♦ WaveLAN®-II: A High-Performance Wireless LAN for the Unlicensed Band

Ad Kamerman and Leo Monteban

In July 1997 the Institute of Electrical and Electronics Engineers (IEEE) completed standard 802.11 for wireless local area networks (LANs). WaveLAN®-II, to be released early in 1998, offers compatibility with the IEEE 802.11 standard for operation in the 2.4-GHz band. It is the successor to WaveLAN-I, which has been in the market since 1991. As a next-generation wireless LAN product, WaveLAN-II has many enhancements to improve performance in various areas. An IEEE 802.11 direct sequence spread spectrum (DSSS) product, WaveLAN-II supports the basic bit rates of 1 and 2 Mb/s, but it can also provide enhanced bit rates as high as 10 Mb/s. This paper discusses various aspects of the system design of WaveLAN-II and characteristics of its antenna, radio-frequency (RF) front-end, digital signal processor (DSP) transceiver chip, and medium access controller (MAC) chip.

Introduction

The first wireless local area network (LAN) products appeared in the market around 1990, although the concept of wireless LANs has existed since the late 1970s.[1] The release of the industrial, scientific, and medical (ISM) bands made unlicensed spectrum available and prompted significant interest in the design of these wireless LANs. The next generation of these products is being implemented on Personal Computer Memory Card International Association (PCMCIA) cards used in laptop computers and portable devices. The major technical issues surrounding wireless LAN systems are size, power consumption, bit rate, aggregate throughput, range of coverage, and interference robustness.

After the NCR WaveLAN® was introduced in 1991 as an Industry Standard Architecture (ISA) card for the 915-MHz band, several new versions appeared for other computer buses, for the 2.4-GHz band, for various network operating systems, and as original equipment manufacturer (OEM) products for companies such as DEC, Toshiba, and Solectek. Throughout the WaveLAN product improvement life cycle, a number of key components—notably the NCR digital signal processing (DSP) application-specific integrated circuit (ASIC) and the Intel Ethernet controller—stayed the same as for the original WaveLAN product. The improvements covered such areas as smaller card design, lower power consumption, and increased software support.

At the beginning of 1998, Lucent Technologies, which played a leading role in driving the Institute of Electrical and Electronics Engineers (IEEE) 802.11 for wireless LANs, will introduce WaveLAN-II. This product has taken a big step forward in performance for the small-size, low-power PCMCIA card-based design.

This paper is organized as follows: In “The IEEE 802.11 Standard,” we define the standard and examine the differences between direct sequence spread spectrum (DSSS) and frequency hopping spread spectrum (FHSS) within the context of this standard. Next, in “WaveLAN-II’s Enhanced Bit Rates,” we consider these higher bit rates. We look at where the different functions are implemented in “Distribution of Functions.” Further, the implementation complexity needed for the radio-frequency (RF) and DSP design is compared to the design in the code division multiple

Panel 1. Abbreviations, Acronyms, and Terms

ACK—acknowledgment
ARF—automatic rate fallback
ASIC—application-specific integrated circuit
BCA—basic coverage area
BCPM—Barker code position modulation
BPSK—binary phase shift keying
CDT—carrier detect threshold (level)
CDMA—code division multiple access
CSIR—co-channel signal to interference ratio
CSMA/CA—carrier sense multiple access with collision avoidance
CSMA/CD—carrier sense multiple access with collision detection
CQ—communications quality (WaveLAN-II proprietary quality scale)
CTS—clear to send
DSP—digital signal processing/processor
DSSS—direct sequence spread spectrum
DT—defer threshold (level)
FHSS—frequency hopping spread spectrum
FSK—frequency shift keying
GFSK—Gaussian frequency shift keying
GSM—Global System for Mobile Communications (European standard for digital cellular telephone)
IEEE—Institute of Electrical and Electronics Engineers
IEEE 802.3—IEEE standard for CSMA/CD-based wired LANs (Ethernet)
IEEE 802.11—IEEE standard for CSMA/CA-based wireless LANs
IF—intermediate frequency
IFS—interframe spacing
IS-95—Interim Standard 95 (U.S. standard for CDMA-based digital cellular telephone)
ISA—Industry Standard Architecture (specifies ISA bus in computers)
ISM—industrial, scientific, medical (frequency band for these applications)
LAN—local area network
MAC—medium access control (layer)
MIPS—million instructions per second
OEM—original equipment manufacturer
PCMCIA—Personal Computer Memory Card International Association (specifies PC card interface)
PHY—physical (layer)
PN—pseudo noise
PPM—pulse position modulation
PSK—phase shift keying
QAM—quadrature amplitude modulation
QPSK—quadrature phase shift keying
RAM—random access memory
RF—radio frequency
RTS—request to send
SCA—shared coverage area
SNR—signal-to-noise ratio
WMAC—wireless medium access control/controller

access (CDMA)-based cellular telephone systems, described in “IEEE 802.11 DSSS vs. IS-95 CDMA.” We also discuss multi-channel roaming and automatic rate selection. “Power Management” describes the power consumption and support for the IEEE 802.11 standard, and “Scalable System” examines configurable carrier detection, configurable defer behavior, roaming thresholds, and infrastructure density. “Conclusion” summarizes the paper and projects how the WaveLAN-II design will respond to future upgrades in data rates.

The IEEE 802.11 Standard

IEEE 802.11 is a standard for wireless systems that will operate in the 2.4- to 2.5-GHz ISM band. Available worldwide, this ISM band allows unlicensed operation of spread spectrum systems. For both the United States and Europe, the 2,400- to 2,483.5-MHz band has already been allocated, while for some other countries, such as Japan, another part has been assigned, as shown in **Table I**. The 802.11 standard focuses on the medium access control (MAC) and physical layer (PHY) protocols for ad-hoc networks and networks based on access points.

In *access point-based networks*, the stations within a group or cell can only communicate directly with the access point. The access point forwards messages to the destination station within the same cell or through the wired distribution system to another access point, from which messages arrive at the destination station. In *ad-hoc networks*, the stations operate on a peer-to-peer level with no access point or (wired) distribution system.

The 802.11 standard supports DSSS with differen-

Table I. Frequency bands and power levels for wireless LANs.

Countries	Frequency range	Maximum RF power level	Rules for DSSS and FHSS
U.S., Canada,* and Latin America (FCC Part 15,247)	902–928 MHz 2,400–2,483.5 MHz 5,750–5,850 MHz	1W (at ERP and maximum 6 dBi antenna gain)	DSSS: Receiver processing gain > 10 dB FHSS: 75 hops or more
Europe† (ETS 300 328)	2,400–2,483.5 MHz	100 mW (at EIRP)	DSSS: Power spectrum density maximum 10 mW/MHz FHSS: 20 hops or more
Japan (MPT Ordinances 78 and 79)	2,471–2,497 MHz	Not specified	DSSS/FHSS: Power spectrum density maximum 10 mW/MHz
Australia	2,400–2,450 MHz	500 mW	

* In Canada, not the 5,750–5,850-MHz band

† In France/Spain, only the 2,445–2,483.5/2,475-MHz band

EIRP – Equivalent isotropically radiated power

ERP – Effectively radiated power

ETS – European Telecommunication Standard

MPT – Ministry of Posts and Telecommunications (in Japan)

tial encoded binary phase shift keying (BPSK) and quadrature phase shift keying (QPSK), FHSS with Gaussian frequency shift keying (GFSK), and infrared with pulse position modulation (PPM). The basic medium access behavior allows compatible PHYs to operate together by using the carrier sense multiple access with collision avoidance (CSMA/CA) protocol and a random back-off time following a busy medium condition. In addition, all directed traffic uses immediate positive acknowledgment (ACK frame), in which the sender schedules a retransmission if it does not receive an ACK.

While CSMA/CA and Ethernet's carrier sense multiple access with collision detection (CSMA/CD) have similarities, their one fundamental difference is the way they handle collisions. In wire-based networks, such as Ethernet, it is not technically complicated to detect if transmissions from two stations are colliding. Detecting collisions in wireless systems that use only one channel is impractical, however, because of the large dynamic range of receive levels. Therefore, 802.11 has chosen CSMA/CA, which uses a collision avoidance scheme. (**Panel 2** compares IEEE 802.11 DSSS with IS-95 CDMA.)

The 802.11 CSMA/CA protocol is designed to reduce the collision probability between multiple stations accessing the medium, at the point in time where collisions would most likely occur. The highest probability of a collision would occur just after the medium becomes free, following a busy medium, because multiple stations would have been waiting for the medium to become available again. Therefore, a random back-off arrangement is used to resolve medium contention conflicts, as shown in **Figure 1**. A very short-duration carrier detect turnaround time is fundamental to this random wait characteristic. The 802.11 standard DSSS uses a slotted random wait behavior based on 20-μs time slots, which cover the carrier detect turnaround time.

In addition, the 802.11 MAC defines an option for reserving a medium using request-to-send/clear-to-send (RTS/CTS) polling interaction and point coordination (for time-bounded services).[2,3] After a busy-medium period, all wireless LAN devices have to wait during an interframe spacing (IFS) period. After waiting a random number of time slot intervals, these devices can attempt to transmit, provided no other transmission was detected in the interim.

Regulations under which radio spectrum may be used and frequencies in which wireless LANs can be deployed have made the 2.4-GHz ISM band available worldwide. Three similar approval regimes exist: the U.S., Europe, and Japan. As Table I shows, there are no worldwide regulations for spectrum occupation or power levels. The IEEE 802.11 standard has focused its standardization efforts on the

2.4-GHz band because it is available worldwide, despite its "regional" variations.

IEEE 802.11 DSSS

DSSS systems spread the signal energy across a relatively wide band by increasing the occupied bandwidth. A DSSS transmitter converts a bit stream into a symbol stream, in which each symbol represents a number of bits, depending on the phase shift keying (PSK) modulation technique. The symbol information is converted into a complex-valued signal that is fed to the spreader. The spreader multiplies its input signal with a pseudo noise (PN) sequence, called a *chip sequence*. This multiplication creates a signal with a wider bandwidth. The in-phase and quadrature components of the spreader output signal are fed to a quadrature modulator. The transmitter front-end provides filtering, conversion to a higher RF, and power amplification.

The 802.11 DSSS, based on the 11-chip Barker sequence +1, -1, +1, +1, -1, +1, +1, +1, -1, -1, -1, has already been used for WaveLAN-I.[4] This 11-chip sequence is used as the PN code sequence, and the symbol duration corresponds to the time of 11 chip intervals. The 11-chip spreading makes the occupied bandwidth larger and increases the effective bandwidth from 1 MHz to 11 MHz. The 802.11 standard specifies two bit rates—1 Mb/s with BPSK and 2 Mb/s with QPSK, with a spectrum that looks the same for both bit rates. The channel center frequencies under 802.11 are 2,412 MHz, 2,417 MHz, 2,422 MHz, ... 2,462 MHz (for the U.S. and Europe). In Europe it also supports 2,467 and 2,472 MHz, but 802.11 supports only the channel center frequency 2,484 MHz in Japan.

IEEE 802.11 FHSS

FHSS systems hop from narrow band to narrow band within a wide bandwidth. FHSS wireless LAN stations send one or more data packets at one carrier frequency, then hop to another carrier frequency to send one or more packets, and continue this hop–transmit sequence, called *slow frequency hopping*. The time these FHSS radios dwell on each frequency is fixed. The hopping pattern appears random, but it is actually a periodic sequence tracked by both the sender and receiver. A FHSS transmitter converts the

Panel 2. IEEE 802.11 DSSS vs. IS-95 CDMA

Interim Standard (IS)-95 for mobile cellular telephone systems, the Telecommunication Industry Association's standard, is based on CDMA. This standardized CDMA technique combines DSSS with a high spreading factor and medium access. During a call setup at an IS-95 cellular system, a full-duplex connection is established, with 28.8 kb/s in the up link and 19.2 kb/s in the down link. Convolutional coding is applied on top of these bit rates for the embedded speech information (1.2 kb/s to 9.6 kb/s). The 28.8-kb/s and 19.2-kb/s match modulation rates over spreading factors of $(64/6) \times 4$ and 64, respectively, which, for both bit rates, result in a chip rate of 1.288 megachips per second (Mc/s).

Power control in the transmitters of mobile stations reduces the interference at the base station receiver side between the transmitted signals from in-cell mobile stations and the interference at other base stations. Likewise, power control in the base station transmitters reduces interference at mobile station receivers. Such reduction of interference allows more in-band activity and increases the capacity for a given bandwidth.

The 802.11 standard DSSS prescribes a single fixed 11-chip spreading code that is used by all stations. The medium access assignment is not made by a unique code, but rather by a listen-before-talk carrier sensing and at a deferred transmission. This transmission can only start after waiting a random number of time slots of 20 µs. Because the medium access is based on sensing carrier activity above a receive level threshold that matches with the target defer range, a transmit power control would conflict with the harmonious defer behavior within the target cell area. The stations and the access point share the single-channel medium and, after getting access to the medium, they can use the full bandwidth of 2 Mb/s for some period. The half-duplex operation at 2 Mb/s is based on a symbol rate of 1 Mbaud and a chip rate of 11 Mc/s. The medium/channel access, medium/channel occupation time, bit rate, bit error rate requirements, real-time aspects, and capacity definitions of packet transmission and circuit-switched telephone services are totally different in nature.

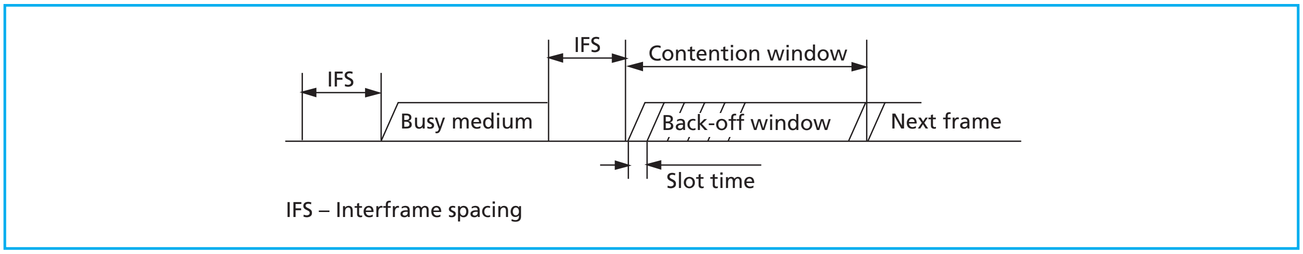

Figure 1.
Basic CSMA/CA behavior.

bit stream into a symbol stream in which each symbol represents one or more bits.

Two frequencies are applied for binary frequency shift keying (FSK) modulation, and four frequencies are applied for quaternary FSK modulation. Frequency hopping is applied to the FSK-modulated signal. The transmitter front-end supplies conversion to a higher RF and power amplification. The 802.11 FHSS uses a GFSK modulation technique with a low modulation index (Gaussian frequency shaping, BT [bandwidth – symbol interval] product = 0.5, modulation index $h = 0.34$ at 1 Mb/s, 0.15 at 2 Mb/s), which gives a relatively narrow spectrum and allows a higher bit rate in the 1-MHz narrow hop channels. However, these FSK conditions increase the sensitivity to noise and other impairments. The 802.11 standard defines hops over channel center frequencies according to a periodic sequence (that looks like a random pattern), within a set of 79 frequencies: 2,402 to 2,480 MHz in the U.S. and Europe. For Japan, however, 802.11 supports a set of only 23 frequencies (2,473 to 2,495 MHz).

DSSS vs. FHSS

We have compared various performance aspects[5] of the 802.11 standard DSSS and FHSS systems. DSSS has a more robust modulation and a larger coverage range than FHSS, even when FHSS uses twice the transmitter power output level. FHSS gives a large number of hop frequencies, but the adjacent channel interference behavior limits the number of independently operating collocated systems. Hop time and a smaller packet size introduce more transmission time overhead into FHSS, which affects the maximum throughput. Although FHSS is less robust, it gives a more graceful degradation in throughput and connectivity. Under poor channel and interference conditions, FHSS will continue to work over a few hop channels a little longer than over the other hop channels. DSSS, however, still gives reliable links for a distance at which very few FHSS hop channels still work.

For collocated networks (access points), DSSS gives a higher potential throughput with fewer access points than FHSS, which has more access points. The smaller number of access points used by DSSS lowers the infrastructure cost.

WaveLAN-II's Enhanced Bit Rates

In addition to the 802.11 bit rates of 1 and 2 Mb/s, a release of WaveLAN-II scheduled for July 1998 will provide bit rates as high as 5, 8, and 10 Mb/s. This allows interoperability with 802.11 DSSS devices at the lower (fallback) bit rate and coexistence with those devices when the proprietary bit rates are higher. The PHY training preamble/header of the transmission, which takes 200 µs, looks the same for each bit rate; it is always based on the modulation for a bit rate of 1 Mb/s. The data portion—which can take a few milliseconds, depending on the frame size (up to 2,300 bytes)—is transmitted at a bit rate of 1, 2, 5, or 8 Mb/s. Because the preamble/header is recognized by any 802.11 DSSS device, such a device then acquires the correct defer behavior. After that device transfers initial information at 1 or 2 Mb/s, the interoperating devices assess each other's capability, which may be 5, 8, or 10 Mb/s. A station then knows if it is useful to switch to a higher bit rate.

The 802.11 standard DSSS uses an 11-chip Barker code sequence to carry 1 or 2 bits of data per pulse use, in the BPSK and QPSK formats, respectively. After the received signal is correlated using a filter matched to the same 11-chip Barker code sequence, the resulting signal gives a main lobe,

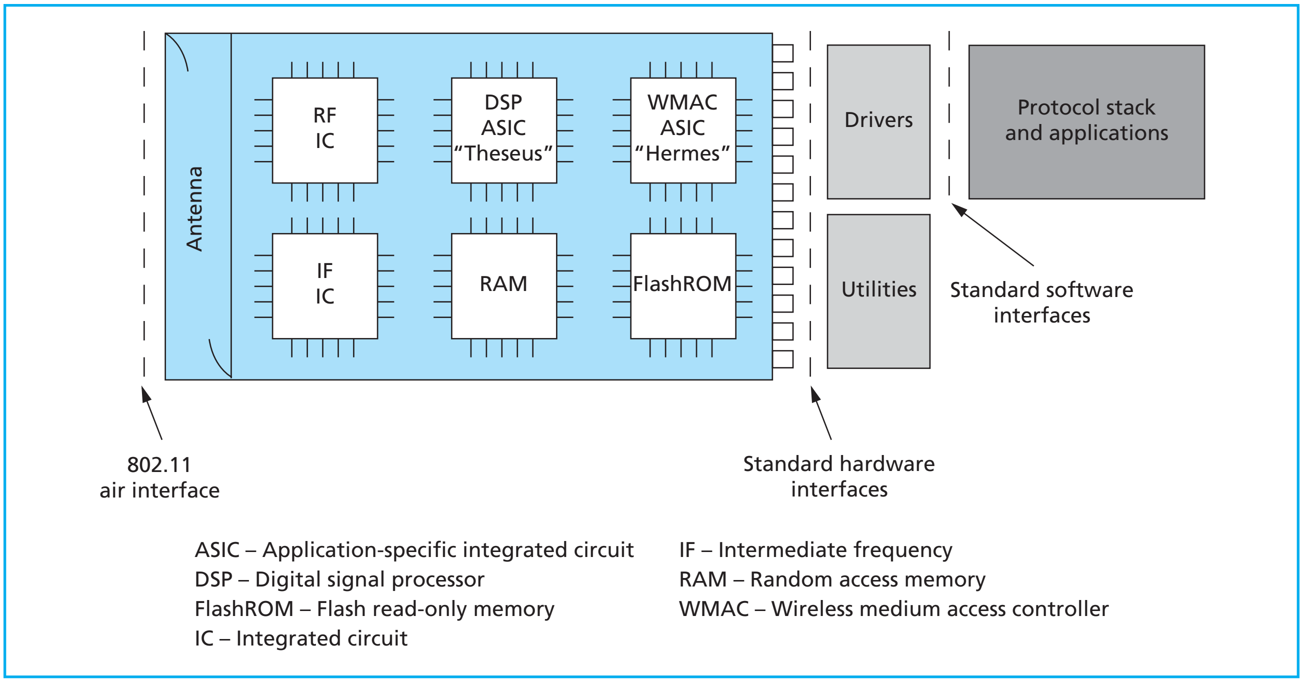

Figure 2.
Major WaveLAN-II functions.

which is 1 chip wide. The 10 additional chip positions within the symbol period give side lobe values 11 times smaller than the main lobe. The position of the 11-chip Barker code sequence within each symbol interval could be modulated in addition to the modulating polarity of the Barker code sequences in the in-phase and quadrature channels.

Bell Labs Research and the Wireless Communications & Networking Division in The Netherlands have jointly investigated modulating the position of the code sequence; this effort has resulted in higher bit rates when no extreme channel degradation exists. If we allow eight positions for the main lobe within the 11-chip symbol interval, and these main lobe positions are chosen independently in the in-phase and quadrature channel, then the information rate is increased by 2 times 3 bits.[6] In this way we can achieve an increase from 2 bits per symbol (QPSK) to 8 bits per symbol (QPSK + disjoint Barker Code position modulation [BCPM], or independent BCPM in the I and Q channels).

In a similar way we can increase the information rate from 2 bits per symbol (QPSK) to 5 bits per symbol (QPSK + joint BCPM); then the BCPM bit rates are 5 Mb/s (QPSK + joint BCPM) and 8 Mb/s (QPSK + disjoint BCPM). The spectra of the transmitted signals at 1, 2, 5, and 8 Mb/s are all the same, because of the spreading of the identical 11-chip Barker code sequence. The differences between these four bit rates with respect to the required signal-to-noise ratio (SNR) for reliable receiver operation are limited to 6 dB, while the bit rate increases eight times. However, the constraints for the transceiver implementation with respect to linearity and compensation for channel degradation are more severe for operation at 5 and 8 Mb/s.

Likewise, as described above, we can combine multilevel modulation with BCPM. At quadrature amplitude modulation (QAM)-16 and disjoint BCPM, the bit rate will be 10 Mb/s.

Distribution of Functions

The WaveLAN-II network card and a set of software modules cooperate to offer an 802.11 LAN connection for a PC. At the highest software interfaces, the 802.11 LAN connection is equivalent to a traditional Ethernet (802.3) LAN. At the air interface, the 802.11 RF/baseband modulation and protocols are

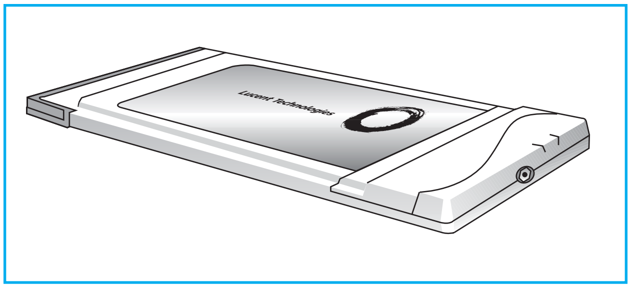

Figure 3.
WaveLAN-II PCMCIA card with internal antenna.

used. **Figure 2** gives a schematic overview of the major functional elements of the hardware and software and describes how the various functions are distributed over these elements.

The WaveLAN-II card, shown in **Figure 3**, is used in laptop computers, its right-hand side extending from the laptop cabinet. The transceiver front-end is mounted in a plastic cover. The slightly thicker part of the card contains the internal antenna, which has excellent properties. The antenna automatically switches from internal to external when connected to an external antenna cable.

Network cards will be available for a number of standard hardware interfaces like PCMCIA and ISA. Figure 2 shows the hardware elements on the network card, including the:

- *Antenna*. The WaveLAN-II card incorporates two L-shaped inverted-F antennas. They are placed in mutually orthogonal fashion (cross-polarized) and spaced at a wavelength of 0.25 to 0.5 from each other, which provides antenna diversity. This antenna represents a "leaky" resonator at 2.45 GHz (+/-100 MHz). Shaping it in such a manner that none of its dimensions exceeds 1/8 of a wavelength makes this antenna "electrically small" and consequently very omnidirectional.
- *RF front-end*. The radio-frequency–intermediate frequency (RF–IF) up/down conversion is based on a dedicated chip design for the 2.4-GHz ISM band. This chip includes a low-noise amplifier. The IF frequency is 352 MHz.
- *IF transceiver*. The IF-baseband up/down conversion is made by a modulator/demodulator chip, which is a special development derived from a Global System for Mobile Communications (GSM) chip.
- *DSP ASIC "Theseus."* The DSP ASIC, which performs the 802.11 DSSS, was developed by Lucent Technologies. It is the heart of the scalable radio section. (See **Panel 3** for a comparison of RF and DSP components.)
- *WMAC ASIC "Hermes."* The "Hermes" is a powerful processor that runs the 802.11 wireless MAC (WMAC) protocol firmware. Beyond the basic frame transmit and receive, some of the firmware functions it carries out are:
  - Power management,
  - Automatic rate fallback (ARF), and
  - Multi-channel roaming and handover management.
- *FlashROM*. The FlashROM provides nonvolatile memory for storing permanent and configurable information to control the behavior of firmware and hardware operation. For example, it holds values used in roaming and scalable system design.
- *RAM*. The random access memory (RAM) volatile memory is used for message buffering and internal firmware administration.

The network card's software elements include:

- *Drivers*. A set of LAN drivers directly control the data flow through the network card. Different drivers offer a variety of high-level software interfaces, such as Microsoft's Network Driver Interface Specification, Novell's Open Data-Link Interface, and Packet Driver.
- *Utilities*. This set of programs, provided for network planning and card diagnosis, uses the network card to collect statistics and to monitor the RF medium.

Single- and Multi-channel Roaming

WaveLAN-II provides roaming within the coverage boundaries of a set of access points that are interconnected by a wired distribution system. The access points send beacon messages at regular intervals. Stations can track the conditions at which the beacons are received per individual access point. The running average of these receive conditions is determined by a communications quality (CQ) indicator, as shown in **Figure 4**. The different zones within the full-range CQ scale refer to various states of activity at which a station tracks or tries to find an access point. When the CQ is poor, the station has to expend more effort to quickly find another access point that gives a better CQ.

The access points are interconnected by a wired distribution system, shown in **Figure 5**. They can use channel frequencies from a basic set of all supported channel frequencies within 802.11 DSSS. A station can search for an access point giving a better CQ by looking at all channel frequencies selected for the interconnected access points. The searching station can initiate an active mode by sending a so-called probe request message referencing the target set of interconnected access points. Each access point will respond to a probe request with a probe response message. This will serve as a "solicited" beacon.

When a station's CQ decreases with respect to its associated access point, the station starts searching more actively. After it has found a second access point that gives a sufficiently good CQ, the station moves into a handover state and reassociates with this second access point. The access points deploy an inter-access point protocol to inform each other of station handovers and to correct any intermediate MAC bridge filter tables.

Power Management

For battery-powered PC devices, the power consumption of a LAN card is a critical factor. WaveLAN-II supports power management for these environments. The 802.11 standard defines power management protocols that can be used by stations. Power management schemes result in a lower consumption of (battery) power compared to traditional operation, where a station is always monitoring the medium during idle periods. To achieve savings in power consumption, a LAN card in a station must have a special low-power state of operation called the DOZE state. In this state the LAN card will not monitor the medium and will be unable to receive a frame. This state differs from the OFF state. The card must be able to make a transition from the DOZE state to the fully operational receive (AWAKE) state in just 250 µs. A transition from the OFF to the AWAKE state will take much more time.

Power management allows a station to spend most of its idle time in the DOZE state, while it is still connected to the rest of the network to receive unsolicited messages. For the latter requirement, the other stations or the access point must temporarily buffer the messages that are destined for a station operating in a power management scheme. Such a station must "wake up" at regular intervals to check on whether any messages are buffered for it.

802.11 Power Management Support

Each frame that is transmitted by a station contains one so-called PM bit to indicate the station's

Panel 3. RF and DSP Components

The WaveLAN-II card has miniaturized RF circuitry for 2.4 GHz and contains ASICs for DSP and MAC. As a reference, Lucent's IS-95 design for the digital cellular telephone is based on RF circuitry for 850 MHz, the Lucent DSP 1629 chip, and an Intel microcontroller.

WaveLAN-II provides half-duplex operation for a 2-Mb/s bit rate with a 1-Mbaud symbol rate and a chip rate of 11 Mc/s with 22-MHz sampling (a factor of 2 oversampling). The receive mode requires much more DSP power than the transmit mode. The receiver processing has to be completed in a 1-µs symbol interval, including the correlator (22-MHz complex-valued signal samples), adaptive matched filter, phase error control, and decision unit.

The current WaveLAN-II DSP implementation requires more than 2,000 multiply–add operations per microsecond, which far exceeds the processing power—in million instructions per second (MIPS)—of present and next-generation general-purpose DSP chips. Therefore, the DSP ASIC "Theseus" is used for WaveLAN-II as a "horsepower" DSP. This "Theseus" chip integrates two analog-to-digital and two digital-to-analog converters, which both run at 22 MHz and have a 6-bit digital representation.

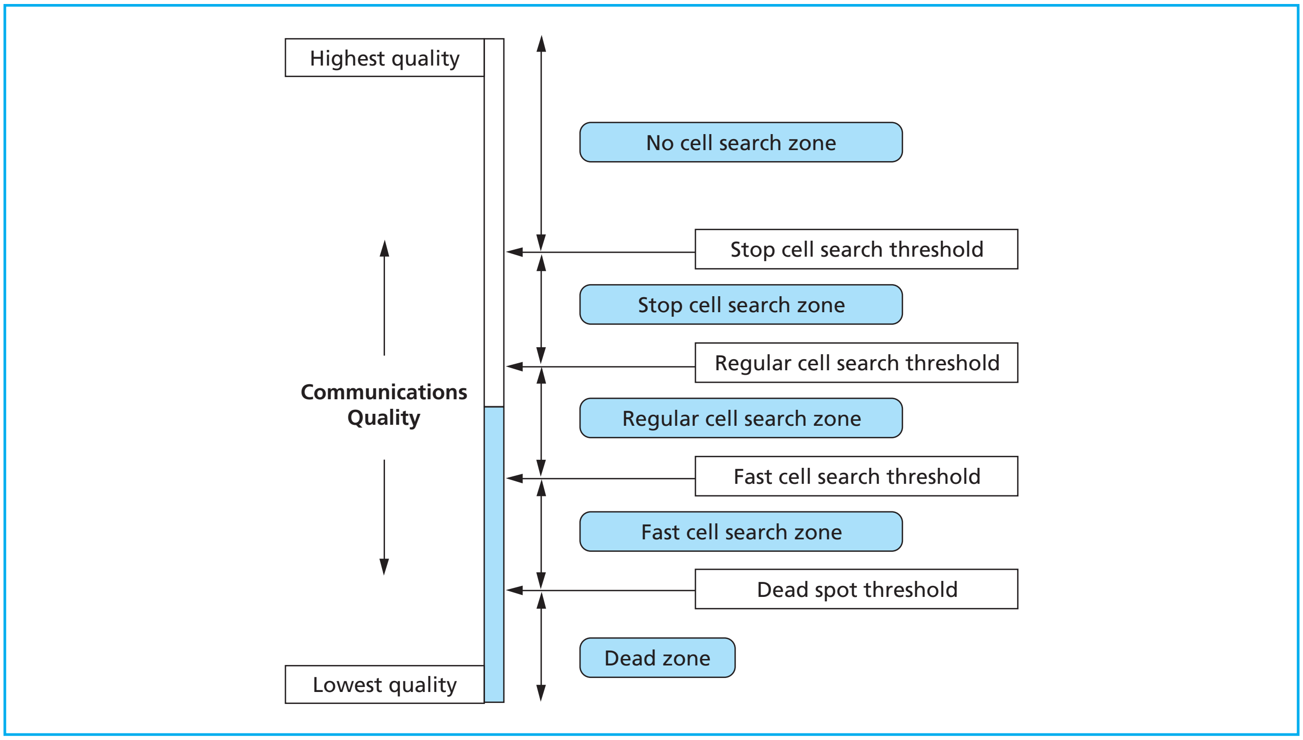

Figure 4.
CQ scale and cell search zones.

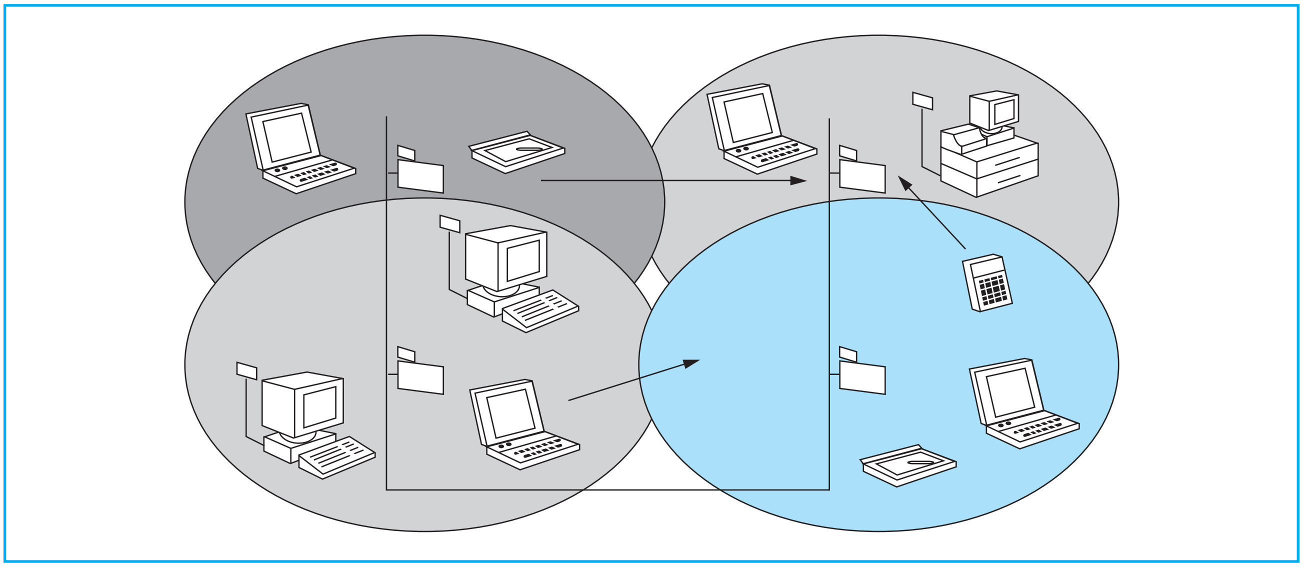

Figure 5.
Wired infrastructure between access points and multi-channel operation with three different "colored" channel frequencies.

mode of operation. *PM-bit* = *PS* indicates that the station is operating in the power saving mode, and *PM-bit* = *A* indicates that the station is in the active mode (continuously active).

In an access point-based network, the access points will learn from the PM-bit whether a station is in active or power saving mode and will buffer unicast messages as required by the station. Access points will send beacon frames on a regular basis (for example, every 100 ms). In each beacon frame, the access point

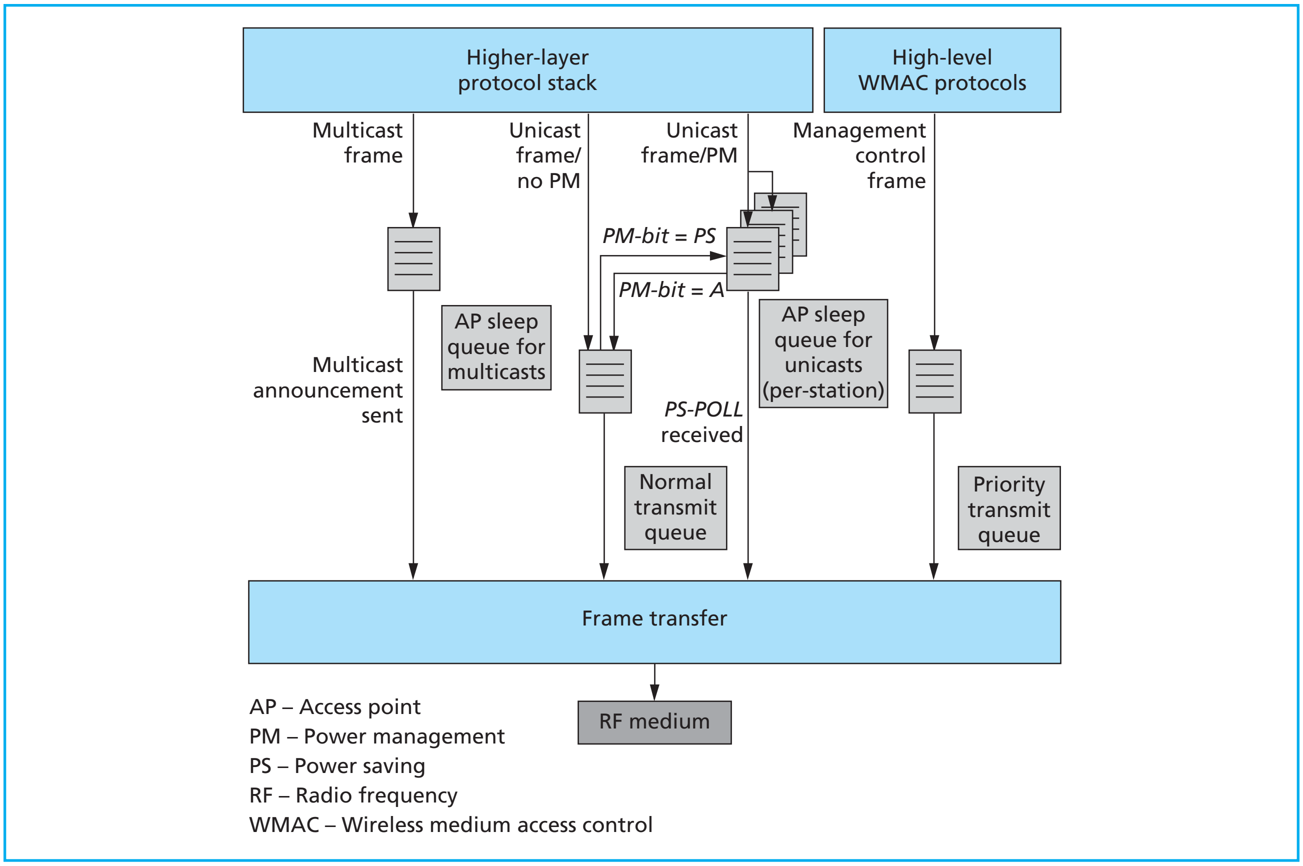

Figure 6.
The 802.11 power management scheme for an access point.

will announce the station whose messages it has buffered. The stations using a power management scheme will wake up with high accuracy just before a beacon transmission and determine, via the beacon, whether any messages are buffered. If no messages are buffered, the station will go into the DOZE state until the next beacon arrives. When one or more messages are buffered, the station will stay in the AWAKE state and will poll the access point for transmission of the buffered messages. At a regular interval, which is a multiple of the beacon transmission interval, another information field is included in the beacon to announce whether any multicast messages are buffered. These messages are transmitted directly following their announcement. Therefore, the stations that need to receive these buffered multicast messages must stay in the AWAKE state following the multicast announcing beacon.

Queue Structure in an Access Point

To support the power management schemes of 802.11, a WaveLAN-II access point will deploy the queuing structure shown in **Figure 6**.

The frame transfer function, shown at the bottom of Figure 6, will interface with the RF medium. This frame transfer block has two input queues, a normal transmit queue and a priority transmit queue. The priority queue schedules internal management and control protocol frames; the normal queue is used for data frames that come from the higher-layer protocol stack. There is a single multicast sleep queue and multiple unicast sleep queues, one for each station under control of this access point.

When a station polls for a buffered message, the buffered message at the head of the unicast sleep queue for that station is passed directly to frame transfer for highest priority transmission. When a station that was in the PS mode sends a message to the access

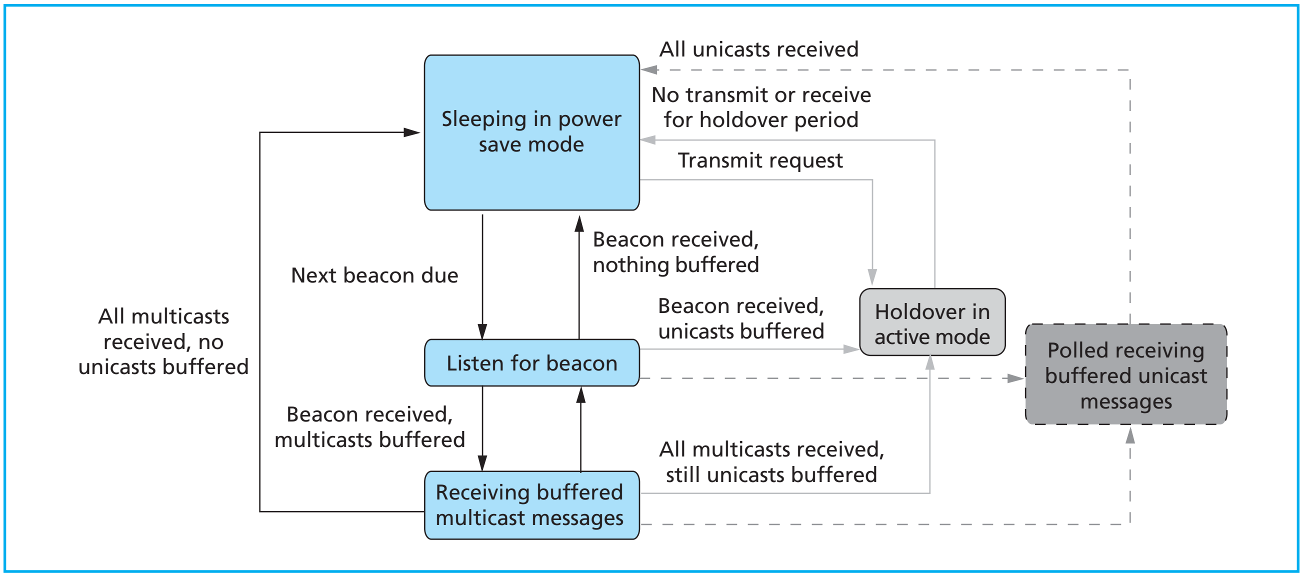

Figure 7.
State diagram for two different power management schemes.

point with the PM-bit set to "*A*", all buffered messages for that station are forwarded to the normal transmit queue. When a station that was in the active mode sends a message with the PM-bit set to "*PS*", all messages for that station in the normal transmit queue are moved to the unicast sleep queue.

Power Management Schemes for a Station in an Access Point-Based Network

The stations in an access point-based network can use the 802.11 power management provisions in a number of schemes. **Figure 7** shows two such schemes—the "standard" 802.11 scheme and the "enhanced" scheme, developed especially for WaveLAN-II. When not involved in actual reception, the receiver will switch between the "sleeping" state and the "listen for beacon" state. An accurate timer will wake up the receiver just in time for the next beacon frame; when nothing is buffered, the receiver immediately returns to the DOZE state. This power management provision, used in both schemes, yields an efficiency rate of more than 99% of idle time spent sleeping. The "standard" 802.11 scheme is shown in solid color on the left side and dark gray on the right side. The "enhanced" scheme is shown in solid color on the left side and light gray on the right side.

When the standard scheme analyzes the beacon and detects the presence of buffered messages, it directs the receiver to stay in the AWAKE state and either wait for the multicast messages or actively poll the access point for unicasts. After all buffered messages are received, the receiver reenters the "sleeping" state.

The main difference between the two schemes is that the enhanced scheme has a special "holdover" state, which is entered when either a transmission is requested or when a buffered unicast message is detected. In this "holdover" state, the station temporarily switches operation from the power save mode to active mode and transmits this information to the access point via the PM bit. The access point then forwards all buffered messages immediately and will stop sleep queue buffering as long as this state lasts. The state will last until there is a designated continuous period—about 0.5 to 3 seconds—of no transmit or receive activity for this station, called the *holdover period*.

The enhanced scheme was developed for the following reasons:

- In many LAN applications, the exchange of messages between a station and a server takes place in bursts. Multiple transmits and receives occur over a short interval, followed by a relatively long period of inactivity.
- Keeping the receiver in the active mode prevents

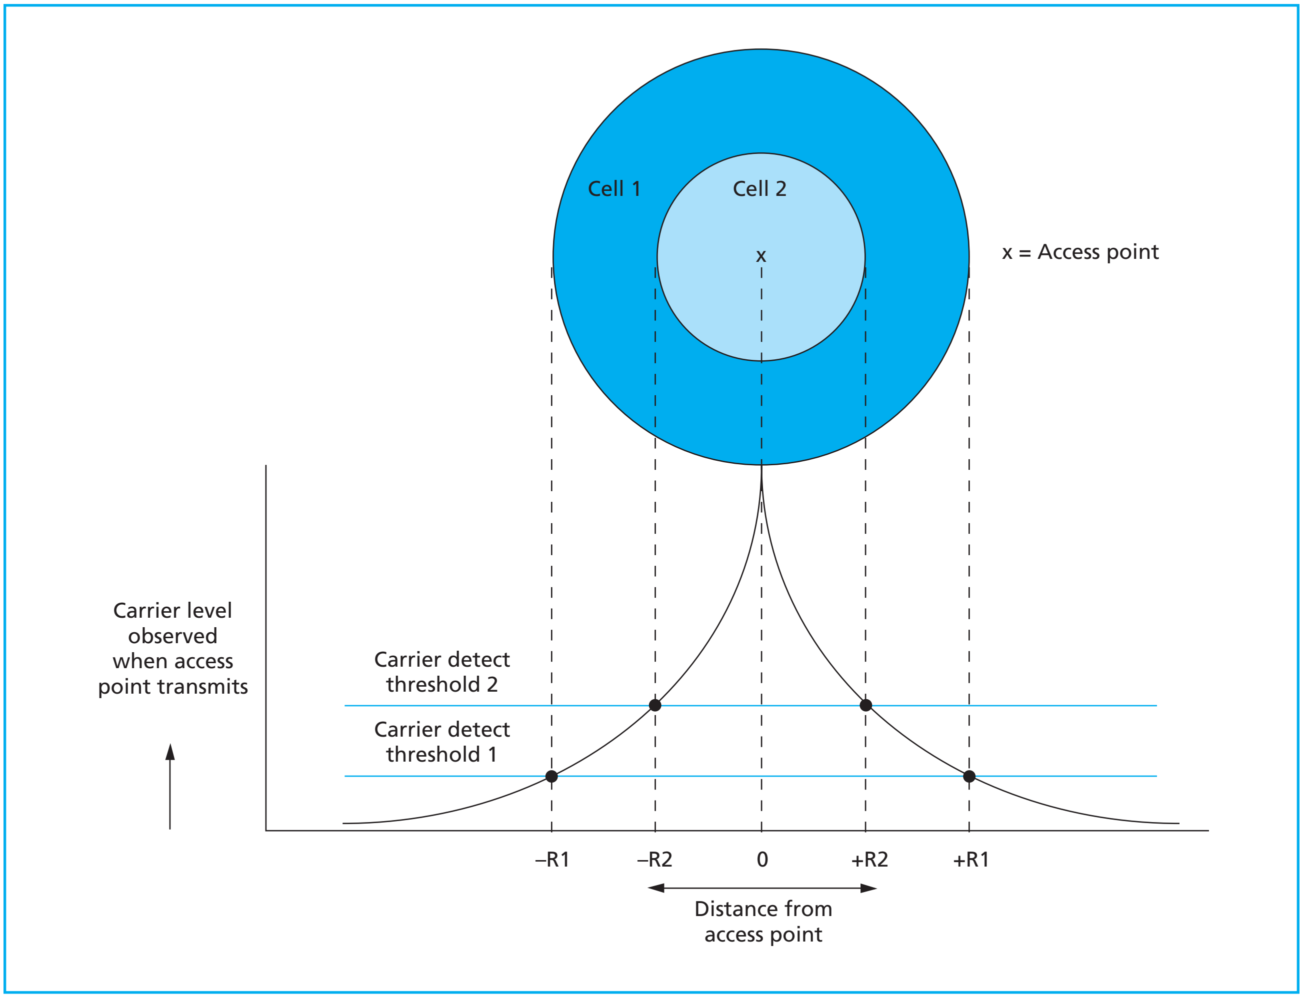

Figure 8.
The effect of carrier detect threshold on cell size.

the formation of queues at the access point, which fill up during large message transfers (for example, when a file is downloaded to a station).

- The PS-POLL mechanism, which polls each message individually, is inefficient.
- When a long sleep time occurs in transaction-oriented systems, it is useful to keep the receiver awake long enough after a transmission to cater to the typical application or protocol stack response time. For example, if the application always responds to a transaction request within (say) 3 seconds, then the holdover can be set at 3 seconds, ensuring that the station will still be awake to receive the response.

WaveLAN-II has three modes of operation for both 3.3V and 5V supplies: the transmit mode, with a used current draw of 300 mA; the receive mode, with a draw of 250 mA; and the DOZE mode, with a draw of 9 mA. Power management can reduce current consumption significantly.

Scalable System

The WaveLAN-II system will be deployable in a broad variety of situations, imposing different and often conflicting requirements on the system behavior. In particular, a stand-alone single cell network will have different requirements for the transmit and receive behavior than an infrastructure network, which contains multiple overlapping cells. To accom-

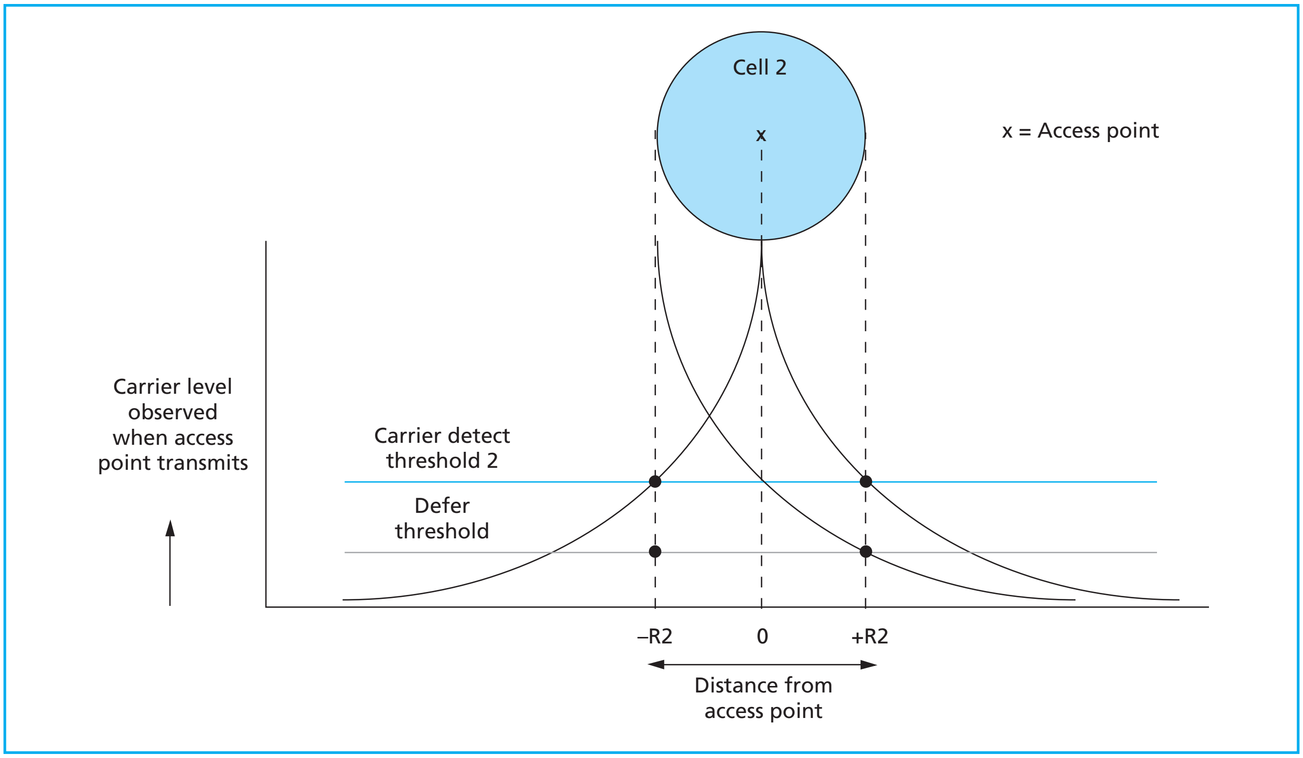

Figure 9.
Ideal relation between the defer threshold and the carrier detect threshold.

modate these variations in operation, the WaveLAN-II products have a number of built-in provisions to create scalable systems optimized for environmental and network usage needs.

The power level of a WaveLAN-II transmitter is 15 dBm (30 mW). When an access point transmits, the carrier signal levels observed by a station will decrease with distance. **Figure 8** illustrates the typical curve for the signal level in two opposite directions. This curve forms the basis for the scalability control elements.

Configurable Carrier Detection

Figure 8 shows two values of the so-called carrier detect threshold (CDT), defined as “the carrier signal level, below which the WaveLAN-II receiver will not do a receive.”[7] The CDT at level 1 crosses the curves of Figure 8 at distances -R1 and +R1, which implies an associated cell size for this CDT value with radius R1. This is shown as a bright color circle above the curves. The smaller cell, Cell 2 (radius R2), is higher (less sensitive) than level 1. The range for meaningful CDT levels has a lower boundary, determined by the sensitivity of the WaveLAN-II receiver circuitry. Setting the CDT to a lower value will result in a number of meaningless receive attempts with a high failure rate. Configurable carrier detection allows the WaveLAN-II cards to be configured at smaller cell sizes than the receiver is capable of handling. Small cell sizes play an important role in reusing the same channel in a relatively small area.

Configurable Defer Behavior

The 802.11 medium access rules (CSMA/CA) are based on the defer and random back-off behavior of all stations within range of each other. The defer decision is based on a configuration entity called the defer threshold (DT). When a carrier signal level is observed above the DT level, the WaveLAN-II card holds up a pending transmission request.

If we consider the example of CDT level 2 (Cell 2) from Figure 8, the ideal DT value produces a double radius, shown in **Figure 9**. A station on one edge of the cell defers to a station at the farthest edge. We can show that by plotting the curve for one edge station and ensuring that the DT level crosses this curve at the other cell edge. Choosing this relation between CDT

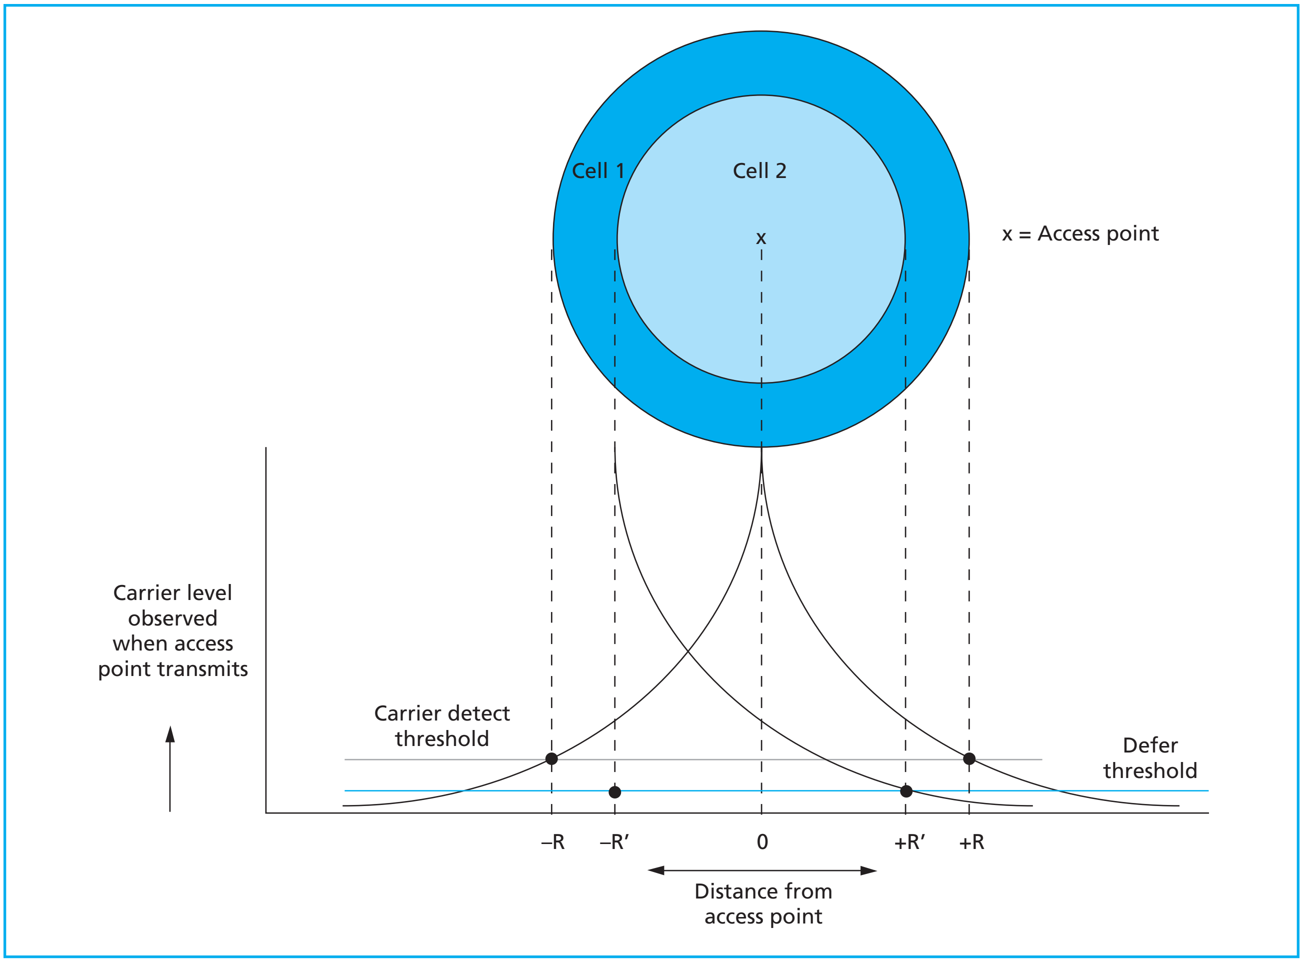

Figure 10.
Large cell characteristics.

and DT levels produces a cell in which all stations defer to each other, and where each station can communicate with the access point, thereby avoiding hidden station problems within the group of stations belonging to this cell.

The range for the DT level has a lower boundary, determined by the sensitivity of the WaveLAN-II carrier detect circuitry. Below a certain level, the signal will not be detected and no defer will take place. The ideal relation, shown in Figure 9, cannot be achieved if the CDT is set to the lowest (and most sensitive) level. In that case the lowest meaningful DT will not guarantee the wanted deferral between two "edge stations," as illustrated in **Figure 10**.

When we choose a low CDT value, we create a large cell size with radius R, shown as the large color circle. When the lowest meaningful DT level is plotted, the area where mutual deferral occurs has a smaller size, shown as the smaller color circle of radius R′. The outer area of the cell (radius R) will not have guaranteed deferral. This combination of thresholds can be used with the RTS/CTS medium reservation mechanism to prevent the hidden station phenomenon. The total cell is referred to as the basic coverage area (BCA). In the smaller color area, called the shared coverage area (SCA), the 802.11 medium-sharing rules will be in effect. In the ideal setup, the SCA is equal to the BCA.

Roaming Thresholds

Creating a cellular infrastructure system with the above-defined thresholds for the low-level receiver and transmitter control requires a proper balance with the roaming thresholds described earlier. The CDT and

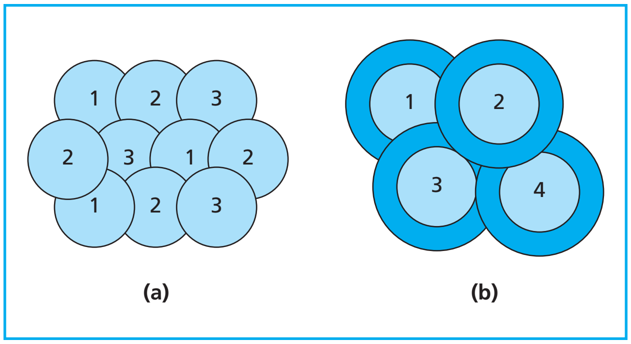

Figure 11.
High-density versus low-density infrastructure.

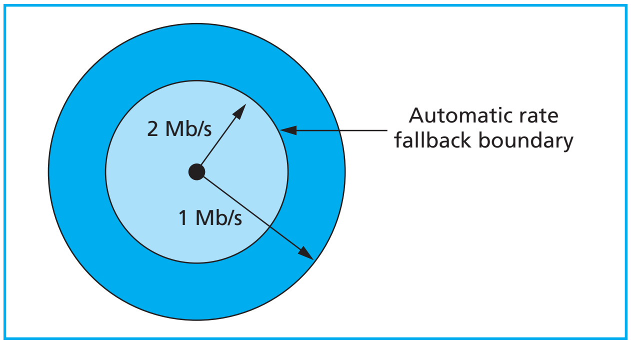

Figure 12.
The relation between data rate and cell regions.

DT determine the transmit/receive behavior of stations and access points that belong to the same cell, and the roaming thresholds determine the moments for deciding to start or stop participation in a cell. A station should base its handover decisions on the currently configured capabilities of the receiver. In particular, when small cell sizes are required, the roaming thresholds must be set to ensure that stations will start searching for a new (better) access point before the receiver becomes physically incapable of receiving messages from the current access point.

Infrastructure Density

The ability to define scalable cell sizes translates directly into the ability to control the density of cells (access points) that cover a certain area. Increasing the number of small cells within an area will increase the use of the same channel, thereby maximizing the overall throughput that can be obtained by fewer large cells.

Case A in **Figure 11** illustrates a high-density coverage setup. In such a setup, the choice of thresholds is such that the SCA and BCA of the cells coincide. At this level of the DT, no defer occurs for traffic at two stations on the same channel one cell diameter apart. If we use three independent channels, the roster can be filled completely. In the example, the raw throughput capacity is 10 times that of a single cell. At 2-Mb/s operation, this gives a 20-Mb/s composite capacity. When a single channel is used (for instance, forced by regulations in Japan), the addition does not work this way because of the influence of adjacent cells. If we use the same example, the overall raw capacity would be on the order of three to four times the single cell.

Case B in Figure 11 illustrates a maximum cell size setup. Here the SCA is smaller than the BCA. The overall capacity is four times that of a single cell (or roughly two times with a single channel). This lower capacity lowers the cost of the total system. Only four access points need to be installed versus the ten required for case A.

Automatic Rate Fallback

The different modulation techniques used for the different data rates of WaveLAN-II can be characterized by more robust communication at the lower rate. This translates into different reliable communication ranges for the different rates, 1 Mb/s giving the largest range. When a WaveLAN-II system is dimensioned for the largest cell size, the 1 Mb/s range is chosen. Stations moving around in such a large cell will be capable of higher-speed operation in the inner regions of the cell. To ensure usage of the highest practicable data rate at each moment, WaveLAN-II includes an ARF algorithm. This algorithm causes a fallback to the lower rate when a station wanders to the outer regions and an upgrade to the higher rate when it moves back to the inner region. **Figure 12** shows the two cell regions associated with the two data rates. The ARF functions come into play when the ARF boundary is crossed in either direction.

Besides enlarging the range, the lower rates will also be more robust against other interfering conditions like high path loss, high background noise, and extreme multipath effects. The ARF scheme will perform a temporary fallback when such conditions appear and an upgrade when they disappear.

The ARF scheme used in WaveLAN-II for IEEE 802.11 operation, which alternates between 2- and 1-Mb/s operation, is based on keeping track of a timing function and missed acknowledgment (ACK) frames. Operation at 2 Mb/s is considered the default. When an ACK is missed for the first time following earlier successful transmissions, the first retry transmission is still performed at the same (2 Mb/s) rate. When the ACK is missed again, the second retry and subsequent transmissions are also performed at the fallback (1 Mb/s) rate. A timer is started to track good ACKs and missed ACKs. When either the timer expires or the number of successively received good ACKs reaches 10, the device attempts to upgrade the rate. A new transmission is sent at the 2 Mb/s rate. When this fails (missed ACK), the system immediately reenters the fallback condition and resumes normal operation.

If higher data rates (5, 8, or 10 Mb/s) are used, the system follows an enhanced ARF scheme, which distinguishes multiple regions and aims at optimum data rates per region.

Conclusion

WaveLAN-II will be one of the early product implementations of the IEEE 802.11 wireless LAN standard, offering the broad suite of functions defined in the standard. The modem design and the choice for the DSSS technology will enable robust, efficient system design and seamless roaming facilities for multichannel wireless LANs. The added features of the product, such as radio scalability, automatic rate fallback, and enhanced power management will further differentiate the product.

A key aspect of the WaveLAN-II design is the ability to upgrade its radio to the higher data rates of 5, 8, and 10 Mb/s without major redesign. This design will allow the introduction of a new class of high-speed radio LANs that can fully coexist with stations following the IEEE 802.11 standard, and that can interoperate with such stations through access points.

References

1. K. Pahlavan, T. H. Probert, and M. E. Chase, "Trends in Local Wireless Networks," *IEEE Communications Magazine*, Vol. 33, No. 3, Mar. 1995, pp. 88–95.
2. R. O. LaMaire, A. Krishna, P. Bhagwat, and J. Panian, "Wireless LANs and Mobile Networking: Standards and Future Directions," *IEEE Communications Magazine*, Vol. 34, No. 8, Aug. 1996, pp. 86–94.
3. H. S. Chhaya and S. Gupta, "Performance of Asynchronous Data Transfer Methods of IEEE 802.11 MAC Protocol," *IEEE Personal Communications*, Vol. 3, No. 5, Oct. 1996, pp. 8–14.
4. B. Tuch, "Development of WaveLAN, an ISM Wireless LAN," *AT&T Technical Journal*, Vol. 72, No. 4, July/Aug. 1993, pp. 27–37.
5. A. Kamerman, "Spread Spectrum Schemes for Microwave-frequency WLANs," *Microwave Journal*, Vol. 40, No. 2, Feb. 1997, pp. 80–90.
6. I. Bar-David and R. Krishnamoorthy, "Barker Code Position Modulation for High-Rate Communication in the ISM Bands," *Bell Labs Technical Journal*, Vol. 1, No. 2, Autumn 1996, pp. 21–40.
7. A. Kamerman, "Throughput Density Constraints for Wireless LANs Based on DSSS," *Proc. IEEE Fourth Int. Symp. on Spread Spectrum Techniques and Applications (ISSSTA '96)*, Mainz, Germany, Sept. 1996, pp. 1344–1350.

(Manuscript approved August 1997)

AD KAMERMAN is a member of technical staff assigned to the WaveLAN-II System Team in the Wireless Communication and Networking Division of Lucent's Network Systems in The Netherlands. He works on wireless LANs, designing their system performance as it relates to the radio-frequency environment, signal processing, and network throughput. Mr. Kamerman received a B.S. and M.S. in electrical engineering from Twente University of Technology in The Netherlands.

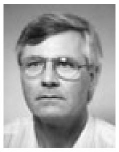

LEO MONTEBAN received a B.S. in mathematical engineering and an M.S. in mathematical engineering and computer science from Delft University of Technology in The Netherlands. A member of technical staff and part of the WaveLAN-II System Team in the Wireless Communication and Networking Division of Lucent's Network Systems in The Netherlands, Mr. Monteban develops software for the WaveLAN-II product family and is responsible for its overall system architecture. ◆