


BLUETOOTH REVEALED

THE INSIDER'S GUIDE TO AN OPEN SPECIFICATION
FOR GLOBAL WIRELESS COMMUNICATIONS

BRENT A. MILLER • CHATSCHIK BISDIKIAN



Library of Congress Cataloging-Publication Data

Miller, Brent A.

Bluetooth revealed / Brent A. Miller, Chatschik Bisdikian.

p. cm.

Includes bibliographical references and index.

ISBN 0-13-090294-2

1. Bluetooth technology. 2. Telecommunication—Equipment and supplies.

3. Computer Network protocols. I. Bisdikian, Chatschik. II. Title.

TK5103.3 .M55 2000

004.6'20--dc21

Editorial/production supervision: *Kathleen M. Caren*

Acquisition Editor: *Mary Franz*

Editorial Assistant: *Noreen Regina*

Marketing Manager: *Bryan Gambrel*

Manufacturing Manager: *Maura Zaldivar*

Cover Design: *Nina Scuderi*

Cover Illustration: *Tom Post*

Cover Design Director: *Jerry Votta*

Series Design: *Gail Cocker-Bogusz*



© 2001 Prentice Hall PTR

Prentice-Hall, Inc.

Upper Saddle River, NJ 07458

All rights reserved. No part of this book may be reproduced, in any form or by any means, without permission in writing from the publisher

The publisher offers discounts on this book when ordered in bulk quantities.

For more information, contact

Corporate Sales Department, Prentice Hall PTR

One Lake Street

Upper Saddle River, NJ 07458

Phone: 800-382-3419; FAX: 201-236-7141

Email (Internet): corpsales@prenhall.com

All products or services mentioned in this book are the trademarks or service marks of their respective companies or organizations.

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

ISBN 0-13-090294-2

Prentice-Hall International (UK) Limited, *London*

Prentice-Hall of Australia Pty. Limited, *Sydney*

Prentice-Hall Canada Inc., *Toronto*

Prentice-Hall Hispanoamericana, S.A., *Mexico*

Prentice-Hall of India Private Limited, *New Delhi*

Prentice-Hall of Japan, Inc., *Tokyo*

Pearson Education Asia Pte. Ltd.

Editora Prentice-Hall do Brasil, Ltda., *Rio de Janeiro*

*To Laurie, Benjamin and Andrew;
Teresa, Eugene and Theodore;
and King Harald Bluetooth*

and mobile telephones rely on wired “landline” telephone systems to carry and route calls between endpoints.

Communicating without wires is not a new concept. Broadcast radio and television are two common examples of wireless communications; others include satellites, cordless and cellular telephones and remotely controlled televisions, garage door openers and automobile door locks. While most of these examples employ communication via radio waves, the use of infrared, a nonvisible spectrum of light, is also relatively common. Bluetooth wireless communication employs *radio frequency* (or RF) technology, using radio waves to communicate through the air in a manner fundamentally similar to broadcast radio or television.

Radio Frequency Wireless Communications

RF technologies employ transmitters and receivers tuned to produce and consume, respectively, radio waves of a given frequency range. The transmitter’s power and the receiver’s sensitivity help to determine the distance over which they can communicate. High transmission power output is used for long-range communications such as broadcast television while short-range communications generally require much less power; thus technologies that are designed to communicate across only a few meters could be employed in small, mobile battery powered devices. Another characteristic that is relevant for communication applications is the ability of radio waves to penetrate many objects. Obstacles reflect light waves used in technologies such as infrared, but radio waves used in RF technologies in general can (with certain limitations) penetrate many obstacles (although in some cases radio waves can diffract or go around objects too). Thus RF technologies can permeate many obstacles such as clothing, bodies, walls, doors and the like. This means that there is no requirement for a “line of sight” between the transmitter and the receiver.

RF technologies use frequency modulation to generate radio waves within a certain frequency spectrum, which encode information and can be intercepted by receivers tuned to the corresponding frequency. FM radio broadcasts, for example, operate in the 88 megahertz (MHz) to 108 MHz frequency spectrum; some cordless telephones operate in the 900 MHz frequency spectrum; Bluetooth wireless communications and other technologies operate in the 2.4 gigahertz (or GHz; one gigahertz equals one billion cycles per second) spectrum.

Because the usable radio frequency space is finite, most governments regulate its use, partitioning frequency ranges and granting licenses to transmit at those frequencies at specified power levels. In the United States, for example, a federal license is required to transmit in the FM radio frequency spectrum except at extremely low power levels that limit the range to no more than about 30 meters. Some frequencies are reserved for use without a license under certain conditions. For example, in the United States unlicensed operation is permitted, with some restrictions, in the 900 MHz and 2.4 GHz frequency ranges (the latter being where Bluetooth wireless communication operates). In fact, through multinational agreement, the 2.4 GHz spectrum requires no license for its use anywhere in the world.¹ The SIG together with other organizations such as the IEEE 802.11 standards body is working with regulatory authorities in some countries to pursue harmonization of the frequency assignments for unlicensed use within the 2.4 GHz spectrum and of the approval process for wireless communications. In general the chosen frequency spectrum can be used globally without license so long as the rules for operating within this spectrum are followed.

RF Communications in the 2.4 GHz Frequency Spectrum

While the 2.4 GHz spectrum is globally unlicensed, there are regulatory requirements and other considerations for its use. These include:

- The spectrum is divided into 79 channels (although in some countries, notably Spain and France, only 23 channels were available for use in the year 2000. Japan began using all 79 channels in 2000).
- Bandwidth is limited to 1 MHz per channel.
- Frequency hopping spread spectrum communications (described below) must be employed.
- Interference must be anticipated and appropriately handled.

Other RF communications technologies also use this spectrum; notable among them are HomeRF™ (an open industry specification for RF communications in home environments; see <http://www.homerf.org>) and the Institute of Electrical and Electronics Engineers (IEEE) 802.11 specification for wireless local area networks (see <http://www.ieee.org>). Microwave ovens also operate within this frequency range. Because this

1. In some countries there are restrictions and only part of this spectrum is available for unlicensed use; these restrictions are discussed elsewhere in the book, notably in Chapter 6.

spectrum is unlicensed, new uses for it are to be expected (for example, a new generation of cordless telephones also uses the 2.4 GHz frequency) and as the spectrum becomes more widely used, radio interference is more likely. Thus the requirement to anticipate and address interference in the 2.4 GHz range is important for all technologies that operate in it.

Each technology using this spectrum has made design choices within the spectrum's constraints that optimize that technology for particular applications or domains. Bluetooth wireless communication is designed to take maximum advantage of the available channel bandwidth and to minimize RF interference and its effects while operating at very low power.

Spread Spectrum RF Communications

Within RF communications, *spread spectrum* refers to dividing the available spectrum based upon frequency, time, a coding scheme or some other method. Messages to be sent are then divided into various parts (packets) that are transmitted across the divided spectrum. Frequency division spread spectrum (or *frequency hopping*), which is the method employed with Bluetooth wireless communication, divides the spectrum into different frequencies, or channels.² A single message packet is transmitted on a selected channel, then the radio selects a new channel (a process called *hopping* to a new frequency) to transmit the next packet, and the process repeats, thereby spreading the message across the available frequency spectrum. Each technology specifies its own method for establishing the frequency hopping pattern. Obviously the receiver(s) of the message must know the hopping pattern to tune to the correct channels in succession to receive each packet and assemble the complete message. This process is called *frequency hopping spread spectrum*, or FHSS.

FHSS introduces additional complexity as compared to using a single statically selected frequency, yet it also supplies some benefits. First, RF interference can be reduced since all radios hop (often randomly or at least pseudorandomly, and often rapidly) from one frequency to another. When all of the participants in the spectrum employ FHSS, interference caused by colliding transmissions on the same frequency is less likely than it would be if each radio used a single channel

2. Contrast frequency hopping spread spectrum with direct sequence spread spectrum, which is not examined here. Direct sequence is another form of spread spectrum RF communication employed in other technologies such as wireless LANs and is outside the scope of this book.

for a long duration. In addition, when collisions do occur, their effects are lessened, since only a single packet is lost and that packet could be retransmitted at a new frequency, where again it is less likely to encounter interference. Second, FHSS can provide a degree of security for communications in that only a receiver that knows the frequency hopping pattern can receive and assemble all the packets of a message. Because the hopping pattern for a given spectrum could be constructed in a number of ways, it could be difficult to deduce and follow an unknown hopping pattern, especially when the spectrum is heavily utilized with many radios. Thus FHSS can be employed to hinder eavesdropping. In fact, this latter characteristic led to the invention of FHSS, usually attributed to George Antheil and Hedy Lamarr (the latter is more famous as an American actress). Their 1942 patent of the frequency hopping concept was motivated by an attempt to find a "secret communication system" using radio waves to control torpedoes during World War II.³

As previously noted, the use of spread spectrum is required in the 2.4 GHz range, largely to minimize interference problems because the spectrum is unlicensed. The design for Bluetooth wireless communication employs relatively rapid frequency hopping (nominally 1,600 times per second) and is described more fully below and in Chapter 6.

Infrared Wireless Communication

RF is not the only form of wireless communication. Infrared technology is used with devices such as notebook computers, personal digital assistants and electronic remote controls. Infrared wireless communication makes use of the invisible spectrum of light just beyond red in the visible spectrum.

In particular, one standard method for infrared communication is specified by the Infrared Data Association (IrDA; see <http://www.irda.org>); this method is commonly used with mobile phones and notebook and handheld computers. IrDA technology is relevant when discussing Bluetooth technology because IrDA is also designed for short-range, low-power unlicensed communications. IrDA also defines

3. The complete story of this invention is fascinating but is outside the scope of this book. Interested readers are referred to, for example, [IAL99] or other accounts easily found via World Wide Web search engines. Furthermore, any rationale or implications of the choice of naming the Bluetooth technology after a Danish king rather than an American actress are not explored here.

As shown in Figure 6.1, a Bluetooth *device*¹ represents the complete physical entity (such as a notebook computer, a digital phone, an information appliance, and so on) that contains applications that can communicate using the Bluetooth wireless technology incorporated in that device. Here we assume that a device is associated with one and only one transport protocol group implementation and air-interface. However, as mentioned in the previous chapter and further elaborated in the next one, the L2CAP layer supports the multiplexing of several higher-layer protocols, like RFCOMM, SDP, and so on, and hence middle-ware protocol stacks and applications, as well.

The design point of the Bluetooth transport protocols is dictated primarily by low manufacturing complexity, associated low cost and fast time to market. For this reason, a low-cost, easy-to-implement, frequency-hopping spread-spectrum radio solution was selected. In addition, the selection of a master/slave architecture for the baseband transmission was dictated by the ad hoc nature of the systems considered. In particular, Bluetooth piconets can form spontaneously among disparate devices with vastly different power and computing capabilities, unlike wireless LAN systems which are designed primarily to support a wireless extension of LAN connectivity services to relatively "power-unconscious" personal computers.

To support ad hoc connectivity with minimal (if any) state maintained by the Bluetooth devices, piconets are dynamically formed in isolation without the support of third-party (infrastructure) control signaling. For as long as needed, the master of a piconet serves as the control point for the communications on the piconet. For the duration of the existence of a piconet, the operation of a master resembles that of a base station of a picocellular system. Thus, the Bluetooth technology permits the spontaneous creation of a temporary picocellular system, where traffic is regulated by a spontaneously created base station, the master, that regulates the traffic to and from the other member of the picocell, the slaves. Recall from previous chapters that the Bluetooth technology permits the creation of multiple, ad hoc picocellular systems that remain operational even in cases of temporal and spatial overlapping. The use of master and slaves reduces the complexity of the design and thus the cost of deploying the Bluetooth technologies.

1. The terms *unit*, *node*, *station*, and so on may (and have) also been used instead of the term *device*. Insofar as possible, however, the use of these other terms is avoided. The reason for this choice will become apparent later.

The rest of this chapter highlights the Bluetooth radio; the link controller and the baseband functions it controls, like piconet creation and the medium access protocol; and the link manager that configures and manages the links between devices. In the reference implementation of a Bluetooth module, the radio and the link controller typically are implemented in hardware while the link manager is in firmware. As such, the radio and the link controller (with its baseband functions) were the first parts of the specification to mature, providing a stable enough hardware specification for chip designers to use in designing their early Bluetooth chipsets. The link manager specification originally was focused primarily on security-related transactions. The SIG's further work on the link manager specification has enriched it over time to take full advantage of the capabilities available in the baseband specification.

The Bluetooth Radio

The Bluetooth system operates in the 2.4 GHz *industrial, scientific, and medical* (ISM) band. The ISM bands are license-free bands set aside for use by industrial, scientific and medical wireless equipment. Regulatory authorities around the world have opened the ISM bands for use by low-power systems that can be operated without the need for a license but nevertheless under strict regulations. In the United States, these regulations are set by the Federal Communications Commission (FCC) and are detailed in the *Code of Federal Regulations part 15* [FCC99]; section 15.247 of the FCC regulations deals with the operational rules of *intentional radiators* operating in the various ISM bands including the 2.4 GHz band.

The 2.4 GHz ISM band is a globally available radio band, albeit not frequency harmonized. Table 6.1 shows the ISM frequency availability in the majority of countries around the globe.

Table 6.1

License-free frequency allocation in the 2.4 GHz band.

2.4 GHz ISM band (MHz)	Frequency channels (MHz) $k = 0, 1, \dots, m-1$	LGB ¹ (MHz)	UGB ² (MHz)
2,400.0 – 2,483.5	$2,402 + k; m = 79$	2.0	3.5

1. "LGB" stands for "lower guard band."
2. "UGB" stands for "upper guard band."

The radio part of the specification consists mostly of a series of design specifications for Bluetooth transceivers, like in-band and out-of-band spurious emissions, frequency accuracy, co-channel and adjacent-channel interference, out-of-band blocking, intermodulation characteristics, and so on. The selection of the radio design specifications is driven by the requirement to allow the development of high-quality, low-cost transceivers that comply with the various 2.4 GHz ISM band regulations around the world.

The numerous design parameters plus the radio testing conditions are not repeated here. Readers interested in the Bluetooth radio design can find an abundance of information in the corresponding part of the specification. Since no new data-exchange protocols nor data-processing algorithms were developed for the Bluetooth radio, the radio part of the specification was the first to be completed. The errata to this portion of the specification are minimal, and the few that surface pertain mostly to the accommodation of new regulations or refinement of the radio testing procedures.

The Bluetooth transceiver is a *frequency-hopping spread-spectrum* (FHSS) radio system operating over a number m of 1 MHz-wide channels. While for the majority of countries $m = 79$, as shown in Table 6.1, regulations in certain countries may further constrain the license-free frequency spectrum for the 2.4 GHz ISM band. Thus, the Bluetooth radio (and the baseband described in the next section) design can accommodate two alternatives that operate over 79 or 23 channels, each one of which is 1 MHz wide. For frequency-hopping systems operating in the 2.4 GHz ISM band, the FCC part 15.247 regulations restrict the maximum peak output power of the radiator to no more than 1 watt (30 dBm). Moreover, at least 75 out of the 79 frequency channels must be used pseudo-randomly with a total residence time in each of the frequencies not to exceed 0.4 seconds within a 30-second period. A Bluetooth radio utilizes the maximum number of channels available, 79 in the United States, and it hops at a high rate, 1,600 hops per second, pseudo-randomly across all these frequencies to achieve high noise resilience. The use of *direct-sequence spread-spectrum* (DSSS) systems, which are also permitted in the 2.4 GHz ISM band, may be prohibitively costly for the low-cost requirement of Bluetooth radios.

Figure 6.2 summarizes the key operations in the Bluetooth radio. The figure also shows a pair of logical interfaces for carrying data and control information between the radio and the rest of the Bluetooth system. Here data relates to all information that is transmitted or received over the air. The control information controls the behavior of the radio

On the transmit side, this includes the carrier frequency to which the transmitter needs to tune prior to transmission of a bit stream of information over the air and the power level that is to be used for this transmission. On the receive side, this includes the carrier frequency to which the receiver must tune for receiving a bit stream of information and (optionally) the strength of the signal being received. Power-supply and time-signaling lines are not shown in the figure. A standardized set of interfaces for the data and control information is not provided in the specification. Thus, chip designers and manufacturers can choose to integrate the radio component with the rest of the components of a Bluetooth module in the way they believe is best for a low-cost, power-efficient system. Table 6.2 summarizes some of the key operational parameters of the Bluetooth radio.




Figure 6.2
The Bluetooth radio operations.

Table 6.2

Key operational parameters of the Bluetooth radio specification.

Modulation	Gaussian frequency-shift keying (GFSK)	BT product ¹ : 0.5; modulation index: 0.28 – 0.35
Symbol rate	1 Msymbol per second	using the binary GFSK, this translates into 1 Mbps raw link speed; bit transmission time: 1 μ sec
Frequency-hopping rate	1,600 hops per second, typical	residence time: 625 μ sec per hop
	3,200 hops per second for inquiries and pages	residence time: 312.5 μ sec per hop
Transmit power	Class 3: 0 dBm (1 mW)	a typical Bluetooth radio; optional power control to below -30 dBm
	Class 2: 4 dBm (2.5 mW)	optional power control as above
	Class 1: 20 dBm (100 mW)	required power control to at least 4 dBm; optional power control as above
Receiver sensitivity	a Bluetooth receiver must attain a raw <i>bit error rate</i> (BER) of 0.1% with an input signal level of -70 dBm or lower	the -70 dBm sensitivity level shall be attained for any input signal generated by any compliant Bluetooth transmitter

1. The term "BT product" is not short for "Bluetooth product." It is a parameter describing the quality of transmitted waveforms expressed as the product of the bandwidth of the modulation filter and the bit time.

The transmit power and receiver sensitivity values are examples of design decisions that were made to reduce cost and power requirements for the Bluetooth radio. An IEEE 802.11 wireless local area network (WLAN) may use up to 30 dBm (1000 mW) of transmit power in the United States (lower power-levels in other parts of the world) and not less than 0 dBm (1 mW). Hence, an 802.11 wireless solution may not be suitable for many power-constrained personal, portable devices. The sensitivity level is also much lower than that of an 802.11 radio receiver². This means that a simpler Bluetooth radio can be built at a reduced cost as compared to an 802.11 radio.

The Link Controller and Baseband

As discussed in the previous section, the radio deals with the acts of sending and receiving data over the air. Outside the scope of the radio's responsibilities are considerations such as what data to transmit and when, what data to wait for and when, and which carrier frequency and transmit power to use. These are the responsibilities of the Bluetooth link controller, described later in this chapter, that executes the baseband communication protocol and related processes.

Figure 6.3 summarizes the key functions of the Bluetooth baseband. These include piconet and device control functions like connection creation, frequency-hopping sequence selection and timing; modes of operation such as power control and secure operation; and medium access functions like polling, packet types, packet processing and link types. These items are detailed later in the chapter as we proceed through the operational phases of a Bluetooth device. The figure also shows a set of logical interfaces for carrying data and control information between the baseband and the rest of the Bluetooth system. For the same reasons mentioned in the radio section, the specification avoids specifying any standardized set of interfaces for the data and control information. Nevertheless, their presence, even on a "logical" plane, aids this presentation, since it allows us to concentrate on the baseband operations in isolation from the rest of the Bluetooth system.

2. The receiver sensitivity for an IEEE 802.11b radio receiver, which uses DSSS, is specified to be -80 dBm for a frame error rate of 8% for a 1024-byte MAC protocol data unit when a 2 Mbps DQPSK modulation is used. For more information on the IEEE 802.11 WLANs see [IEEE99]




Figure 6.3
The baseband functions.

Due to the vast and diverse areas covered by the baseband specification, the remainder of this section tries to combine the information flow in the specification with that of more traditional communication protocol standards and related articles. The discussion progresses step-wise, identifying the key components over which Bluetooth communications actually occur. We begin with the definition of a piconet over which Bluetooth devices can exchange information packets. We then describe fundamental helper elements that assist in the creation and maintenance of the piconet, such as the Bluetooth clock and frequency-hopping selection process. Using these helper elements, the presentation continues with the description of the procedures that Bluetooth devices follow to create and join a piconet. Finally, with a piconet in place, we focus on protocols and processes like medium access control,

error control, security, power-saving operation and so on, that are exercised on the Bluetooth devices and on information packets sent and received over the piconet.

The Piconet

For devices to communicate with each other using Bluetooth wireless technology, they need to be part of a *piconet*. Simply stated, a piconet comprises a shared communications channel through which members of the piconet communicate. In the FHSS space in which Bluetooth radios operate, this communication channel consists of a well-defined sequence of frequency hops pseudo-randomly selected from the set of frequencies in Table 6.1 at a nominal rate of 1,600 hops per second. The members of the piconet are able to follow the successive hops of the frequency-hopping sequence in a synchronized manner. Piconets are formed as needed and endure for as long as participating devices need to communicate. The baseband operations define how a frequency-hopping sequence for a piconet is created, how devices learn how to follow it and hence join the piconet, and how to send and receive information packets communicated in a coordinated manner between devices in the piconet.

Bluetooth piconets are formed in a rather ad hoc manner. They are formed on demand among devices that want to communicate with each other without relying on the services of a dedicated support entity, such as a *base station* in a cellular network or a corporate or home WLAN. The baseband protocol establishes the rules according to which these ad hoc connections are established such that devices communicate in a coordinated and efficient manner.

The frequency-hopping sequences that define the communication channels in piconets are highly uncoordinated, or, to be more precise, they are created in a manner that makes them appear highly uncoordinated. Thus, owing to the frequency-hopping nature of transmissions in a piconet, multiple piconets may exist in time and space with minimal interference among themselves. Multiple piconets overlapping, at least partially, in time and space are referred to as *scatternets*. This opens the possibility of interpiconet communications when devices become members of multiple piconets.

Each piconet has one and only one *master* and one or more *slaves*. These roles are temporary ones and they are meaningful only while Bluetooth devices are members of a piconet. Certainly, Bluetooth devices may be built to operate only as masters or only as slaves, but this is a host application and usage scenario issue rather than a Bluetooth

specification issue. The specification generally assumes that a Bluetooth device is capable of acting both as a master and as a slave, depending upon the role required to accomplish a given usage case. In the case of a scatternet, a device that participates in more than one piconet can be the master of at most one of these piconets, while it can be a slave in several of them. Through a detailed process, described later in this chapter, a master and a slave may exchange their roles. It is also possible for an "old" piconet based around an original master to be migrated to a new piconet with a new master. Piconet migration as well as communication across scatternets are not discussed in detail here, because they are not very mature processes in the version 1.0 specification, which does not define any usage scenarios that address communication across piconets. Nevertheless, the specification contains necessary considerations that could allow these processes to be accomplished.

The primary role of a master for a piconet is to define:

- which frequency-hopping sequence the members of this piconet shall follow;
- when frequency hops shall occur, thus defining the timing foundation for timed events in the piconet;
- which particular frequency is the "current" frequency; and
- which slave will be transmitted to and/or which slave will be permitted to transmit next (recall that transmission on the Bluetooth air-interface is through polling of the slaves).

The first three items are strongly associated with how a piconet is formed and maintained and how devices join it. The last item is associated with how transmissions occur in a piconet.

Figure 6.4 shows the operational states for a Bluetooth device. In the *connected* state, the device is a member of a piconet. On the other hand, when a device is not associated with any piconet or participates in no action that could result in its forming or joining a piconet, it is said to be in the *standby* state. The *standby* state is the default operational state for a Bluetooth device. In this state, a device typically idles with only its native clock operating in a low-power mode.

To move to the connected state, a device goes through the inquiry and page states, which are instantiated differently but in a complementary manner within a potential master and a potential slave. In the *inquiry* state, a device learns about the identity of other devices in its vicinity; these other devices must be in an *inquiry scan* state to listen for and subsequently respond to inquiries. In the *page* state, a device explicitly invites another device to join the piconet whose master is the invit-

ing device; the other device must be in the *page scan* state to listen for and subsequently respond to pages. As Figure 6.4 shows, a device may bypass the inquiry state if the identity of a device to be paged is already known. The figure also implies that while a device is a member of a piconet, it may still perform inquiries and pages for additional devices to join this or some other piconet. In the latter case, a scatternet (multiple piconets overlapping in time and space as discussed in Chapter 2) eventually could be created.




Figure 6.4
Operational states of a Bluetooth device.

To become a member of a piconet, a Bluetooth device needs to know how to recreate the frequency-hopping sequence that defines that piconet and which frequencies from this sequence will be visited and when. Also, to participate in communications over the piconet, the device needs to know how to formulate, read and write information packets on the piconet. All of these operations, as well as nearly every other operation in a Bluetooth device, are related to the following two *fundamental elements*:

- the Bluetooth device address
- the Bluetooth device (or native) clock

Any process in the Bluetooth baseband is intimately related to these two elements. Two baseband processes are notable and we refer to them as the *fundamental processes*, which are those that generate:

- the frequency-hopping sequence, and
- the access code.

These fundamental elements and fundamental processes are detailed below.

The Bluetooth Device Address (*BD_ADDR*)

The Bluetooth device address (*BD_ADDR*)³ is the most static entity of a Bluetooth device. The *BD_ADDR* is a single 48-bit address electronically “engraved” on each device. A device’s *BD_ADDR* is globally unique among Bluetooth devices. To guarantee uniqueness, a numbering authority assigns *BD_ADDR*s. The *BD_ADDR* is an IEEE 48-bit address, similar to the *medium access control* (MAC) address of IEEE 802.xx LAN devices.

The 48-bit address field, shown in Figure 6.5 from the least significant bit (LSB) to the most significant bit (MSB), is partitioned into three parts: the *lower address part* (LAP), the *upper address part* (UAP), and the *non-significant address part* (NAP). The 24 bits of the UAP and the NAP constitute the *organization unique identifier* (OUI) part of the address that is assigned by the numbering authority to different organizations. The LAP is assigned internally by various organizations. The various parts of the *BD_ADDR* are involved in nearly every operation of the baseband from piconet identification, to packet header error checking, to authentication and encryption key generation. These items are discussed in more detail later in this chapter.




Figure 6.5
The Bluetooth 48-bit device address (*BD_ADDR*).

The Bluetooth Clock

Each Bluetooth device has a free-running 28-bit (native) Bluetooth clock. The Bluetooth clock is never adjusted and is never turned off. The clock ticks 3,200 times per second or once every 312.5 μ sec, representing a clock rate of 3.2 KHz. Notice that this is twice the nominal frequency-hopping rate of 1,600 hops per second. The clock has an accuracy of ± 20 ppm⁴. In low-power modes like standby, hold, and park (introduced in Chapter 2 and detailed later in this chapter), lower

3. The use of the notation *BD_ADDR* in the specification is the main reason for choosing the term “device” in Figure 6.1 over any of the other legitimate alternatives identified in footnote 1 of this chapter.
4. The abbreviation “ppm” stands for “parts per million” and it is a metric of clock accuracy, where the smaller the value the more accurate the clock.

power consumption can be achieved by using a low-power oscillator to drive the clock at a reduced accuracy of ± 250 ppm. The Bluetooth clock is illustrated in Figure 6.6.




Figure 6.6
The Bluetooth native clock.

The Bluetooth clock wraps around just short of once per day. The Bluetooth clock plays a fundamental role in deciding when a device can or cannot transmit or listen for a transmission, and at which frequency and for what types of information packets it transmits or listens. A slave device uses the value of the Bluetooth clock of a master to accomplish piconet communications as discussed below. The significance of the time intervals shown in Figure 6.6 will become apparent as we proceed through the rest of this chapter.

The Frequency-Hopping Sequences

For devices to communicate with each other, they must transmit and receive on the same frequency at the same time. The *frequency-selection module* (FSM) contains the procedure for selecting the next frequency to be used under various operating conditions. The algorithmic and implementation details of FSM are beyond the scope of this book. They are covered extensively in chapter 11 of the Baseband part of the specification.




Figure 6.7
The frequency-selection module (FSM).

In Figure 6.7, the notation $(x|y)$ means that there are two permissible alternatives, x or y , only one of which is entered into the module based upon the circumstances. The notation $V[m:n]$, $m \leq n$, denotes the use of bits m through n of the bit field V ; m is the least significant bit (LSB) of the two.

Depending upon the country of use, the FSM is set by the manufacturer to operate in a 23- or 79-channel frequency hop mode as explained in the radio section of this chapter.

Given the country frequency-hop mode, the *clock input* determines which frequency from the current frequency-hopping sequence is to be used, and when. In other words, the clock input determines the *phase* of the frequency-hopping sequence. The actual frequency-hopping sequence is determined through the *address input*. In each of the three active operational states shown in Figure 6.4, a different combination of clock and address inputs is supplied to the FSM, as explained immediately below.

Normal Piconet Operation

During normal piconet operation, the *channel-hopping sequence* is used. For the channel-hopping sequence, the address input to the FSM for all devices in the piconet consists of the 28 least significant bits, $BD_ADDR[0:27]$, of the master's Bluetooth device address. The channel-hopping sequence has a very long period, but the (23 or 79) hop frequencies are distributed equally over short periods of time to satisfy the regulatory requirements for the frequency-hopping sequence described earlier in this chapter.

During normal piconet operation, a new frequency from the channel-hopping sequence is selected every 625 μsec . This interval is the period for bit c_1 of the Bluetooth clock shown in Figure 6.6; in this case, the LSB of the clock, c_0 , is not used. The time, 625 μsec , between two successive ticks of bit c_1 of the clock is referred to as a *slot*. For devices in the connected state, the slot boundaries coincide with the ticks of bit c_1 of the master's clock. Furthermore, all the devices in the piconet utilize the current value of the master's clock to drive their FSMs.

During the residence time at a frequency, a device may transmit a single packetized piece of information, referred to as a *baseband packet data unit* (BB_PDU)⁵. BB_PDUs are strictly constrained within the resi-

5. The specification refers to the baseband PDUs as such or simply as packets. The BB_PDU nomenclature is introduced here for consistency with use of the terms PDU and packet elsewhere. Nevertheless, the term packet is used whenever appropriate for ease of reading.

dence time at a frequency. However, the residence time at a frequency may occupy multiple slots, thus permitting multi-slot BB_PDU transmissions to occur as discussed below. Following a multi-slot transmission, the next frequency selected is the one that would have been used if single-slot transmissions had occurred instead. That is, the frequencies from the channel-hopping sequence are selected on a slot basis, although the frequency hops themselves occur on a BB_PDU basis.

Page Operation

One device "invites" another device to join its piconet through the use of a page. The device issuing the page is called a paging device, and the device listening (or scanning) for pages is called the paged device. A paging device selects a new frequency at which to transmit a page every 312.5 μsec , which is the period of bit c_0 of a Bluetooth clock. During page scans, when a device listens for transmitted pages, a new listening frequency is selected every 1.28 seconds, which is the period of bit c_{12} of the clock. Note that a paging device changes frequencies at a much higher rate than a paged device. While the paged device uses its own clock to drive its FSM, the paging device uses its best estimate of the clock of the paged device to drive its own FSM. The paging device estimates the clock value of the paged device based upon the most recent communication between the devices. In the worst case, the paging device could use its own clock.

During the page operation, the *page-hopping sequence* is used. To generate this sequence, the paging and paged devices use the 28 least significant bits of the address of the paged device—that is, the LAP and part of the UAP of the address shown in Figure 6.5—as the address input to their respective FSMs. For each device, its page-hopping sequence is a well-defined, periodic sequence composed of 32 (resp.⁶ 16) frequencies uniformly distributed over the 79 (resp. 23) frequency channels permissible in the 2.4 GHz band in various countries. The period of the page-hopping sequence is 32 (resp. 16) hops.

Inquiry Operation

Through inquiries a device "searches" for other devices in its vicinity. Just as in page operation, an inquiring device selects a new frequency at which to transmit an inquiry every 312.5 μsec . Inquired devices, executing inquiry scans, select a new listening frequency every 1.28 seconds.

6. The abbreviation "resp." stands for "respectively."