Originally developed to let WLANs share spectrum with radar and satellite systems, 802.11h defines mechanisms that can reduce both interference and power consumption.

Daji Qiao and Sunghyun Choi

New 802.11h Mechanisms Can Reduce Power Consumption

With the rise of portable computing and communication devices such as laptops and PDAs, millions of people worldwide have come to rely on wireless networks for communication. For indoor broadband wireless networking, the wireless local area network, standardized as Std. IEEE 802.11 WLAN, has become the dominant technology. Known as hotspots, public WLANs have sprung up in conference venues, airport lounges, cafes, and coffee shops. Near the end of 2003, a total of 31,580 hotspots served 1.53 million users around the world ("Public Wireless LANs: Hotspots—Finally Heating Up?" http://www.datamonitor.com, report no. DMTC0921, Sept. 2003). By January 2006, the number had grown to 100,355 hotspots in 115 countries ("Worldwide Wi-Fi Hotspots Hits the 100,000 Mark," http://www.jiwire.com/press-100k-hotspots.htm).

The original 802.11 standard—which the IEEE 802.11 working group began in 1991, first published in 1997, and revised in 1999—specifies the protocols for a WLAN's medium access control (MAC) sublayer and physical (PHY) layer. During the past few years, the working group has focused on standardizing new specifications to enhance WLAN performance. Figure 1 shows relationships among the existing and emerging specifications; the arrows indicate enhancements springing from established standards ("Overview of Emerging IEEE 802.11 Protocols for MAC and Above," Sunghyun Choi, *SK Telecom—Telecommunications Rev.*, Nov. 2003). For example, 802.11e is the quality-of-service amendment to 802.11 MAC. The "Hotspot Standard Rundown" sidebar lists descriptions (and publication details, when available) for the standards in Figure 1.

The emerging standard 802.11h, which involves amendments to both MAC and PHY, specifies mechanisms for dynamic frequency selection (DFS) and transmit power control (TPC) for WLAN devices. Although the standard seeks to satisfy European regulatory requirements, the two mechanisms have numerous other applications that will be useful in all countries—automatic frequency planning, power consumption reduction, range control, interference reduction, and quality-of-service enhancement.

IEEE 802.11H FOR SPECTRUM AND TRANSMIT POWER MANAGEMENT

The original purpose of 802.11h was to extend WLAN operation in Europe to the 5-GHz band,

1520-9202/06/$20.00 © 2006 IEEE

where WLAN devices need DFS and TPC to coexist with the band's primary users—radar and satellite systems. According to European regulations, when a WLAN device in the 5-GHz band detects a radar signal on its current operational frequency channel, it must switch to another channel. When a WLAN device detects a satellite signal, it must limit its transmit power to 3 dB below the regulatory maximum (ordinarily, it could transmit at up to the regulatory maximum level). To meet these requirements, 802.11h defines DFS and TPC mechanisms on top of 802.11 MAC and 802.11a PHY.

Figure 2 diagrams the 802.11h layer management model that enables the DFS and TPC functions. Policy decisions, such as channel switching and TPC, reside in the station management entity (SME); the associated protocols reside in the MAC sublayer management entity (MLME). Interactions between entities take place through service access points (SAPs), across which primitives are defined and exchanged.

Although 802.11h defines the mechanisms or protocols for DFS and TPC, it does not cover the implementations themselves. Both functions involve implementation-dependent algorithms. So an 802.11h-compliant device needs a TPC algorithm, for example, to correctly determine a frame transmission's transmit power level.

Dynamic frequency selection

An 802.11 WLAN's basic building block is the basic service set, the BSS, which consists of a set of stations controlled by a single coordination function and operating in the same frequency channel. An 802.11a/h BSS occupies the 20-MHz channel: the US has 12 channels available for 802.11a, and Europe has 19. The WLAN uses DFS for dynamically switching a BSS from one operational frequency channel to another. Besides fulfilling regulatory requirements, a BSS might have many other reasons to change its operational frequency channel. It might do so, for instance, when the condition of its current channel becomes poor because of interference from neighboring devices. In this context, DFS enhances the WLAN's quality of service.

In an infrastructure BSS, which consists of an access point (AP) and multiple stations associated with it, the AP determines when to switch and which channel to switch to. For this purpose, the AP monitors the status of the current channel, and it might also request that other stations measure and report the status of the current channel and those at other frequencies. There are three types of measurement:

- *basic measurement* determines whether another BSS, a non-802.11 OFDM (orthogonal frequency-division multiplexing) signal, an unidentified signal, or a radar signal is using the measured channel;
- *clear-channel assessment (CCA)* measures the fractional duration of the channel's busy period during the total measurement interval; and
- *received power indication (RPI)* measures the histogram of the quantized measures of the received energy power levels as seen at the antenna connector during the measurement interval.

Based on its own measurement and reports from the associated stations, the AP continues to monitor the channel status so that the BSS can switch channels at the proper instance. Channel switching occurs immediately before a target beacon transmission time (TBTT), which the AP

Hotspot Standard Rundown

- **802.11 Working Group Web site, http://www.ieee802.org/11**
- **Main standard, 802.11, MAC and PHY for WLANs: IEEE Std. 802.11-1999 (R2003), Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications**
- **802.11a, Orthogonal Frequency-Division Multiplexing (OFDM) in the 5 GHz Band: IEEE Std. 802.11a-1999 (R2003), Supplement to IEEE Std. 802.11-1999: High-Speed Physical Layer in the 5 GHz Band**
- **802.11b, Complementary Code Keying (CCK) in the 2.4 GHz Band: IEEE Std. 802.11b-1999 (R2003), Supplement to IEEE Std. 802.11-1999: Higher-Speed Physical Layer Extension in the 2.4 GHz Band**
- **802.11e, Quality of Service: IEEE 802.11e-2005, Amendment 8 to IEEE Std. 802.11-1999: Medium Access Control (MAC) Quality of Service (QoS) Enhancements**
- **802.11g, Orthogonal Frequency-Division Multiplexing (OFDM) in the 2.4 GHz Band: IEEE Std. 802.11g-2003, Amendment 4 to IEEE Std. 802.11-1999: Further Higher Data Rate Extension in the 2.4 GHz Band**
- **802.11h, DFS and TPC: IEEE Std. 802.11h-2003, Amendment 5 to IEEE Std. 802.11-1999: Spectrum and Transmit Power Management Extensions in the 5 GHz Band in Europe**
- **802.11i, Security: IEEE Std. 802.11i-2004, Amendment 6 to IEEE Std. 802.11-1999: Medium Access Control (MAC) Security Enhancements**

specifies, so that normal communication can take place for the following beacon interval at the new operational frequency channel. The AP transmits the beacon frames periodically. DFS also defines a channel-quieting operation, because European regulations require the station to become silent once it detects a radar system in its operational frequency channel. We explain an example DFS algorithm in a later section.

Figure 1. IEEE 802.11 MAC/PHY and its enhancements.

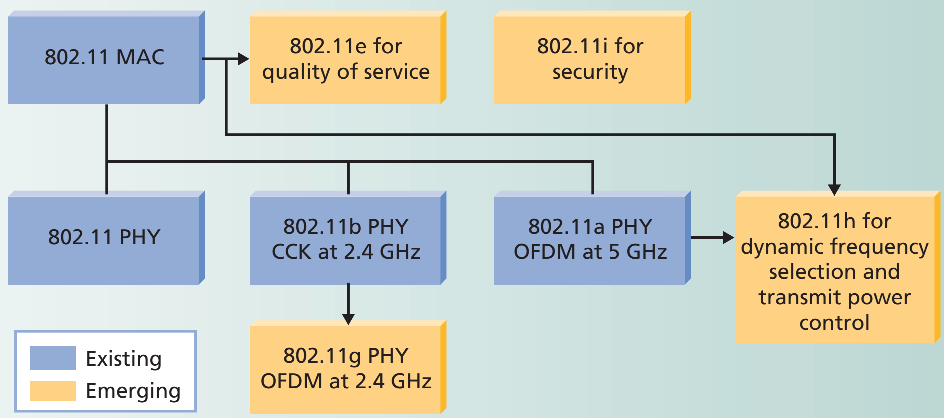

Transmit power control

For wide-area cellular systems, such as IS-95 CDMA (code-division multiple access) and 3G W-CDMA (third-generation wide-band CDMA), TPC has three critical purposes:

- ameliorating the near-far problem (if all transmissions use the same power, those closer to the receiver drown out the more distant ones)—specifically, for CDMA uplink transmissions;
- minimizing interference to and from other cells (cochannel interference); and
- improving system performance on fading channels by compensating for fading dips.

Until recently, TPC has not been considered as critical to success for 802.11 WLANs; most of today's 802.11 devices use fixed transmit power for frame transmissions. In recent years, however, because 802.11a devices must support TPC to meet European regulatory requirements, researchers have zeroed in on finding other applications of TPC in 802.11 WLANs. Because most WLAN devices are battery powered, the idea of using TPC to preserve battery energy and thus extend battery life is especially popular. In the multicell WLANs common in offices and public-access environments, TPC can also reduce intercell interference, because it improves error performance in a given area.

Standard 802.11h specifies two TPC-related functions. First, the AP in an infrastructure BSS or a wireless station in an independent BSS advertises the regulatory and local maximum-power level for the current frequency channel in the Beacon and Probe Response frames. They do so using Country and Power Constraint elements. The local maximum specifies the actual maximum power level allowed in the BSS, which is less than or equal to the regulatory maximum. The stations in the BSS can use any transmit power less than or equal to the local maximum value.

Figure 2. Layer management model in 802.11h.

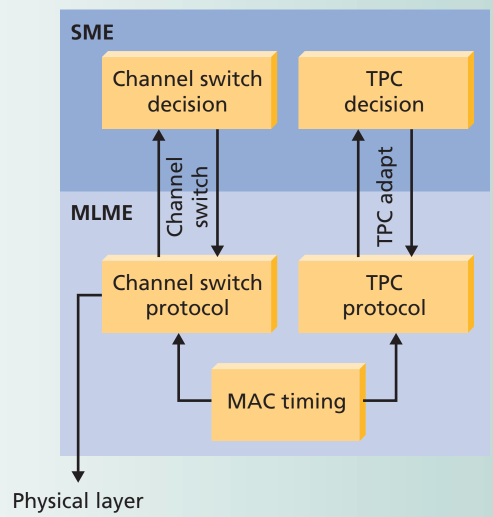

Second, 802.11h provides a transmit-power reporting mechanism. It defines a TPC Report element that contains Transmit Power and Link Margin fields. The Transmit Power field simply contains the transmit power used to transmit the frame containing the TPC Report element, while the Link Margin field contains the link margin, calculated as the ratio of the received signal strength (of the corresponding TPC request frame) to the minimum desired by the station. TPC Report elements are included in TPC Report frames in response to TPC Request frames. In addition, the AP in an infrastructure BSS or a wireless station in an independent BSS autonomously includes a TPC Report element in any Beacon frame it transmits.

Figure 3. Finite-state machine for a DFS algorithm.

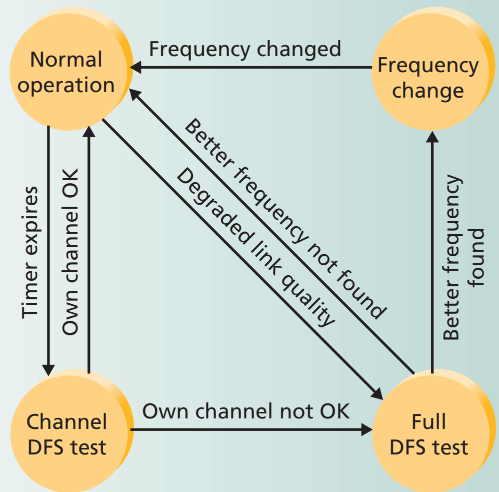

Such a TPC Report element has a Link Margin field set to zero along with the frame's transmit power information.

An immediate impact of such a transmit-power reporting mechanism is that it makes energy-efficient frame transmissions feasible in 802.11h WLANs. In general, to determine the proper transmit power level for a given frame, a wireless station must estimate the link quality between itself and the receiver. With 802.11h, a simple estimation scheme could work as follows. Whenever a wireless station receives a frame including the TPC Report element, it can use the transmit power information in that element and the receiver signal strength information obtained via RSSI (receiver signal strength indicator) to estimate link quality (in terms of path loss) from the sending station to itself by performing a simple subtraction.

AN EXAMPLE DFS ALGORITHM

In general, a DFS algorithm for an infrastructure-based 802.11h system consists of the following two phases: startup and regular. At startup, the AP performs a full DFS measurement on all frequency channels. Based on the measurement results, the AP selects a starting frequency channel that is unoccupied by the primary users for its BSS. Moreover, the AP tries to avoid selecting the frequency channels that are already used by other secondary users—other operating 802.11a/h systems.

Once the algorithm chooses the starting frequency, the AP begins broadcasting Beacon frames so that the wireless stations can detect its presence and associate themselves with the AP. The BSS then starts and enters the normal-operation state of the regular phase. In the regular phase, the DFS algorithm is better described with the finite-state machine shown in Figure 3 (IEEE 802.15-01/072, "Liaison Statement on the Compatibility between IEEE 802.11a and Radars in the Radio Location and Radio Navigation Service in the 5250-5350 MHz and 5470-5725 MHz Bands," IEEE 802.15 Working Group, Jan., 2001); it has four different states: normal operation, channel DFS test, full DFS test, and frequency change.

The BSS remains in normal operation until either the measurement timer expires or the link quality degrades, which suggests possible interference. If the timer expires, the state changes to channel DFS test, during which the AP reassesses the current operating-frequency channel. If the channel status is good enough to continue operation, the state reverts to normal operation; otherwise, it changes to full DFS test, during which the AP performs a full DFS measurement on all frequency channels. During normal operation, if the link quality degrades, state changes directly to full DFS test.

In full DFS test, the AP measures (or asks wireless stations to measure) all the frequency channels. Then, based on the measurement results, the AP makes the channel switching decision. If the current frequency channel is the best in terms of link quality, the BSS continues operating on the current channel, and state reverts to normal operation. Otherwise, the state changes to frequency change, in which the AP wakes up all the sleeping wireless stations, if any, announces what the new frequency channel is, and when operation in the new frequency channel will start. Finally, after the frequency channel switches successfully, the state returns to normal operation.

MISER: AN INTELLIGENT TPC MECHANISM

Because an 802.11h wireless station can use the transmit-power reporting mechanism to estimate link quality between itself and the receiver, it can also apply intelligent TPC to minimize communication energy consumption. So, based on the estimation scheme, we have developed an optimal low-energy transmission strategy for 802.11h called MiSer ("MiSer: An Optimal Low-Energy Transmission Strategy for IEEE 802.11a/h," Daji Qiao and colleagues, *Proc. ACM MobiCom 03*, ACM Press, 2003, pp. 161-175).

MiSer challenge

According to the 802.11 standard, a WLAN device can operate in four modes: transmit, receive, idle, and doze. It consumes the most power in transmit mode and very little energy in doze mode. In idle mode, a WLAN device must sense the medium, so it consumes a similar amount of power as when it is in receive mode. Researchers have proposed numerous power management policies to send a WLAN device into doze mode adaptively to save battery energy ("Dynamic Power Management for Portable Systems," Tajana Simunic and colleagues, *Proc. ACM MobiCom 00*, ACM Press, 2000, pp. 11-19). An alternative way to conserve energy is to apply TPC, which lets a

WLAN device use the minimum-required power level in transmit mode. Thus TPC can complement other power-management policies.

MiSer design

Because 802.11 DCF (distributed coordination function) systems necessarily involve contention, MiSer's effectiveness relies on the condition that applying TPC to data transmissions will not aggravate the *hidden-nodes problem* or interference in the network. The hidden-nodes problem occurs when a wireless node cannot hear from one or more of the other nodes so that multiple nodes might attempt to transmit simultaneously, causing frame collisions and signal interference with each other. To deal with such problems, MiSer exchanges RTS/CTS (request to send/clear to send) frames before each data transmission attempt. More importantly, it transmits the CTS frames at a stronger power level, which not only allows data frames to be transmitted at appropriate transmission rates and transmit power levels to save energy, but also ameliorates the potentially aggravated interference that TPC might cause.

Figure 4. Snapshot of MiSer's rate-power combination table.

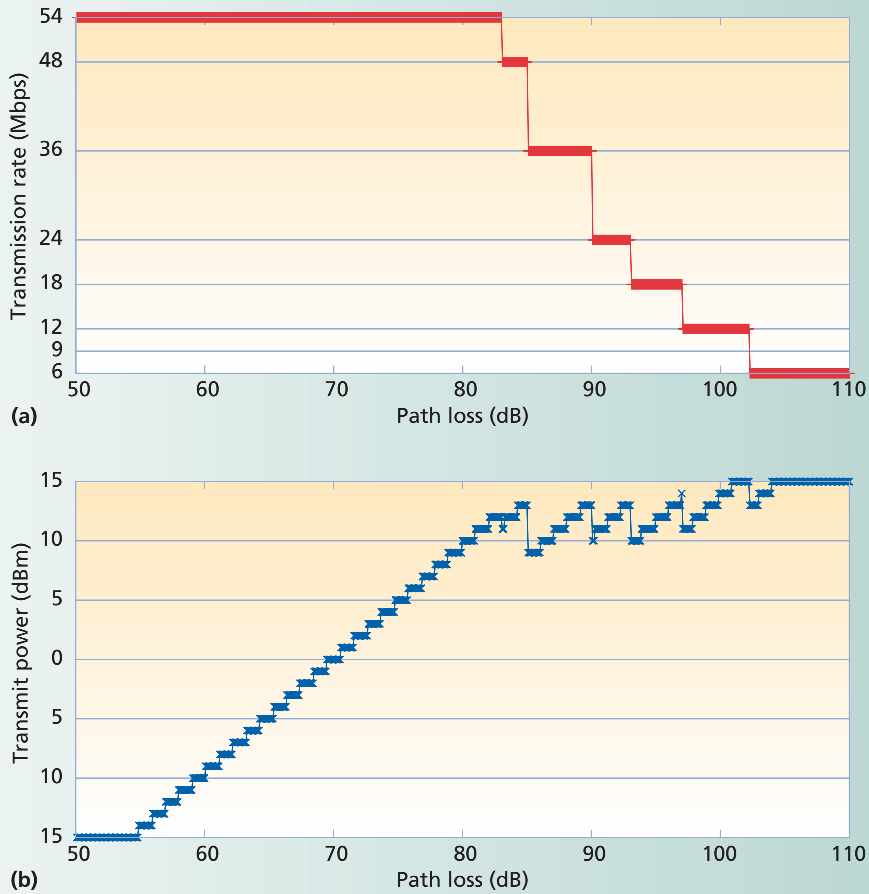

MiSer uses a simple, table-based approach to determine the most energy-efficient combination of transmission rate and transmit power for each data frame. The basic idea is that, offline, a wireless station computes a rate-power combination table indexed by the data transmission status; each table entry is the optimal rate-power combination in the sense of maximizing energy efficiency under the corresponding data transmission status. MiSer characterizes data transmission status as a quadruplet of data payload length, path loss between the transmitter and the receiver, and two frame retries. It defines energy efficiency as the ratio of the expected delivered data payload to the expected total energy consumption. At runtime, the wireless station estimates the path loss between itself and the receiver, updates the data transmission status, and then selects the proper transmit rate and power for the current data transmission attempt through a simple table lookup.

Figure 4 shows a snapshot of MiSer's rate-power combination table for a certain data payload length and fixed numbers of remaining transmission attempts. The two parts of the figure show the transmit rate and power that achieve the most energy-efficient data communications, under different path loss conditions. For example, when the path loss is 80 dB, the most energy-efficient strategy is to transmit at 54 Mbps with 9 dBm power. Other snapshots resemble shifted versions of Figure 4. For example, when a data frame carries a larger payload or fewer transmission attempts remain for a data frame, the graphs shift left and a more conservative combination—lower rate, higher power, or both—becomes optimal under the same path loss condition.

MiSer performance

Simulation results show MiSer's clear superiority to the two-way Data-Acknowledgment frame exchange mechanisms in the presence of hidden nodes. In addition, compared with other four-way RTS-CTS-Data-Acknowledgment frame exchange mechanisms, MiSer delivers about 20 percent more data per unit of energy consumption than a scheme using the transmission rate adaptation alone without TPC. It also significantly outperforms single-rate TPC schemes, thanks to the excellent energy-saving capability of transmission rate adaptation.

As happens with any popular technology, along with WLAN's worldwide success has come a clamor for the technology's evolution. Among the past few years' remarkable enhancements to 802.11, the spectrum and transmit power management tools of 802.11h provide a particularly fertile bed for innovation. MiSer, our simple application example for saving significant communication energy for 802.11h devices, is but one of many possible exciting new directions for 802.11 technologies. ■

Daji Qiao *is an assistant professor at Iowa State University. Contact him at daji@iastate.edu.*

Sunghyun Choi *is an assistant professor at Seoul National University. Contact him at schoi@snu.ac.kr.*

For further information on this or any other computing topic, visit our Digital Library at http://www.computer.org/publications/dlib.