The Family Dynamics of 802.11

The 802.11 family of standards is helping to move wireless LANs into promising new territory.

Three trends are driving the rapid growth of wireless LAN (WLAN): The increased use of laptops and personal digital assistants (PDAs); rapid advances in WLAN data rates (from 2 megabits per second to 108 Mbps in the past four years); and precipitous drops in WLAN prices (currently under \$50 for a client and under \$100 for an access point).

As a result, 802.11 technology [see 802.11 Standard, ISO/IEC 8802-11:1999 (E), ANSI/IEEE; and *802.11 Handbook, A Designer's Companion*, by B. O'Hara and A. Petrick, IEEE Press, 1999] is sure to become ubiquitous. The WLAN market is expected

to grow from \$1.79 billion in 2001 to \$3.85 billion in 2004 [see "It's Cheap and It Works: Wi-Fi Brings Wireless Networking to the Masses," by G. Paulo, *In-Stat MDR*, December 2002]. Industry analysts have predicted that by the end of 2004, 70 percent of laptops will come with 802.11 technology already embedded [*Wireless Applications Proliferate*, by Chris Kozup, META Group, Feb. 24, 2003].

Numerous other wireless technologies—including Bluetooth, broadband wireless access, cellphones, and 3G networks—are also evolving, but none is experiencing the rapid growth of 802.11. Because the 802.11 standard covers only

BILL MCFARLAND AND MICHAEL WONG,
ATHEROS COMMUNICATIONS

The Family Dynamics of 802.11

layers 1 and 2 of the networking stack (the physical and media access control layers), we will not cover the higher-level networking aspects of wireless LANs in as much detail here.

ONCE AND FUTURE APPLICATIONS

The largest driver for 802.11 products is “traditional” networking at home and in the office. In these networks, the traffic is primarily TCP/IP and looks much like the traffic over wired LANs. In offices, wireless LANs have generally been installed as overlay networks, on top of wired networks, to provide connectivity in conference rooms and cafeterias, as well as to allow Internet access. Early generations of 802.11 technology have not had sufficient throughput and overall system capacity to allow offices to go completely wireless. The emergence of 5-gigahertz 802.11a, however, permits moderate-sized offices to “unwire.”

In the home, the wireless network serves primarily as a convenient way to provide mobile access to an external Internet connection such as DSL or a cable modem. In most cases, homes with wireless networking have no wired network.

The 802.11 standard provides for ad hoc networking, which allows users to form a networking group without any infrastructure. Although this is a convenient feature, it is not widely used. Networking without access to the Internet is only of modest importance.

Several new applications promise huge growth potential for 802.11 technology. Wireless hotspots are in their infancy, but will mature quickly, aided by the low cost of 802.11 technology. It's easy to envision hotspots located in hotels, coffee shops, airports—anywhere that people have a chance to pause long enough to read e-mail or surf the Web.

Consumer electronics applications will use 802.11 extensively. Most of the large consumer electronics companies are developing wireless video connections between set-top boxes and flat-panel TVs. Future applications will

include video streaming from mobile devices such as digital camcorders and cameras.

Although never intended for this application, 802.11 is also being considered for use in wide-area mobile networks. The low cost and ubiquity of clients makes creating such a system attractive. Estimates show much lower infrastructure costs than 3G, not to mention operation in free unlicensed spectrum. Significant technical hurdles will have to be cleared, however, before a wide-area 802.11-based system will work.

A REGULATORY PUSH

Government regulation is a major factor in wireless communication. Since the mid-'80s, more and more spectrum has been allocated for free and unlicensed use. The most important unlicensed allocations are at 2.4 GHz and 5 GHz. Spectrum from 2.400 to 2.4835 GHz has been available in most countries for many years. In 1997 the U.S. government allocated 5.15 to 5.35 and 5.725 to 5.825 GHz. Europe and Japan made similar allocations.

Just recently, a grand compromise was reached in the U.S., involving industry, the military, and the government, that will allocate 5.15 to 5.850 GHz for unlicensed use. Much of the world is expected to follow this example, freeing 700 MHz of spectrum for WLAN use in the next few years.

The newest development in regulation is the U.S. government's allowance of ultra-wideband (UWB) technology in the frequency band from 1.99 to 10.6 GHz. UWB spreads radio transmissions over a wide bandwidth, keeping the power spectral density across any given frequency range low. Using spread spectrum-type processing, UWB receivers can detect these very weak signals, while causing virtually no interference to normal receivers already operating in these bands. Although such a large allocation is an opportunity, the required low transmission power levels (two to three orders of magnitude lower than the power allowed in the other unlicensed bands), and the availability of spectrum only in the U.S., will limit the commercial viability of these systems for some time.

WE HAVE STANDARDS

A number of WLAN standards have developed over the years, including HiperLAN 1 and 2 and HomeRF. Only 802.11, however, has had any significant commercial impact, and it is now the only WLAN standard that matters. It is actually a family of standards that is constantly being extended. These extensions are considered supplements to the original 802.11 standard and are identified by one-letter suffixes.

The following is a brief summary of the extensions that are completed or in progress:

- **802.11**—Original standard provided for 1- and 2-Mbps PHY layer, CSMA/CA MAC, adopted 1997
- **802.11a**—Enhancement to provide 54 Mbps in the 5-GHz band, adopted 1999
- **802.11b**—Enhancement to provide 11 Mbps in the 2.4-GHz band, adopted 1999
- **802.11d**—Changes for international regulatory compliance, adopted 2001
- **802.11e**—Enhancements to the MAC layer to provide Quality of Service (QoS) through prioritized CSMA and advanced polling techniques, still in debate
- **802.11f**—Recommended practices for inter-access point communication, nearly completed
- **802.11g**—PHY layer enhancement to provide 54 Mbps in the 2.4-GHz band, nearly completed
- **802.11h**—Enhancements to 802.11a to help with regulatory compliance in Europe, nearly completed
- **802.11i**—Enhancements for greater security, still in debate
- **802.11j**—Changes to meet Japanese regulatory requirements, just starting
- **802.11k**—Improved WLAN system management, just starting
- **802.11m**—A generalized cleanup and editing of the existing standard
- **802.11?**—Standard for greater-than-100-Mbps throughput, project not approved yet, will take three-plus years to complete

Through these supplements, the 802.11 standard can grow to address applications that were never considered when it was originally created. In particular, the QoS enhancements in 802.11e will allow the standard to serve the needs of voice and video, allowing 802.11 to become a cellular voice system or a consumer electronics network.

THE IMPORTANCE OF MODULATION

The most important aspect of the physical layer is the modulation system, because it determines the raw data rate, as well as much of the implementation complexity and resulting cost. One of the key aspects of the modulation system is how it overcomes multi-path (echoes),

inherent in wireless transmissions. Multi-path arises from reflections of the transmitted radio waves off objects in the environment. The reflections arrive with different time delays, and the resulting overlapping of earlier and later portions of the transmitted waveform can make reception difficult.

The currently popular 802.11b supplement (with raw data rates of up to 11 Mbps) uses complementary code keying (CCK). The input data stream is expanded according to a particular code, and the new data stream is transmitted using phase modulation of a single carrier. For example, if data enters the system at 11 Mbps, each block of 8 input bits is coded into 16 output data bits (commonly called *chips* to distinguish them from the original data bits). These 16 chips are then transmitted two at a time using eight-quadrature phase shift keying

By using **multiple** carrier waves, each at a slightly **different** frequency, it is possible to transmit a **high aggregate data** rate while the **symbol rate** on each carrier remains **low**.

(QPSK) symbols. Each symbol is a period of time in which the phase of the carrier sinusoid is at one of four discrete values, thereby conveying the two chips of information. Thus, 22 Mcips per second are transmitted, which are decoded at the receiver to recover the original 11-Mbps data stream.

The basic concept of expansion coding the input bits is sound. The redundancy allows for error correction at the receiver and spreads the frequency spectrum, making the system more resilient to some of the effects of multi-path. Unfortunately, the specific code selected was not a very wise choice. It achieves only half the benefits that a better code might have provided.

The other disadvantage to CCK modulation is the high symbol rate on the airwaves. This means that multi-path echoes can cause multiple symbols to overlap in time at the receiver. The standard way of combating this is to use adaptive equalizers, but they require a relatively long header for training. Because the echoes change significantly even with slight movement in the environment,

The Family Dynamics of 802.11

each packet must contain the overhead of this training header.

The computational complexity of an adaptive equalizer grows as the square of the data rate. Therefore, while an adaptive equalizer is workable for an 11-Mbps data rate, it would be impractical for the 54- or 108-Mbps rate seen in recent WLANs.

To avoid these limitations, 802.11a and 802.11g use Orthogonal Frequency Division Multiplexing (OFDM) [see *OFDM for Wireless Multimedia Communications*, by R. Van Nee and R. Prasad, Artech House Publishers, 2000]. The basic concept behind OFDM is to create very long symbols, such that multi-path echoes create overlapping during only a small portion of the symbol, which can be avoided by the receiver.

The time in which overlapping occurs as a result of multi-path is called the *guard interval*. It is set to 800 nanoseconds for 802.11a and g.

Providing a high data rate with a slow symbol rate is difficult, however. The solution is to transfer the data in parallel. By using multiple carrier waves, each at a slightly different frequency, it is possible to transmit a high aggregate data rate while the symbol rate on each carrier remains low.

Figure 1 describes how OFDM modulation works in both the time and frequency domains. In the frequency domain, 802.11 consists of 52 separate car-

rier frequencies, spaced as close as possible to minimize the overall bandwidth of the signal. The data is modulated on each carrier by varying the phase and amplitude of the carrier wave. In addition, an error-correcting code is applied to the data before modulation. A variety of data rates is supported by varying both the amount of error-correction coding and the complexity of the modulation. The data rate of each device in an operating WLAN is continuously adapted to provide the highest possible throughput, given its location and physical environment.

In the time domain, the transmitted waveform can be seen as the sum of 52 sinusoids, each at a different frequency, and each modulated in phase and amplitude to convey the data.

The advantage of this approach is that it does not require a long training header to correct for multi-path, and a Fast Fourier Transform (FFT) can be used to generate the required waveform very efficiently. Given data rate R , the FFT requires only $R/2 \cdot \log_2(R)$ computations, as opposed to the R^2 computations for an equalized system. Finally, the error-correcting code is much stronger. The end result is that 802.11a and 802.11g can far outperform 802.11b in real environments.

Two new methods promise further improvements in data rate and range. Both are "smart-antenna" tech-

nologies in which multiple antennae direct WLAN signals. Figure 2 shows a beam-forming [“Antenna Systems for Broadband Wireless Access,” by R. Murch and B. Letaief, *IEEE Communications Magazine*, April 2002, pp.76-83] access point (AP) using four antennae. In theory, one beam can be created per antenna element. In practice, more antenna elements are needed. The beams can be directed at different stations. If the stations are sufficiently separated from each other physically, it is possible for them all to communicate simultaneously without interfering with each other. In addition, the focusing of the power into beams can increase the range over which a connection can be maintained.

Closely related are multiple input multiple output (MIMO) systems. As shown in Figure 2, these systems have multiple antennae at both the AP and the station. Assuming that the environment has plenty of multi-path reflections, beams can be formed that will not interfere with each other. Each of the beams can then carry a different set of information, multiplying the data rate by $\min(\# \text{ Tx antennae}, \# \text{ Rx antennae})$. In practice, the increase in throughput is less, but research has demonstrated very high-capacity systems are possible [“Layered Space-Time Architecture for Wireless Communication in a Fading Environment When Using Multiple Antennas,” by G. J. Foschini, *Bell Labs Technical Journal*, Vol. 1, No. 2, Autumn 1996, pp. 41-59].

FIG 2

FIG 3

Both of these smart-antenna techniques are difficult because of rapid changes in the wireless channel as objects move. Algorithms for finding and then tracking the required beam directions are the key technology. In addition, the 802.11 MAC generally depends on the ability of all stations in a cell to hear all the other stations in the cell. Although they extend range, beam-forming tech-

The Family Dynamics of 802.11

niques create hidden nodes that cannot hear each other. Dealing with this efficiently within the 802.11 framework is a challenge.

THE MAC LAYER

The basic 802.11 media access mechanism is listen before talk, with slotted random backoff (CSMA/CA) and packet-by-packet acknowledgment (automatic repeat request, ARQ). Because it is difficult to build radios that can transmit and receive at the same time, detecting collisions is also difficult. Instead, a positive acknowledgment is sent for each correctly received packet. The sender knows to retransmit the packet if an acknowledgment is not received. Wireless networks often operate with packet loss rates of 10 percent because of noise and interference, so quick retransmission at the physical level is necessary.

Figure 3 shows how this basic mechanism is being enhanced in 802.11e to provide QoS. Different traffic streams can be given different priority access to the wireless channel by adjusting their initial waiting periods (the AIFS) and their random backoff contention window (CW min and CW max). This mechanism maps well to 802.1p QoS traffic commonly carried on wired networks, and it works well when there are a few high-priority streams.

When a large number of high-priority streams are coming from multiple sources, a more centralized scheduler may provide better performance. The 802.11e standard also provides a polling-based mechanism called the Hybrid Control Function (HCF). In this method, the access point informs all stations to hold transmissions until they are polled and informs them of how long they may transmit when polled. HCF allows the access point to schedule all transmissions on the medium, effectively converting the 802.11 system to a Time Division Multiple Access (TDMA) network.

UPPER PROTOCOL LAYERS

Security is a critical issue for any wireless communication system. Security can be broken into two categories: link-

layer encryption and the system-level security mechanisms that handle key distribution.

The original security system defined in 802.11 was called Wired Equivalent Privacy (WEP). Several holes have since been discovered in WEP, and it is no longer considered secure. The 802.11i supplement is taking several approaches to fix the problem. The following list summarizes the link-layer encryption options:

- **WEP**—The original simple encryption method, based on RC4. It is no longer considered secure. It provides privacy only, no integrity check.
 - **WPA** (also called TKIP, Temporal Key Exchange Protocol)—This is a temporary “fix” for WEP that scrambles the key between packets and adds a message integrity check to prevent spoofing.
 - **Advanced Encryption Standard (AES)**—From the National Institute of Standards and Technology (NIST), this is a standard symmetric block cipher intended to replace the Data Encryption Standard (DES). It is considered to be “military strength” and includes an integrity check. This is the long-term solution chosen by 802.11i.
- Because 802 standards are restricted to the PHY and MAC layers, 802.11i is adapting existing solutions for authentication and key handling:
- **802.1X**—A general-purpose and extensible framework for authenticating users and generating/distributing keys for encryption.
 - **Simple Secure Network (SSN/RSN)**—A recipe for authentication based on 802.1X, bridging between an 802.11 base station and an authentication server in the network. Radius servers are used as authentication servers.

A final important aspect of security for WLANs is the use of virtual LANs (VLANs). Because of their use in hotspots and as access mechanisms for visitors, users on a given WLAN may need to be given different privileges. The 802.11 family of standards includes no specific support for VLANs, but many networks have been configured and successfully run with VLANs over 802.11. Closely related to VLANs is the need for unique billing solutions for users of hotspots.

WLANs will be a huge driver for improvements in mobile IP and traffic routing. Stations in a WLAN continuously scan for access points that have stronger signals to associate with. Even if the user is sitting still, changes in the environment may cause a station to roam from one access point to another. Service disruption under these conditions is disconcerting to the user. Any attempt to use 802.11 as a wide-area mobile network will tax traffic routing and mobile IP solutions as never before.

TABLE 1

Theoretical Maximum Throughput and System Capacity of Various 802.11 Standards

WLAN Mode	Maximum Link Rate	Maximum UDP Rate	Maximum TCP Rate	System Capacity	
				# of Channels	Maximum UDP Capacity
802.11a Turbo	108 Mbps	55.1 Mbps	42.7 Mbps	6	330.6 Mbps
802.11a	54 Mbps	30.7 Mbps	24.0 Mbps	13	399.1 Mbps
802.11g (11g-only)	54 Mbps	30.7 Mbps	24.0 Mbps	3	92.1 Mbps
802.11g (with 11b present and idle)	54 Mbps	19.6 Mbps	14.5 Mbps	3	58.8 Mbps
802.11g (with 11b present and active)	54 Mbps	11.2 Mbps	9.2 Mbps	3	33.6 Mbps
802.11b	11 Mbps	7.1 Mbps	6.1 Mbps	3	21.3 Mbps

There is growing interest in mesh networks in which WLAN access points forward traffic through the net rather than to a wired backbone. The 802.11 standard provides a basis for this by defining packet types that allow the forwarding of traffic from one AP to another. The solution built into 802.11, however, is limited to a single hop. People have been able to build on this to create complete mesh networking systems, but there is no standardization effort within 802.11 on this topic.

PERFORMANCE PARAMETERS

The four most important performance parameters of any wireless system are user throughput, total system capacity, range, and power dissipation. Table 1 shows a comparison of the throughput and system capacity for the various 802.11 standards. The 108-Mbps "turbo" mode is not officially standardized, but is provided by a number of vendors.

As in so many cases, users do not see the "advertised"

The Family Dynamics of 802.11

numbers. In discussing the data rates of the standards, most people quote the peak raw data rate on the airwaves. Any real data stream, however, has an effective throughput that is degraded by the time spacing between the packets and the time required to send acknowledgments. The net throughput is further degraded by adding higher-layer protocols such as TCP/IP. Even the user throughputs shown in this table can be optimistic because they assume 1,500-byte packets. Shorter packets are less efficient and would result in lower throughputs.

The case of 802.11g is particularly complicated. In a “pure” 802.11g environment, its performance is the same as 802.11a. If any legacy 802.11b devices are present, however, 802.11g must protect its OFDM packets from being stepped on by the 802.11b devices that do not understand OFDM. To preserve the listen-before-talk collision avoidance system, the device sending OFDM packets must precede them by a short 802.11b packet announcing that it will be using the channel. This message is called a clear to send (CTS) packet. Sending a separate packet using the old slow modulation system significantly

degrades the performance of the 802.11g devices.

The situation becomes worse when 802.11b devices are transmitting data in the environment. The MAC protocol gives these devices a nearly equal chance of transmitting at any given time. Because the channel is shared among all the devices in time, the long slow packets of the 802.11b devices reduce the total throughput that can pass through the channel. The values shown in Table 1 assume one 802.11g station and one 802.11b station, each trying to transmit as much as possible. The throughput shown is the total net throughput for both devices.

Within a physical area, each cell (called a Basic Service Set, or BSS, in 802.11) is operated on a different frequency channel. If two BSSs are on the same frequency channel and are physically close, they will interfere with each other. The result is that the throughput of a single BSS will be shared between the two overlapping BSSs. Therefore, there is a limit to the total system capacity provided in a given physical space. This limit can be approximated by multiplying the throughput of a single channel by the number of different frequency channels available. As shown in Table 1, the overall system capacity of 802.11a can be far higher than 802.11g or b, as much more spectrum is available at 5 GHz. Therefore, only 802.11a is likely to have sufficient system capacity to “unwire” large office buildings.

A number of theories have been put forth about the relative ranges of 802.11a, b, and g. Many have speculated that 2.4 GHz will propagate farther than 5 GHz. Propagation measurements, however, show that any difference is small [see “Indoor Propagation for Wireless LANs,” by D. Dobkin, *RF Design*, September 2002]. In

TABLE 2 System-level Power Consumption of Various WLAN Cards

Chipset Vendor	WLAN Mode	Idle (Sleep)	TCP Uplink (Tx)	TCP Downlink (Rx)	Metric
Atheros	802.11a	0.75 W	2.48 J/MB	2.56 J/MB	3.16 J/MB
Cisco	802.11a	1.04 W	2.48 J/MB	2.88 J/MB	3.69 J/MB
Intersil	802.11a	2.06 W	2.40 J/MB	2.48 J/MB	4.09 J/MB
Atheros	802.11g	0.92 W	3.20 J/MB	3.20 J/MB	4.02 J/MB
Broadcom	802.11g	1.78 W	2.64 J/MB	3.28 J/MB	4.79 J/MB
Intersil	802.11g	1.31 W	2.48 J/MB	3.04 J/MB	4.12 J/MB
Agere	802.11b	0.14 W	15.28 J/MB	17.28 J/MB	17.14 J/MB
Atheros	802.11b	0.87 W	8.40 J/MB	8.08 J/MB	10.94 J/MB
Intersil	802.11b	0.46 W	16.40 J/MB	14.88 J/MB	17.28 J/MB

Not intended to be an exhaustive survey of WLAN cards.

practice, the differences in the quality of the radio design and digital signal processing far outweigh any differences in propagation.

Figure 4 shows throughput vs. range measurements taken in a typical office environment with a mixture of cubicles and closed offices. As expected, the throughput falls off with distance between the transmitter and receiver. This relationship is not monotonic, however, because of specific obstacles in the environment, such as concrete walls or large metal shelves. The figure shows that, based on the sampling of commercially available equipment, 5-GHz 802.11a actually has the greatest range of any of the technologies.

Similarly, many people have claimed that the high data rates and high carrier frequencies of 802.11a would make it power hungry. In practice, power consumption has to be viewed from the system level and needs to be normalized to the amount of information that is transferred.

Measurements show that in a typical laptop, the WLAN card itself accounts for approximately 10 percent of the total power draw; the added processing and activity on the laptop's processor resulting from the wireless communication consumes about 20 percent of the total power draw; and the remaining 70 percent of the power draw is for the screen and other functions completely unrelated to the wireless connection.

Commercially available WLAN cards have similar power consumption—for the cards themselves. However, the driver software, and the way data is moved from the card to the host, can have a dramatic effect on the overall power consumption. Some cards use programmed I/O to move the data, which requires a high degree of activity of the processor, and, therefore, high system power consumption. Other cards use direct memory access (DMA), moving the data into the host's memory with little or no processor activity.

WLAN cards are all able to “sleep” in a very low-power dissipation state when they are not sending or receiving data. For a given amount of data to be transmitted, cards that can transmit data at a higher rate can spend a much greater percentage of the time in sleep mode. Therefore, a fair comparison of power consumption should describe the energy consumed per megabyte transferred, rather than just the power consumed while actively transmitting or receiving. It is interesting to note the following relationship:

$$\text{Energy/ MB} = (\text{Joules per second})/(\text{Mbps}) = \text{power dissipation (in Watts)/data rate (in Mbps)}$$

Table 2 summarizes the measured added system-level power consumption from the wireless link for a variety of commercially available cards. The power consumption was measured by actually monitoring the power drain from the laptop's battery with and without the wireless cards present. The power consumption has been normalized to the measured TCP/IP throughput for each card. We created a “real world” metric that represents a wireless link that spends 70 percent of its time idle (sleeping), 20 percent of its time receiving, and 10 percent of its time transmitting. This breakdown was found to be typical in actual usage tests.

FUTURE CHALLENGES

At first, 802.11 WLANs appeared as Ethernet links to software applications. Virtually any application that runs over standard wired networks will run unaltered over WLAN connections. However, several second-order effects should be considered.

The available bandwidths of these links can vary tremendously. A single 802.11b access point shared among many users in a hotspot will result in performance not much better than a telephone modem. The same access point when used by few users is commensurate with broadband access such as DSL or cable modem. Throughputs achieved by properly deployed 802.11a networks can approximate wired Ethernet networks. The key is that workplace applications should not assume the huge bandwidths provided by switched 100BaseT, or Fast Ethernet, networks. Although this may seem like a step backward, the value of mobility outweighs the loss in data rate for most users.

The reduced capacity of wireless networks relative to switched 100BaseT networks will also increase the importance of Quality of Service techniques in making the best use of the available capacity. Although wireless networks will have QoS capabilities as described previously, the networks cannot apply these techniques if the applications and wired network backbones do not assign and maintain traffic classifications and priorities. The use of 802.1p/q prioritization and other standardized methods of traffic classification will therefore become more important, just to ensure that the WLAN endpoints will be able to allocate their available capacities effectively.

Capacity limitations can be minimized through proper planning and deployment of WLAN access points. Current tools for planning WLAN installations, frequency planning (assigning frequency channels to each of the APs), and adjusting transmit powers to minimize interfer-

The Family Dynamics of 802.11

ence (power control) are fairly crude. Although 802.11a with its many channels can relieve this issue, demand will drive the development of better planning and deployment tools.

Wireless networking also differs from wired networking in the degree of mobility that the upper-layer protocols and software must accommodate. Significant periods of service interruption are common. In particular, including the necessary security key exchanges, service can be interrupted by up to 50 milliseconds each time a user is handed off from one AP to another. Therefore, applications must be tolerant of occasional significant service interruptions.

Fortunately, most applications are designed for such disruptions resulting from the behavior of traffic over the Internet. Mobile IP and other methods for mobility will increase in importance. Implementing roaming from wide-area networks such as 3G cellular systems to WLANs poses a particularly difficult problem.

Another interesting trend is the increasing number of devices that will be served by the Dynamic Host Configuration Protocol (DHCP). Complicating this will be the eventual movement of servers and even switches into the wireless fabric, at which time these devices will most likely use DHCP as well. Working in a world in which IP addresses are not static will become a greater requirement for software applications over time.

Security is another area in which wireless LANs behave differently from wired LANs. Because WLAN signals can be sniffed from great distances, the need to encrypt sensitive information is critical. Application developers need to keep these two points in mind: First, many users who buy WLANs do not enable the security features; second, in many cases the WLAN needs to be open to provide access for visitors. The use of VLANs and virtual private networks (VPNs) is highly recommended in workplaces that will be using WLAN technology. Unfortunately, such systems are complicated to deploy and manage. The wireless revolution will drive improvements, but in the

interim many users will accept the less flexible link-layer security provided within 802.11

A PROMISING WIRELESS FUTURE

The 802.11a standard is just beginning to be deployed widely and will significantly enhance the WLAN space by providing the highest throughput, system capacity, and range, while consuming the lowest energy per bit transferred of any member of the 802.11 family of standards. New modulation techniques such as antenna beam forming and MIMO will increase data rates beyond 100 Mbps and increase ranges to several kilometers.

Enhancements to the 802.11 standards will provide wireless QoS, enabling high-quality voice, video, and audio. All of these improvements will create opportunities for WLAN beyond today's corporate and home data networks. With these opportunities, however, come new challenges. In particular, mobility will demand new methods for traffic routing and mobile IP; ubiquity will require better security and VLAN capabilities; and the desire for wireless multimedia services will stimulate the need for QoS throughput in the wired and wireless network. Q

BILL MCFARLAND is director of algorithms and architecture at Atheros Communications, a fabless semiconductor company making chipsets for WLAN products. He joined Atheros in 1999 and manages a team developing digital signal processing algorithms, defines digital and analog radio architectures, and represents Atheros in regulatory and standardization efforts. Prior to joining Atheros, McFarland spent 14 years at Hewlett-Packard Labs working on high-speed digital test equipment and fiber optic communications links, as well as serving as manager of the wireless circuits research group. He has published more than 25 papers, holds eight patents, and earned a bachelor's degree from Stanford University and a master's degree from the University of California at Berkeley.

MICHAEL WONG is a systems engineer at Atheros Communications. Prior to that he was a graduate student in electrical engineering at Stanford University. As a Stanford Graduate Fellow, he specialized in wireless networking and architectures. He was awarded his bachelor's degree from the University of Toronto in 2001 and master's from Stanford in 2002.

© 2003 ACM 1542-7730/03/0500 \$5.00

BIO