Lesson 137: Wireless LANs (802.11b)

by Jonathan Angel

We've covered wireless LANs before, as regular readers will have noticed (see "Wires Not Included," June 1999 and "Lesson 112: 802.11 and Spread Spectrum," December 1997). Recent approval of the IEEE's 802.11b standard for 11Mbit/sec wireless networking, however, has been followed by a wave of new product announcements from Apple, Compaq Computer, Lucent Technologies, and others.

This tutorial will discuss what's new about 802.11b and provide some implementation examples, then explain how it compares with HomeRF (proposed for home-based wireless networks, at www. homerf.org) and Bluetooth (for Personal Area Networks, or PANs, at www.bluetooth. com). For information on wireless WAN technologies and standards, as opposed to the wireless LAN architectures this tutorial will discuss, see "Wide World of Wireless."

Resources

For information about the 802.11 standard and the P802.15 study group, see http://grouper.ieee.org/groups/802/11/ and http://grouper.ieee.org/groups/802/15/, respectively.

The Wireless LAN Alliance is at www.wlana.com**,** while the Wireless Ethernet Compatibility Alliance (WECA) is at www.wirelessethernet.org. Other members of the nonexclusive 2.4GHz club are HomeRF (www.homerf.org) and Bluetooth **(**www.bluetooth.com**).** The HiperLAN/2 global forum may be reached at www.hiperlan2.com**.**

NOT ALL TRIPLETS ARE CREATED EQUAL

Every article on 802.11 (including our earlier two) inevitably points out that it is actually three standards in one. To be more specific, 802.11, ratified in 1997, gave birth to a single MAC standard for the lower portion of the Data-Link layer, plus three possible Physical (PHY) layers.

There is nothing new and noteworthy about the MAC layer in 802.11b, so I gladly refer you to the December 1997 tutorial for in-depth information. Briefly, however, wireless networking uses Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA). This modifies the Ethernet 802.3 standard so it can work via radio with a "virtual carrier sense" feature, based on brief Request to Send (RTS) and Clear to Send (CTS) packets.

The three PHY layers defined by the 802.11 standard are, confusingly, not interoperable. They are: infrared (never implemented by anybody); Frequency Hopping Spread Spectrum (FHSS) radio; and Direct Sequence Spread Spectrum (DSSS) radio.

FHSS was employed in most early 802.11 networks. Originally conceived during World War II by actress Hedy Lamarr and composer George Antheil (would I kid you about this?), it employs a narrowband carrier, changing its frequency in a pattern known only to both the sender and receiver. As intended, this makes information difficult to intercept.

While FHSS lives on in some products, as I will explain later, it is not part of 802.11b. The only triplet anointed as part of the new standard is DSSS. This type of signaling uses a broadband carrier, generating a redundant bit pattern (called a "chip") for every bit of data to be transmitted. While seemingly wasteful of bandwidth, DSSS copes well with weak signals: Data can often be extracted from a background of interference and noise without having to be retransmitted, making actual throughput superior.

Proponents of DSSS point to its superior range, plus its ability to reject multipath and other forms of interference. In fact, DSSS can reject noise from a microwave oven, for example, with relative ease, though it would still be swamped if deployed in the vicinity of a hospital's MRI scanner.

In any case, the 802.11b version of DSSS transmits data at a nominal 11Mbits/sec (actual rates vary according to distance from another transmitter/receiver). It is downwardly compatible with 1Mbit/sec and 2Mbit/sec wireless networking products, provided they also use DSSS and are 802.11-compatible.

With few exceptions, 802.11b is a worldwide standard. It uses the 2.4GHz to 2.48GHz Instrumentation, Scientific and Medical (ISM) frequency band, dividing this into as many as 14 different channels. In the United States, 11 channels are available for use.

Vendors must tailor their hardware access points (discussed later) to use legal channels in each country they ship to. Wireless NICs, however, can often adapt themselves automatically to whatever channels are being employed locally. Therefore, it is possible to travel with an 802.11b client and make connections in any country.

DEPLOYING A SYSTEM

It is difficult to plan a wireless network just by looking, or even by measuring distances. The antennas typically can, at the power levels permitted, transmit and receive for distances of about half a mile. This figure, however, only applies to outdoor, line-of-sight transmission.

Indoors, it is difficult to predict how a building's contour will affect propagation of radio waves. According to Harris Semiconductor, whose chipset is used in many 802.11b devices, range in an open plan "cube farm" may be from 200 feet to 500 feet. In a closed-wall office environment, it may be as low as 100 feet.

The metal found in an office building's floor can cut a signal by as much as 30 decibels (dB). Therefore, every floor in such a building will require one or more transmitters.

The simplest type of wireless LAN is a peer-to-peer setup that might be used in a conference room or at a trade show. Here, all stations are kept within a circle with a radius of approximately 300 feet, and direct communication between stations is possible.

To create this type of network, an administrator would install wireless NICs, setting their drivers to the ad hoc mode of operation, then selecting a radio channel for the workgroup. (In the United States, there is enough spectrum for three channels to coexist in one location, but channels must be 25MHz apart to avoid interference.)

In 802.11 lingo, this workgroup would be known as a Basic Service Set (BSS). A mechanism known as the Distributed Coordination Function (DCF), basically the "virtual carrier sense" function described earlier, provides best-effort delivery of data within a single, peer-to-peer BSS.

A more typical wireless network, however, is an "infrastructure" network—one that operates as an adjunct to a preexisting wired network. Here, Access Points (APs) are employed to act as a bridge (and usually a router), moving traffic between the wireless and wired networks (see "Gaining Access").

A hardware AP is a self-contained unit, typically featuring one or more Ethernet ports, plus either a built-in radio or a PC Card slot. (For the sake of versatility and easy upgrading, much 802.11b equipment employs PC Card-based transceivers, whether these will be installed in a portable computer or in a stationary piece of equipment.) A software AP is a functional, more affordable equivalent, using an existing computer that has been equipped with both wired and wireless NICs to perform bridging and routing.

As well as providing a gateway between network types, an AP has several other functions. For a start, the AP can provide Point Coordination Function (PCF), an optional connection-oriented mode. By broadcasting a beacon signal, the AP can temporarily silence ordinary terminals in order to provide point-to-point transmission of time-sensitive data, such as voice.

The primary functions of an AP, however, are authentication and association. The AP performs authentication to determine if a given wireless device is permitted to join the network, and can be based on MAC address, password, or some other parameter. Association is a handshaking relationship between the wireless device and the AP. It is designed to ensure that the client connects to only one AP at any given time.

ROAM SWEET ROAM

An Extended Service Set (ESS) is a logical collection of more than one BSS. Via an ESS, multiple APs can work together so that computers can roam from one to another while still staying in the same network.

To create this type of network, an administrator would install APs and wireless NICs, setting drivers to Infrastructure mode and making sure that all components are set to use the same ESS ID number (ESSID). To avoid interference, each AP should be set to a different channel.

Each 802.11 device associates with one AP initially, but a wireless network would be of limited use if stations were unable to roam. Fortunately, clients can switch from AP to AP in a way that is transparent to the user.

Logically, there are several ways roaming can take place, depending on the way APs have been set up. The simplest case is when different APs have the same ESSID and are on the same subnet of the same LAN. Slightly more complexity results when different APs have the same ESSID but live on different subnets. Here, DHCP re-registration is required, unless a Mobile IP solution is being used (see "Mobile IP Hits the Street," November 1999). Multiple APs can also form different logical networks on a single LAN via the use of different ESSIDs.

Given the nature of radio-based communications, eavesdropping is always a possibility. Therefore, the 802.11 standard includes a shared-key encryption mechanism known as Wired Equivalent Privacy (WEP). When a client tries to connect to an AP, the AP sends a challenge value to the station. Upon receiving this, the client uses the shared key to encrypt the challenge and send it to the AP for verification.

While useful, WEP only allows for 40-bit encryption; some vendors of 802.11b equipment offer optional 128-bit encryption or plan to make it available as a firmware upgrade. A few also sell wireless NICs that have been

manufactured not only with a unique MAC address but also with a unique public/private key pair. Administrators can require that all allowable hardware address/public key combinations be entered into APs in advance. Alternately, they can simply configure APs to keep track of the combinations they encounter and subsequently reject any mismatches. This way, an attacker can be prevented from breaking into a network via MAC address masquerading.

CUTTING ACCESS COST

While valuable for the infrastructure network, features such as multiple ESSIDs, roaming, and 128-bit encryption increase the cost of hardware APs to around $1,000 each. Vendors such as Apple and Lucent, however, offer simplified APs without roaming for less than half that price. Not to be confused with prestandard wireless networking solutions, these "budget" APs have the advantage of being fully compatible with other 802.11b equipment.

This means users could purchase a single wireless NIC, then use it both in a corporate setting and in a home office. In order to popularize the advantages of 802.11b compatibility, the Wireless Ethernet Compatibility Alliance (WECA, at www.wirelessethernet.org) recently announced a labeling program known as "Wi-Fi."

Of course, 802.11b is not the only entrant into the 2.4GHz wireless networking melee. A rival of sorts is the HomeRF Shared Wireless Access Protocol (SWAP) system, which has been designed for consumers. It uses FHSS transmission and eliminates the more complex parts of 802.11 (such as PCF and RTS/CTS). An advantage here is that a single connection point can support both voice services via Time Division Multiple Access (TDMA) and data services via CSMA/CA.

Another contender, Bluetooth, uses the 2.4GHz band for localized connection between different devices on a PAN. These might include a PC and a handheld device, a phone and a headset, or a notebook computer and a printer. While there are grounds for concern about interference between 802.11b, HomeRF, Bluetooth, and the many other devices using the same spectrum (such as baby monitors and garage door openers), some observers seem to believe all these can coexist. A coexistence study group exists within the IEEE (P802.15) and presentations about this topic have been fairly encouraging.

Eventually, wireless LANs will migrate into the relatively wide-open spaces offered in the 5GHz band, where they will be able to exchange data at up to 54Mbits/sec. Just as portable computers have always lagged behind their desktop cousins in terms of speed and affordability, wireless networks will always lag behind what copper and fiber can offer. Once again, however, there's no question of which best suits the needs of a mobile worker.

Jonathan Angel, senior editor, can be reached at jangel@mfi.com.