section also contains a DL-MAP message followed by one UL-MAP message per each associated uplink channel. Furthermore, after the last UL-MAP message, the frame control section may also contain downlink and uplink channel descriptor (DCD and UCD) messages.

The downlink data burst sections are used to transmit both data and control messages to the SSs. The data to be transmitted is always FEC encoded and transmitted with the correct modulation format for the

receiving SS. Again, the sequence of data transmitted in the TDM section is ordered by its modulation robustness whereas the data transmitted in the TDMA section does not require any special sequencing. The DL-MAP message contains information about the number(s) of the PSs where a burst profile change occurs. If there is insufficient downlink data to fill the entire downlink subframe, the transmitter will automatically shut down when it is finished transmitting the data bursts that it has.

Downlink Physical Medium Dependent Sublayer The downlink physical layer coding and modulation scheme for this particular part of the IEEE 802.16 standard is shown in Figure 11–9.

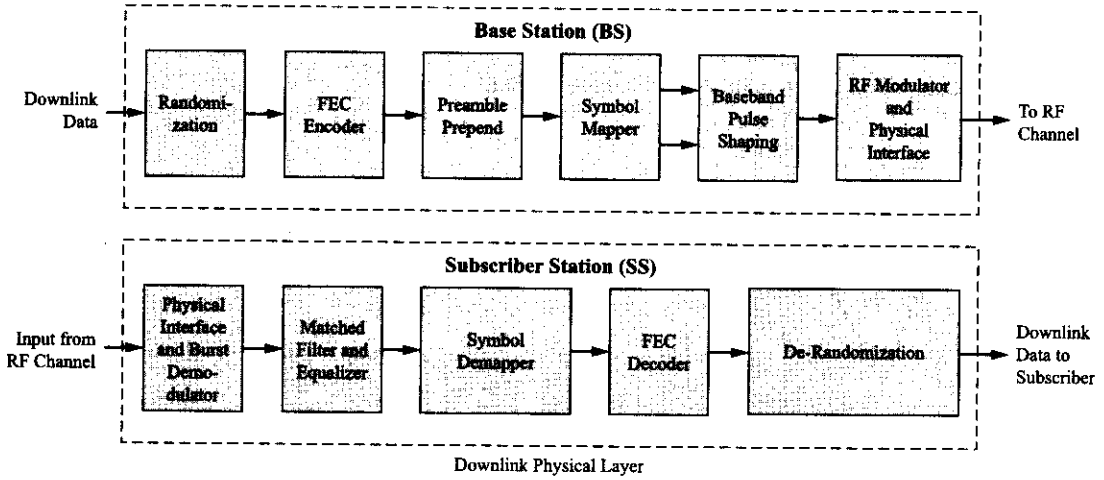

Figure 11–9 Block diagram of the IEEE 802.16 downlink physical layer coding and modulation scheme (Courtesy of IEEE).

The downlink burst profile of the user data is communicated to the SSs through the use of MAC messages during the frame control portion of the downlink subframe. Also, since it is possible for the SS to employ optional modulation and FEC encoding schemes, this information is communicated to the base station during the subscriber registration operation. Randomization of the data to be transmitted is performed to ensure the ability of the system to recover the clock signal from the incoming data stream. Several different complex forward error correction schemes are employed by the system depending upon the level of protection desired, the type of modulation to be used, data block size, and so on. Reed-Solomon, convolutional codes, and block turbo codes are all part of the FEC encoding suite that may be used. Various portions of the downlink subframe typically use different types of FEC encoding.

IEEE 802.16 Modulation Techniques To optimize the system for maximum reliable data transmission rates, the IEEE 802.16 physical layer allows for the use of several multilevel single-carrier modulation schemes. For use in the 10–66 GHz frequency range, the standard Wi-Max base station must be able to support QPSK, 16-QAM, and optionally 64-QAM modulation. In each case, the baseband I and Q signals must be prefiltered (shaped) by a square-root raised cosine filter before undergoing RF modulation and being upconverted into a microwave/millimeter wave passband signal.

Uplink Operation

The uplink subframe structure is shown in Figure 11–10. During the uplink subframe period, three classes of data bursts may be transmitted by the subscriber station. They are bursts that are transmitted during contention periods reserved for initial maintenance, bursts that are transmitted during contention periods defined as request intervals and reserved for responses to multicast and broadcast polls, and bursts that are

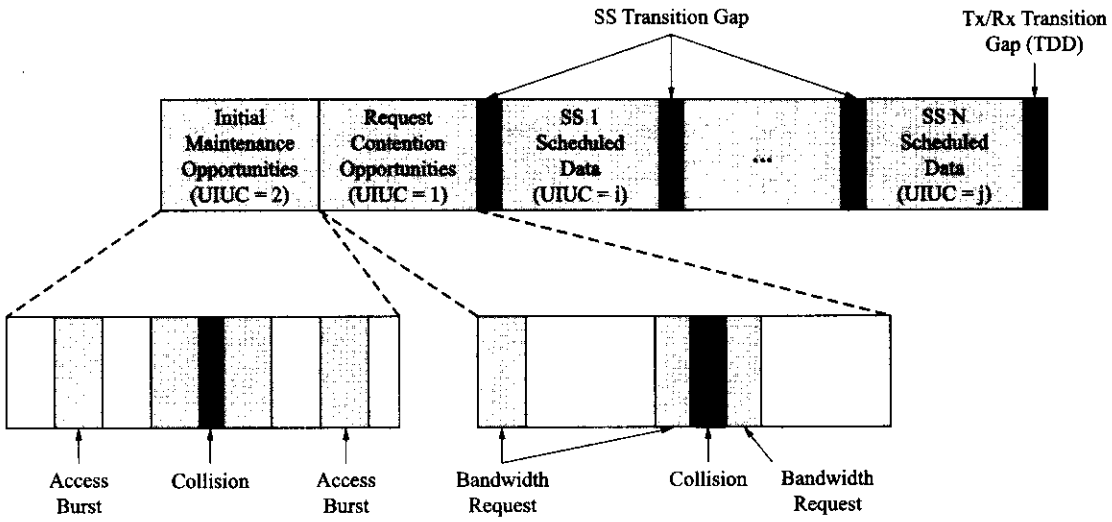

Figure 11–10 IEEE 802.16 uplink subframe structure indicating SS access burst activity (Courtesy of IEEE).

transmitted during intervals defined by data grant information elements (IEs) that have been specifically allocated to individual SSs. In any given uplink subframe, these burst classes may be present. The quantity and order of the bursts within the frame are set by the base station uplink scheduler and indicated by the UL-MAP as contained in the frame control portion of the downlink subframe transmitted by the base station.

The contiguous physical slots allocated for individual SS scheduled data transmissions are grouped together by SS. Within the scheduled SS timeslots, the SS transmits the uplink data using the burst profile that has been previously specified or assigned by the base station to the individual subscriber station. Furthermore, uplink SS transition gaps (similar to the downlink timing gaps) are used to separate the transmissions of the various SSs that occur during the uplink subframe. The transition time gap allows for the base station to ready itself for the reception of a preamble from the next transmitting SS. This uplink burst preamble, similar to the downlink burst preamble, allows the base station to synchronize itself with the new SS.

The uplink physical layer medium dependent, sublayer coding and modulation scheme are depicted in Figure 11–11. As the reader can see, it is very similar to that shown for the downlink in Figure 11–9. The randomization process and FEC schemes are basically identical, with only some slight FEC modifications implemented due to the smaller expected size of the uplink traffic payload. The baseband prefiltering and modulation schemes are identical to what has already been discussed for the downlink.

Miscellaneous Radio Subsystem Control Techniques

The subscriber station downlink demodulator is typically used to provide a reference clock that will be in synchronization with the base station downlink clock (typically locked to GPS time). This reference clock will then be able to be used by the subscriber station for accurate timing of system functions.

At the high frequencies used by the air interface for IEEE 802.16, frequency control becomes very important. Temperature and aging can cause frequency errors that are multiplied at microwave and millimeter wave frequencies. In an effort to reduce the complexity of the subscriber stations (i.e., keep the cost low) the uplink and downlink carrier frequencies will be used to reference one another. When a subscriber station initially attaches itself to a servicing base station, it undertakes an initial ranging process for timing, frequency, and power calibration. Once the initial frequency calibration has been performed, periodic measurements of the base station frequency offset value will be made and sent to the SS through the use of a MAC message.

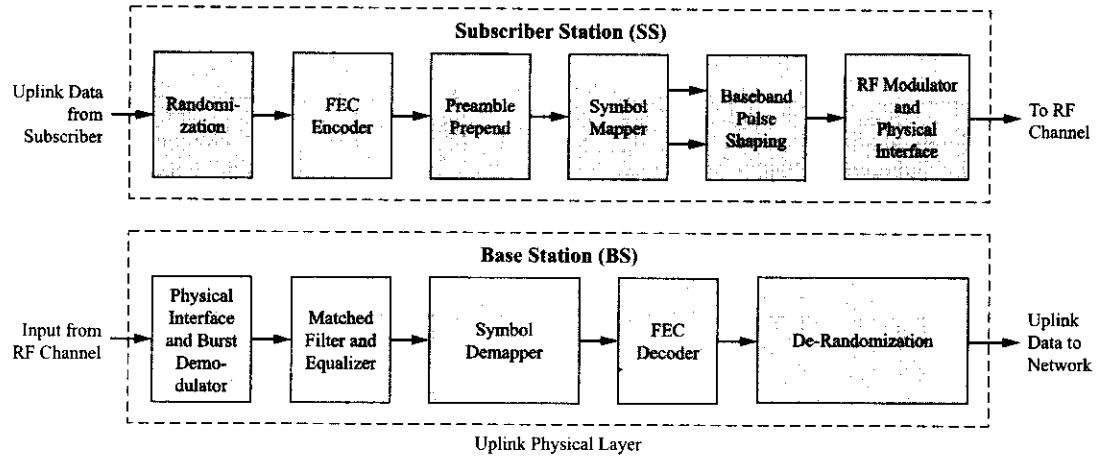

Figure 11–11 Block diagram of the IEEE 802.16 uplink physical layer coding and modulation scheme (Courtesy of IEEE).

A power control algorithm will be supported by the standard for the uplink channel. This algorithm will require an initial calibration procedure and then the ability to perform a periodic adjustment procedure without data loss. The base station must have the ability to make accurate measurements of the power level of the SS's received burst signal (RSSI). Using this information, the received signal level will be compared to a reference level and a calibration message can be sent to the SS via a MAC sublayer message. The power control algorithm should be able to compensate for power fluctuations or fades that occur at a maximum rate of 10 dB/second and have depths of up to 40 dB. The SS should be able to control its output power in fixed steps and through the feedback process (MAC messages) using the base station's RSSI measurements.

The standard calls for some minimum hardware performance requirements for operation in the 24–32 GHz range. The SS must be able to output +15 dBm and have certain maximum BER specifications for the different modulation types allowed by the standard. The BS output power should not exceed either +14 dBW/MHz or local regulatory requirements and should conform to ETSI modulation spectrum mask standards when used in ETSI territories.

Baud Rates and Channel Bandwidths for 10–66 GHz Systems

Within the 10–66 GHz range there is a large amount of spectrum that can be used for point-to-multipoint systems. Unfortunately, there are many different regulatory requirements worldwide. Therefore, the standard does not specify certain frequency bands within the 10–66 GHz range. What the standard does specify is the channel size and the corresponding symbol rate and bit rates for the different modulation schemes employed in this frequency range. Table 11–1 shows these values for a recommended frame size of 1 ms.

Table 11–1 IEEE 802.16 recommended transmission parameters for a 1-ms frame size (Courtesy of IEEE).

Channel Bandwidth MHz	*Symbol Rate (MBaud)*	*Bit Rate (mbps) QPSK*	*Bit Rate (mbps) 16-QAM*	*Bit Rate (mbps) 64-QAM*	*Recommended Frame Duration (ms)*	*Number of PSs per Frame*
20	16	32	64	96	1	4000
25	20	40	80	120	1	5000
28	22.4	44.8	89.6	134.4	1	5600

11.5 IEEE 802.16A PHYSICAL LAYER DETAILS FOR 2–11 GHZ

The IEEE 802.16a-2003 follow-on standard specifies the physical and MAC layers of the air interface for both point-to-multipoint and optional mesh (multipoint-to-multipoint) broadband wireless access systems. This new standard provides the necessary functionality for wireless access to data, video, and voice services with a specified quality of service. The MAC layer can support multiple physical layer implementations that are each suited to a certain type of operational environment. Each of these air interface technologies provides for the use of adaptive antenna systems, ARQ, and space time coding diversity operation, and one option also allows for wireless mesh network operation. Furthermore, these physical layer specifications may be used in both the licensed bands that have been designated for public network access or with an additional MAC layer modification (DFS support) in the license-exempt bands in the frequency range between 2–11 GHz.

The extension of the IEEE 802.16 standard to cover the lower- microwave-frequency bands requires additional physical layer functionality. In particular, these lower-frequency bands provide an environment where near-LOS and non-LOS operation is possible. At the same time, this fact presents other problems because of the distinct possibility of significant multipath propagation. Therefore, support for sophisticated power management techniques, interference mitigation schemes, coexistence, and multiple antenna technologies becomes important. Additionally, the support of an optional mesh topology means that the standard can now support a form of multipoint-to-multipoint radio air interface. This requires new features and enhancements to both the physical and MAC layers. In this new standard, ARQ operation has been introduced to support system operation over poorly behaving and lossy channels and during optional mesh operation. Lastly, it should be noted that the use of license-exempt bands introduces the probability of additional interference and coexistence issues and at the same time limits the maximum allowed radiated system output power. In an effort to detect and avoid interference under these conditions, dynamic frequency selection (DFS) has been introduced to this specification under the IEEE designation of wireless high-speed unlicensed MAN or WirelessHUMAN.

To summarize, implementations of the standard for licensed frequencies in the 2–11 GHz range may use any of the new physical layers in IEEE 802.16a. Any implementations of this standard in the unlicensed bands in the 2–11 GHz range may also use any of these new physical layer specifications if they also comply with the DFS protocols outlined by the new standard. The next several sections will address these new physical layer specifications.

WirelessMAN-SCa Physical Layer

The IEEE 802.16 WirelessMAN-SCa physical layer option is designed specifically for NLOS operation in the 2–11 GHz frequency range. When used in licensed bands, the permitted channel bandwidth is limited to the local regulatory allowed bandwidth divided by any power of 2 but in any case no less than 1.25 MHz. In an effort to be informative but also brief, the coverage of this topic will focus on the differences introduced by this new specification and downplay the similarities between it and other previously introduced wireless systems.

Figure 11–12 shows a block diagram of both the downlink and uplink transmitting process. This process is similar to that used by the WirelessMAN-SC specification. The data to be transmitted is randomized and typically undergoes various types of complex FEC encoding procedures to increase its resistance to bit errors during transmission. There is also a non-FEC option that uses ARQ for error control. The types of FEC encoding available include a form of concatenated FEC using Reed-Solomon and pragmatic trellis coded modulation (TCM) with optional byte interleaving and FEC options that include the use of block turbo codes (BTCs) or convolutional turbo codes (CTCs). Broadcast messages must use the concatenated form of FEC whereas nonbroadcast messages may use all of the different optional forms of FEC encoding made available by the system.

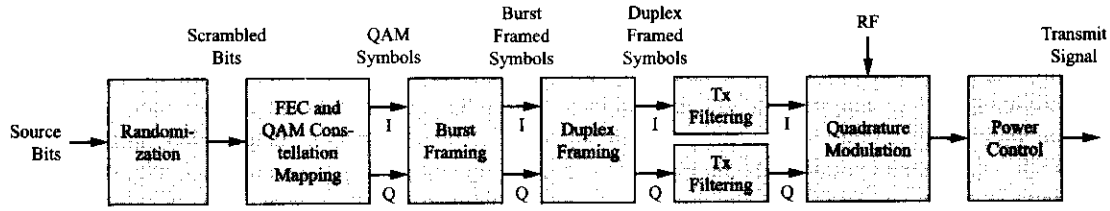

Figure 11–12 WirelessMAN-SCa standard block diagram of transmitting process for both the downlink and uplink (Courtesy of IEEE).

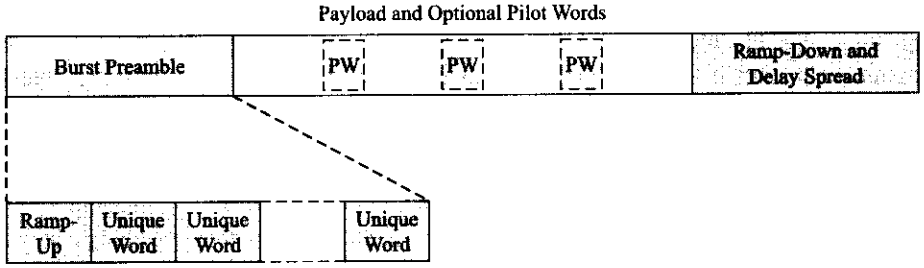

Figure 11–13 WirelessMAN-SCa standard burst frame format (Courtesy of IEEE).

Both downlink and uplink data is then formatted into bursts that employ a framed burst format. The fundamental burst frame is shown in Figure 11–13. This burst frame consists of framing elements that enable improved equalization and channel estimation performance. This is especially important when transmitting data over an NLOS radio link that experiences extended delay spread due to multipath. The burst consists of three framing elements: a burst preamble that includes some ramp-up time and a sequence of unique words, a payload that might include optional pilot words, and a ramp-down and delay spread time period. The use of the beginning and ending frame elements and the optional pilot words within the payload all serve to enhance system operation in the 2–11 GHz air interface environment.

The next step in the transmission process is duplex framing. The WirelessMAN-SCa specification calls for the support of at least one of these duplexing modes (i.e., FDD or TDD). As discussed previously, FDD provides separate channels (carrier frequencies) for the downlink and uplink transmissions. TDD multiplexes the downlink and uplink messages over the same carrier during different time intervals. Figure 11–14 illustrates an example FDD system with burst time division multiplexing downlink payloads. As shown by the figure, the downlink and uplink frames are of equal duration and they repeat at MAC-defined intervals. The downlink payload(s) may not exceed the length of the downlink subframe but they do not have to fill the entire subframe either. The first burst in each downlink subframe is a burst preamble that is directly followed by a frame control header (FCH). The FCH is a broadcast message that may contain DCD, UDC, and MAP information. For this FDD format, time division multiplexed downlink payload data may follow the FCH. Each downlink burst consists of a framed burst as already discussed. Depending upon the type of downlink bursts to be transmitted, time gaps may be used between the individual downlink bursts for improved system operation.

The uplink subframe is similar to that used by the WirelessMAN-SC standard as shown by Figure 11–14. The UL-MAP in the downlink FCH governs the location and the burst size and profile for the bandwidth grants to individual subscriber stations. The burst profile selection is typically based upon the effects of distance, interference, and other environmental factors experienced by the uplink transmissions from the SS to the base station.

The time division duplexing mode multiplexes the downlink and uplink data on the same carrier frequency as shown by Figure 11–15. This operation is also similar to that already described for

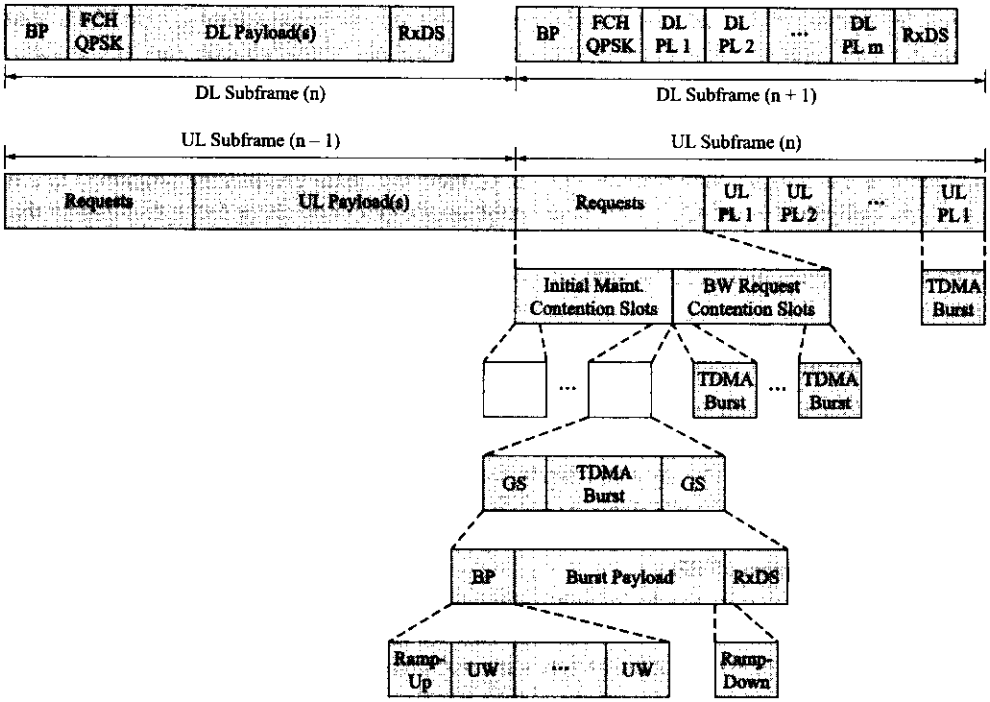

Figure 11–14 WirelessMAN-SCa standard burst formats for FDD operation (Courtesy of IEEE).

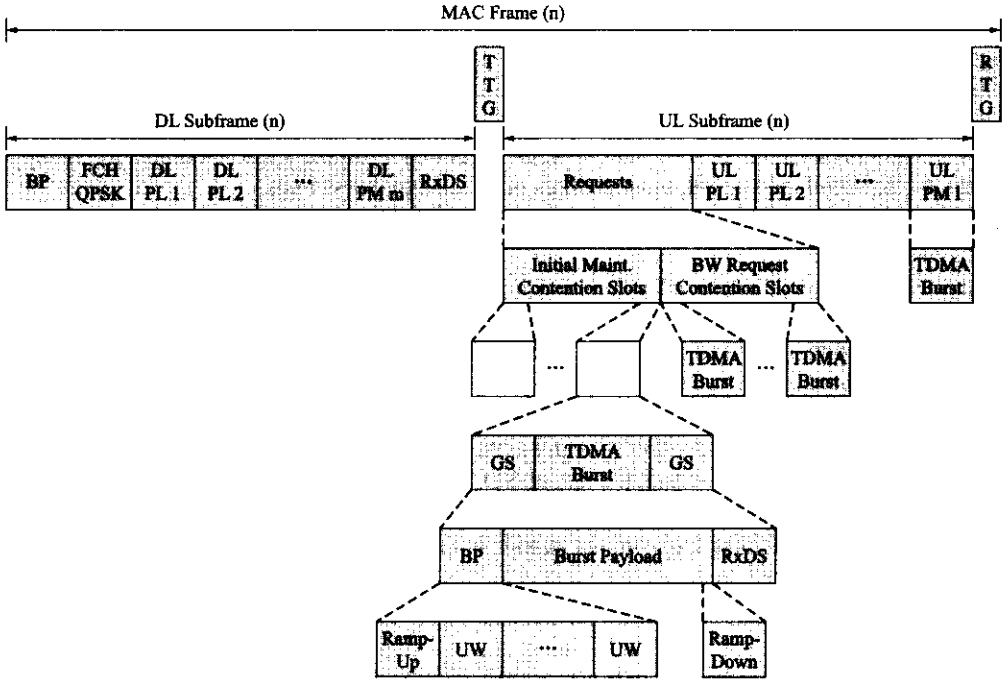

Figure 11–15 WirelessMAN-SCa standard burst formats for TDD operation (Courtesy of IEEE).

WirelessMAN-SC operation. The downlink and uplink subframes share a constant-length MAC frame. The percentage of the MAC frame allocated to the downlink and uplink subframes may vary and is set by information provided in the FCH. A transmitter/receiver (Tx/Rx) timing gap (TTG) is provided between the downlink-to-uplink changeover and a receiver/transmitter (Rx/Tx) timing gap (RTG) is provided between the uplink-to-downlink changeover.

The last steps in the transmission process consist of baseband data prefiltering (shaping) and then the application of the individual I and Q data streams to the quadrature modulator for upbanding to the carrier frequency. The type or level of modulation to be performed is dictated by the particular burst profile. The RF output of the quadrature modulator is applied to an amplifier equipped with a sophisticated power control system.

WirelessMAN-SCa Power Control and Modulation Formats

The WirelessMAN-SCa standard must support power control on the uplink using both initial calibration and periodic adjustments. The base station must be able to make accurate power measurements of the RSSs of an SS's uplink burst signal. The measurements can then be compared against a reference level, and MAC calibration messages may be sent back to the SS to correct any differences. The power control system should be able to respond to power fluctuations that do not exceed a rate of change of 30 dB/second and depths of change of at least 10 dB.

This standard shall provide for the following modulation types on both the downlink and uplink radio links: BPSK (optional on downlink), QPSK, 16-QAM, 64-QAM (optional on uplink), and 256-QAM (optional). As stated, not all of the modulation types are mandatory.

Channel Quality Measurements

To aid in the successful implementation of the WirelessMAN-SCa standard by the air interface components of the system, RSSI and carrier-to-interference-and-noise-ratio (CINR) measurements are taken by the system. The system may require that downlink RSSI measurements be taken by the SS. When mandated by the base station, the SS will collect data on the RSSI from the received downlink burst preambles. From a sufficient number of successive RSSI measurements the SS will derive values for the mean and standard deviation of the RSSI. The base station can ask for these values through the issuance of an REP-RSP message to the SS. In a similar fashion, the mean and standard deviation of the CINR may be measured and calculated by the SS and reported to the base station when requested to do so. These measurements may be used to determine the downlink and uplink burst profiles on a burst-by-burst basis.

Antenna Diversity Schemes

This standard supports the use of advanced antenna system (AAS) diversity schemes to enhance system operation. In particular, a form of space time coding (STC) transmit diversity is supported. In this scheme, logically paired blocks of data separated by delay spread guard intervals are jointly processed at both the transmitter and receiver. The use of frequency domain equalizer techniques to obtain estimates of the transmitter data streams is typically employed for this form of diversity. This subject is beyond the scope of what the author hopes to accomplish with this text and therefore will not be discussed any further at this time. What does need to be pointed out is that the use of STC diversity requires modifications to the framing bursts and in particular the STC frame preamble. All of the details of this type of system operation are laid out in the WirelessMAN-SCa standard and the reader is encouraged to research this topic if further understanding of the topic is desired.

System Requirements

The standard lists many system requirements in terms of frequency accuracy, timing jitter, power level control, spurious emissions, and receiver sensitivity. These last specifications are all given in terms of system BER for the different modulation schemes that may be employed during system operation.

WirelessMAN-OFDM Physical Layer

The IEEE 802.16 WirelessMAN-OFDM physical layer option is based on a form of orthogonal frequency division multiplexing (OFDM) modulation that was designed specifically for NLOS operation in the 2–11 GHz frequency bands. When used in licensed bands, the permitted channel bandwidth is limited to the local regulatory allowed bandwidth divided by any power of 2 but in any case no less than 1.25 MHz. The use of OFDM provides the transmitted data stream with multipath immunity as well as tolerance to symbol time synchronization jitter. OFDM operation consists of the simultaneous transmission of many orthogonal frequency carriers. For this particular implementation of OFDM, the number of carriers is 256 of which 200 are actually used. Of these 200 carriers, 8 are used for the transmission of pilot signals and therefore data is transmitted over the remaining 192. Figure 11–16 depicts the system's OFDM output signal (that, in essence, is a transmitted symbol). A description of this signal would indicate that it contains three types of carriers: data, pilot, and guard or null carriers. The data carriers provide the means by which data is transferred over the air interface, the pilot carriers are used to enhance system operation, and the null carriers are not used for transmission but instead form a lower and upper guard band for the OFDM signal/symbol. There are 28 lower-frequency and 27 upper-frequency guard carriers and an additional null or DC carrier (at the channel center frequency) that are not used to transmit any signals. The carriers are given indices that run from –128 to +127 with pilots carriers located at –84, –60, –36, –12, 12, 36, 60, and 84 (every 24 indices). As an additional feature of this OFDM scheme, the 192 remaining carriers are additionally subdivided into four subchannels of 48 carriers each. Default system operation is for the subscriber station to use all 192 OFDM carriers during normal data transfers. However, if the subscriber stations support subchannelization transmissions on the uplink, either two or four SSs could transmit simultaneously during the uplink subframe using either two or one subchannels apiece, respectively.

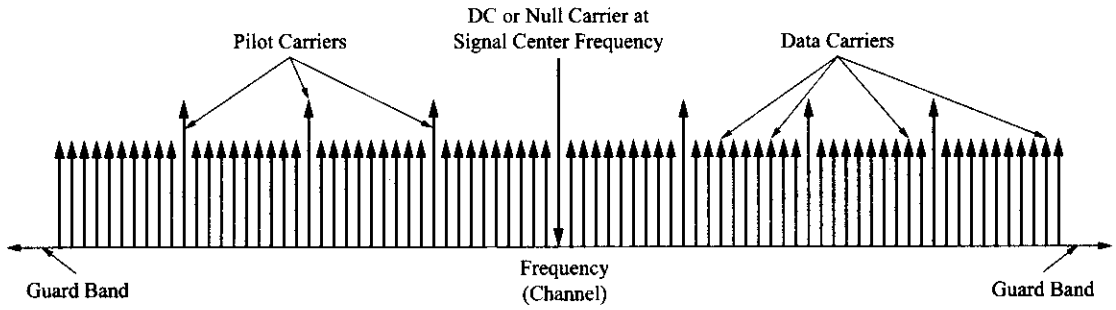

Figure 11–16 OFDM output signal (Courtesy of IEEE).

Channel Coding

Channel coding for the WirelessMAN-OFDM specification is similar to the coding schemes already discussed for the other IEEE 802.16 standards. Three steps are typically involved at both the transmitter and receiver: data randomization, FEC encoding, and interleaving. The order of the complementary operations at the receiver is just the opposite of what was done at the transmitter. The FEC operation is similar to that described earlier, with optional block and convolutional turbo coding possible. After channel coding has been performed, the data bits are grouped and applied simultaneously in smaller groups to constellation mappers (quadrature modulators) that support QPSK, 16-QAM, and optional 64-QAM modulation schemes. For OFDM operation, one should recall how the data to be transmitted is distributed over all the carriers (refer back to Chapter 8 for an example of this operation if needed). An OFDM symbol consists of all the carriers transmitted together simultaneously. The number of bits per symbol depends upon the number of carriers and the modulation coding level. For example, with 192 carriers and 16-QAM modulation used for each carrier, one obtains a data rate of 4 bits/carrier × 192 carriers = 768 bits/symbol.

Frame Structure/Point-to-Multipoint Operation

For use in the licensed bands, a WirelessMAN-OFDM system may use either FDD or TDD. In the unlicensed bands, TDD operation must be used. Figure 11–17 shows the OFDM frame structure for TDD operation. The basic frame interval contains the physical layer PDUs of the base and subscriber stations plus time gaps and guard intervals to facilitate the Tx/Rx and Rx/Tx turnaround times. For TDD, recall that the BS and the SSs all use the same carrier frequency, hence the necessity for these time intervals with a minimum specified duration of 5μs each. As complicated as Figure 11–17 appears, it merely illustrates the typical type of system operation already discussed previously in this chapter. The OFDM physical layer supports frame-based transmission. A frame consists of both downlink and uplink subframes that are not necessarily of equal duration. A downlink subframe contains only one downlink PDU whereas the uplink subframe consists of various contention, maintenance, and bandwidth request intervals and one or more uplink PDUs that come from a single or multiple SSs.

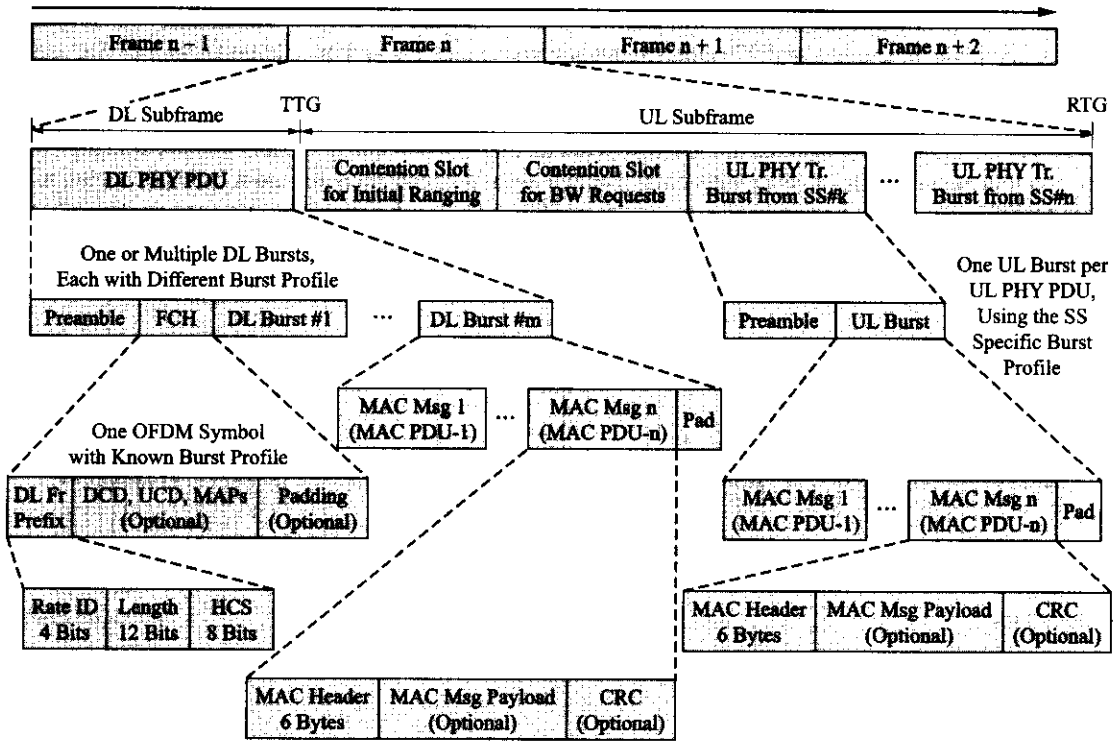

Figure 11–17 OFDM frame structure for TDD operation in the unlicensed bands (Courtesy of IEEE).

The downlink PDU starts with a long preamble that is used by the SSs for synchronization. The preamble is followed by an FCH burst that consists of one OFDM symbol that is transmitted using QPSK modulation for robustness. The FCH contains a downlink frame prefix to specify the burst profile and length of the first downlink burst and a header check sequence (HCS). The FCH may also contain MAC control messages like DCD, UCD, and MAPs. The FCH is followed by one or more downlink bursts that each can have a different profile as indicated by the DL-MAP. The first downlink burst typically contains broadcast MAC control messages. For both the downlink and uplink transmission schemes, each burst will consist of an integral number of OFDM symbols that carry MAC messages or MAC PDUs. Padding may be used to create the integral number of OFDM symbols.

For FDD systems there is no need for the time gaps used during TDD operation since separate carrier frequencies are being used to transmit downlink and uplink information. Figure 11–18 shows the framing structure for this situation. As shown by the figure the downlink and uplink frames are of equal duration. Again, as complicated as Figure 11–18 appears, the operation depicted is extremely similar to what has been already presented.

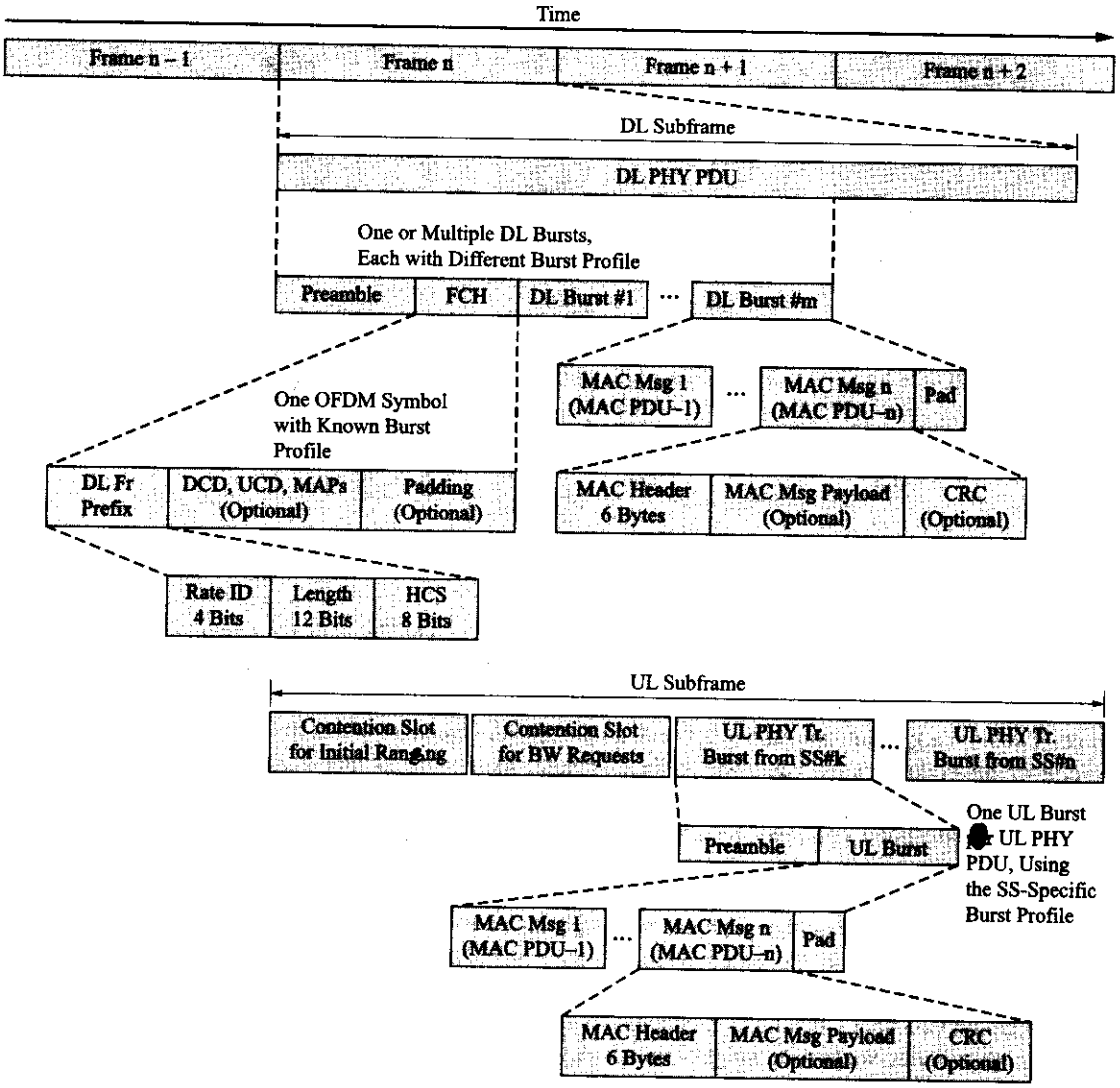

Figure 11–18 OFDM frame structure for FDD operation (Courtesy of IEEE).

Frame Structure/Mesh Operation

Within the WirelessMAN-OFDM specification an optional frame structure has been defined (for TDD operation only) to support the operation of mesh networks. **Mesh networks** are used to provide NLOS operation. The basic difference between point-to-multipoint (PMP) operation and mesh operation is that PMP only allows radio links between the base station and the SSs. In the mesh mode, network traffic can be routed through other SSs and is also allowed directly between SSs (see Figure 11–19) thus providing the

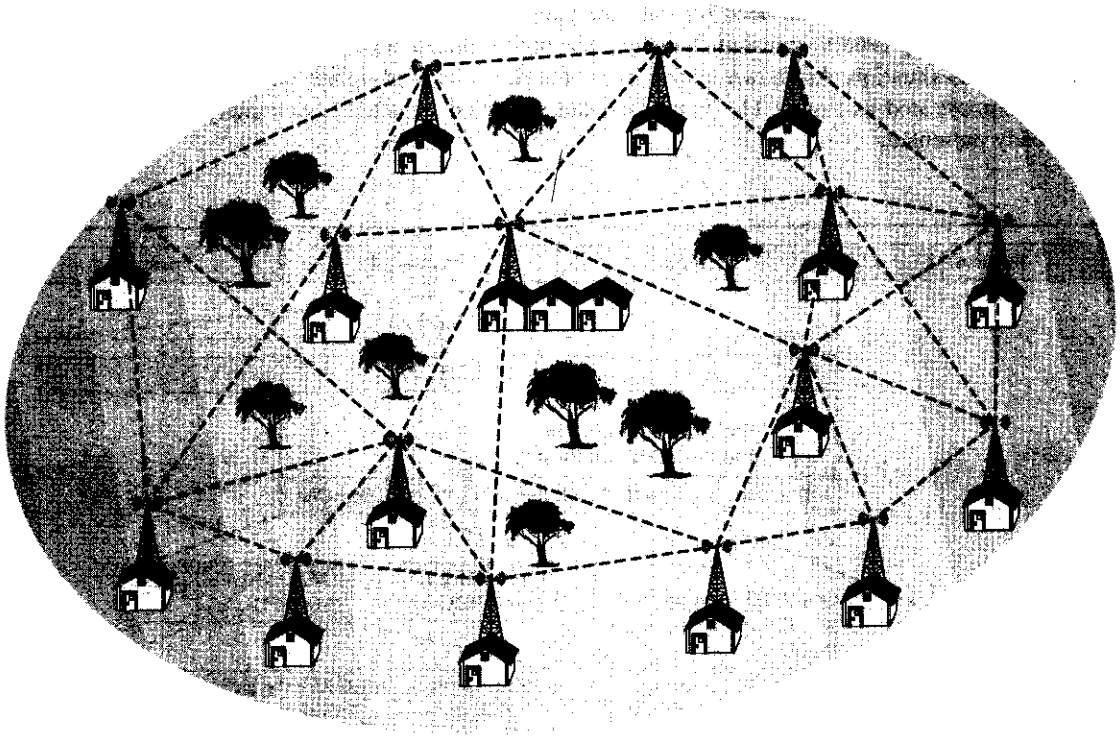

Figure 11–19 A typical wireless mesh network (Courtesy of IEEE).

required connectivity to circumvent many severe NLOS situations. For a mesh network, a base station that services the mesh (has a backhaul connection to the PDN) is typically termed a mesh BS. All other mesh nodes are termed mesh SSs. Several new terms describe the downlink and uplink operations of a mesh network. Traffic away from the mesh BS and traffic in the direction of the mesh BS are now used when describing mesh data transfer operations. Furthermore, there is the notion of mesh neighbors. The stations that a node has a direct link to are known as neighbors. Neighbors of a node form a neighborhood. An extended neighborhood consists of all the neighbors of the neighborhood (i.e., two hops away from the node).

Operation within a mesh network needs to be coordinated. Using a form of distributed scheduling, all the nodes including the mesh BS within a two-hop neighborhood are required to coordinate their transmissions. This process entails the broadcasting of their schedules (i.e., requests and grants) to all of their neighbors. In an optional scheme, the schedule may be determined through uncoordinated requests and grants between two nodes as long as this does not cause other problems (i.e., data collisions or interference) between other nodes within the neighborhood. Another form of scheduling uses a centralized system. In this case, the mesh BS gathers information about all the mesh SS requests within a certain hop range. The mesh BS then provides grant information back to the individual mesh SSs. In a mesh network, all mesh transmissions are done in the context of a link between two nodes. QoS for the link is provisioned on a message-by-message basis. Functions like traffic classification and flow regulation are performed by upper-layer protocols. In mesh networks various antenna systems (from omnidirectional to narrow beam) may be used depending upon the particular circumstances of the node.

To facilitate mesh network operations, the mesh frame consists of a control and data subframes as shown by Figure 11–20. There are two possible types of control subframes: a network control subframe that is used to create and maintain organization between different systems and a schedule control subframe that is used to coordinate the scheduling of data transfers between systems.

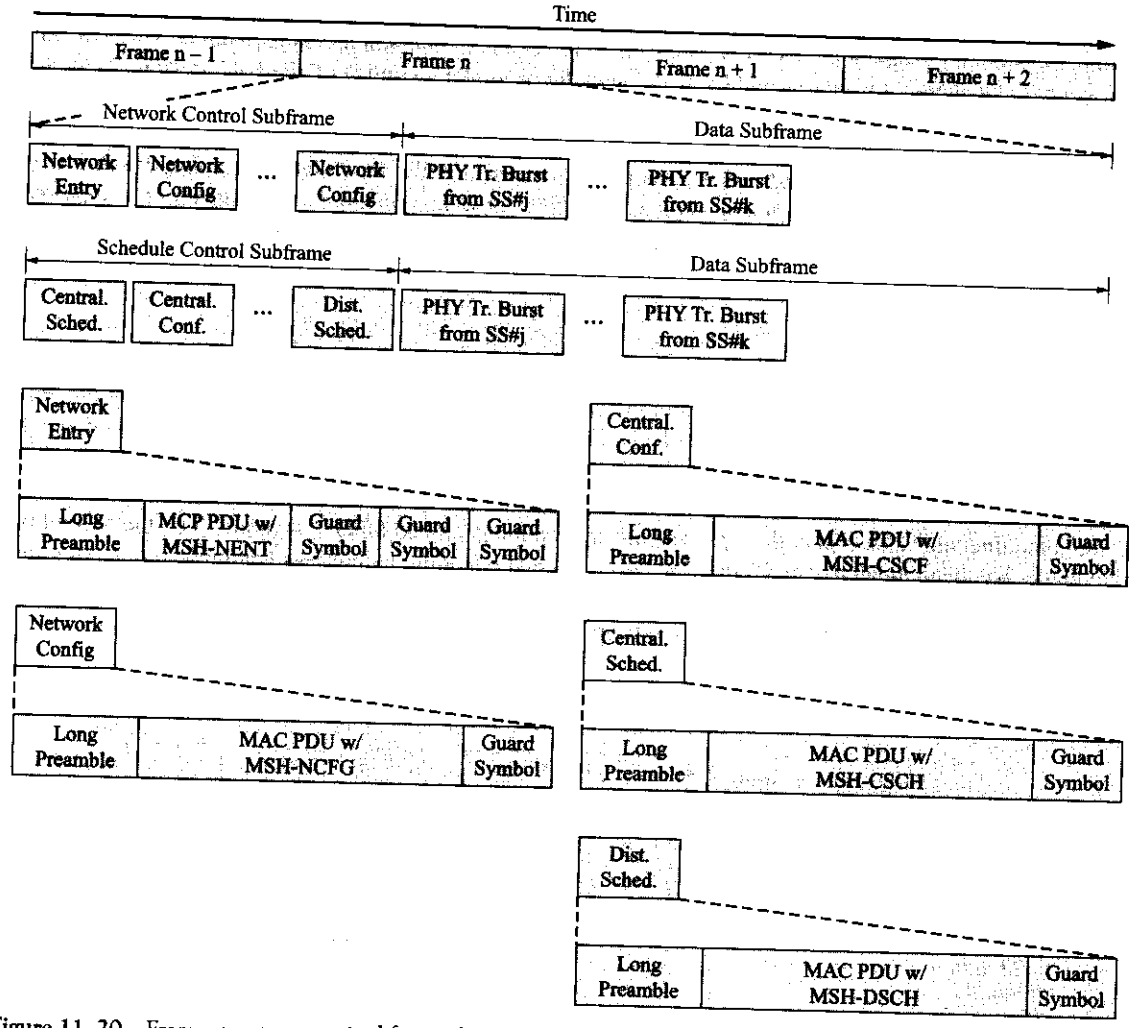

Figure 11–20 Frame structure required for mesh network operation (Courtesy of IEEE).

The network control subframe consists of network entry and configuration bursts. Within a mesh network, mesh network configuration (MSH-NCFG) messages provide a means of communicating information between the network nodes, and mesh network entry (MSH-NENT) messages provide the means for a new node to gain network access and synchronization to a mesh network. The schedule control subframe consists of distributed and centralized scheduling bursts and a centralized scheduling configuration burst. Mesh distributed scheduling (MSH-DSCH) messages, mesh centralized scheduling (MSH-CSCH) messages, and mesh centralized scheduling configuration (MSH-CSCF) messages are all used to facilitate the scheduling of traffic bursts between nodes in a mesh network that occur during the data subframe.

System Control Mechanisms

The WirelessMAN-OFDM specifications call for several system control procedures that are used to enhance system operation. Among these procedures are network synchronization, ranging operations, and power control. A few brief comments about each of these topics will be provided here. The standard recommends that all BSs be time synchronized to a common timing signal. This common timing signal would

typically be provided by a GPS receiver. In the event of the loss of the network timing signal, base stations would continue to operate and would resynchronize to network time when it is recovered. The associated GPS-derived frequency reference could also be used for frequency control of the base station. Ranging is performed during an initial SS registration procedure and then periodically during regular data transmissions. The first ranging operation provides coarse synchronization that must meet an acceptance criterion for new subscribers. The measured range parameters are stored at the base station and then transmitted to the SS for use during operation. During normal operation these measurements are periodically updated to fine-tune the system. If synchronization is lost, the ranging process is performed again during the reregistration process. A form of system power control is employed that is identical to that already discussed for WirelessMAN-SCa operation.

Transmit Diversity Options

Similar to WirelessMAN-SCa operation, an optional form of transmit diversity is supported by the WirelessMAN-OFDM specification. Figure 11–21 illustrates the fundamental concept employed by this technology. There are two transmitting antennas used on the base station side of the radio link and one antenna used on the SS side of the link. Both antennas transmit two different OFDM data symbols (a pair) at the same time and then repeat the process with different but related OFDM symbols (i.e., the second pair of symbols has been generated from the first pair). Special decoding techniques (multiple-input single-output channel estimation) are used at the receiver to achieve what is referred to as second-order diversity. This type of diversity is used to enhance system operation.

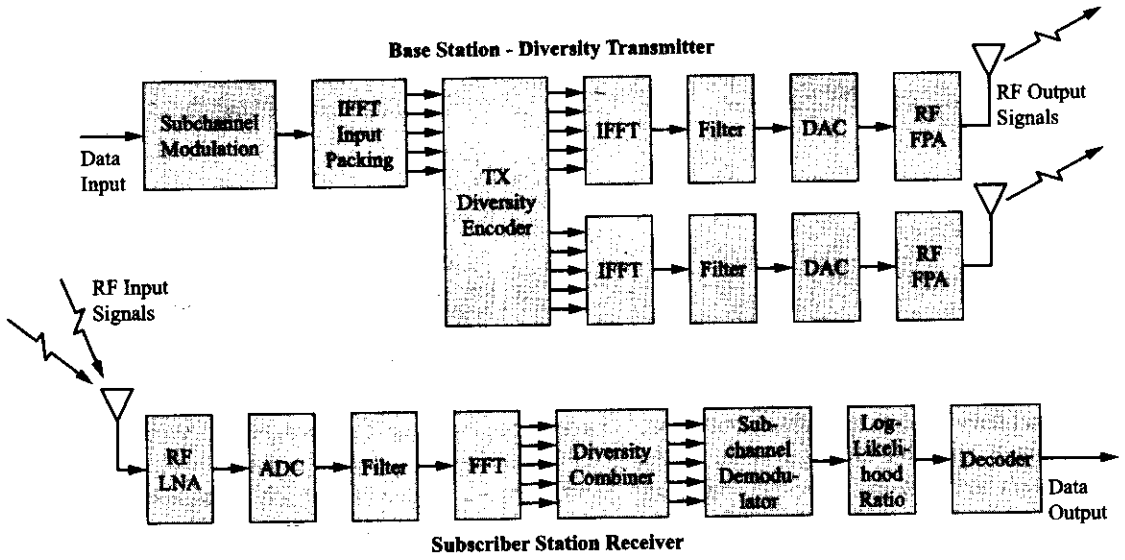

Figure 11–21 Transmit diversity supported by the WirelessMAN-OFDM standard (Courtesy of IEEE).

WirelessMAN-OFDMA Physical Layer

The IEEE 802.16 WirelessMAN-OFDMA physical layer option is a wireless access technique based on OFDM modulation. It is designed specifically for NLOS operation in the 2–11 GHz frequency range. When used in licensed bands, the allowed channel bandwidth is limited to the regulatory provisioned bandwidth divided by any power of 2 but in any case no less than 1.25 MHz. As was the case for WirelessMAN-OFDM, the WirelessMAN-OFDMA symbol consists of data carriers, pilot carriers, and null carriers. Figure 11–22 illustrates the OFDMA symbol structure in the frequency domain. Each OFDMA symbol consists of 2048

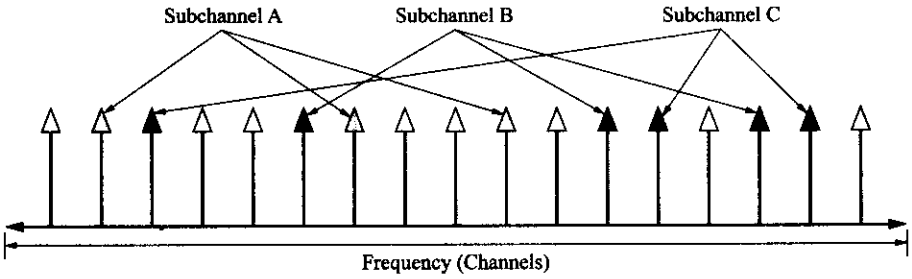

Figure 11–22 WirelessMAN-OFDM symbol structure (Courtesy of IEEE).

carriers of which there are 173 lower-frequency and 172 upper-frequency guard or null carriers, a DC or null carrier at the channel center frequency, 166 pilot carriers, and 1536 data carriers. Furthermore, the 1536 data carriers are subdivided into thirty-two subchannels of 48 data carriers, each within each OFDM symbol.

A few comments about the use of these OFDM subchannels are appropriate. On the downlink, a subchannel may be intended for a particular SS or a group of different SSs. On the uplink, a subscriber station may be assigned a single subchannel or several to many subchannels. This technique allows variable bandwidth allocation and provides the ability for several SS transmitters to transmit their data simultaneously within the same uplink frame. The division of the OFDMA symbol into logical subchannels allows for system scalability, multiple SS access, and advanced antenna array signal-processing capabilities.

OFDMA Slot Definition

In all IEEE 802.16 systems discussed to this point, data is transmitted within a certain physical timeslot or timeslots. However, the OFDMA physical layer slot happens to be two-dimensional. That is, a physical data burst in an OFDMA system is allocated to a group of contiguous subchannels in a group of contiguous OFDMA symbols. This concept may be difficult to grasp initially so several diagrams will be used to illustrate this point. Figure 11–23 shows the basic concept for an OFDMA system with only twelve subchannels per OFDM symbol.

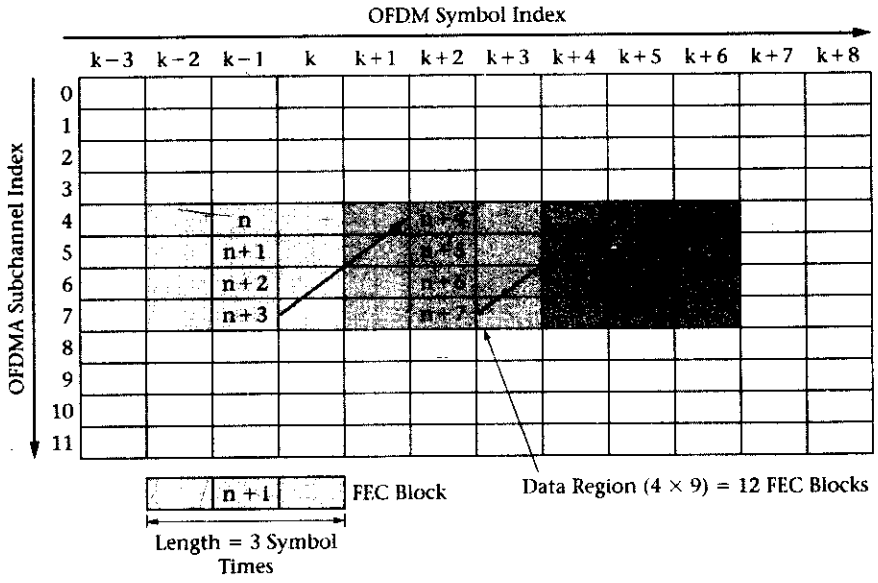

Figure 11–23 OFDMA data region example (Courtesy of IEEE).

As shown by the figure, MAC data is mapped on to an OFDMA "data region" using the following rules. The data is segmented into blocks that are compatible in size to FEC blocks, each FEC block spans one OFDMA subchannel in the subchannel direction and three OFDM symbols in the time direction, and the first FEC block is mapped into the lowest-numbered subchannel in the lowest-numbered OFDM symbol. Furthermore, the mapping is continued in the subchannel direction until the edge of the data region is reached. Then the mapping is continued in the lowest numbered subchannel in the next OFDM symbol. Figure 11–23 shows a mapping example. In this case, the data region dimension is four by nine (4×9). In other words, the OFDMA slot consists of four subchannels that are transmitted for nine consecutive OFDM symbols. The use of this technique spreads the signal out over both time and frequency thus enhancing data transfer over an NLOS radio link.

OFDMA Frame Structure

As is typical for IEEE 802.16a, OFDMA operation over licensed bands supports both FDD and TDD operation and operation over license-exempt bands must use TDD. A typical TDD point-to-multipoint frame structure for OFDMA operation is shown in Figure 11–24. The entire frame structure consists of both downlink base station and uplink subscriber station transmissions. Each downlink burst consists of integer multiples of three OFDMA symbols. The Tx/Rx time gap (TTG) and the Rx/Tx time gap (RTG) are inserted between the downlink and uplink subframes and at the end of the uplink subframe, respectively. After the TTG the base station looks for the first OFDMA symbols of an uplink burst (preamble symbols). Similarly, after a RTG the subscriber stations look for the first symbols of the downlink burst (i.e., QPSK modulated data). For FDD operation there is no need for the TTG and RTG time intervals since downlink and uplink frames are transmitted over different frequencies.

For the downlink subframe a DL frame prefix is transmitted first. This first FEC block of the downlink frame contains information about the FCH and the beginning of the DL-MAP message as shown by Figure 11–24. The DL frame prefix is transmitted with the most robust modulation profile (i.e., QPSK) and is used to provide information about the coding and modulation of the DL-MAP message, the DL-MAP

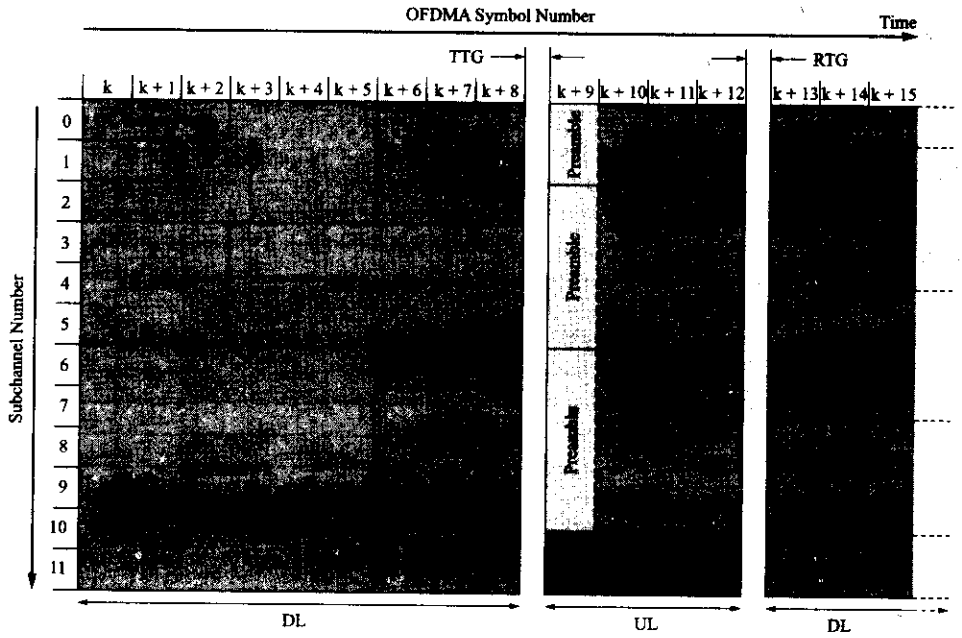

Figure 11–24 Typical TDD frame structure for OFDMA (Courtesy of IEEE).