

Network Computing

MAY 27, 2004 | WWW.NWC.COM *For IT By IT*

WORKSHOP WIRELESS

Wi-Fi 101

By Dave Molta

In the fall of 1999, the Wireless Ethernet Compatibility Alliance was born, introducing the Wi-Fi brand name at a media event. When a reporter objected to the term *wireless Ethernet*, he was basically told “It’s close enough.” After all, they both use a contention-based access scheme and share other characteristics at the MAC layer.

But the similarities end there. Although Wi-Fi and Ethernet share some common framing and field characteristics at Layer 2, they differ dramatically at the physical layer. Whereas Ethernet basks in the warm and sunny confines of a guided, protected medium—an unshielded twisted pair or fiber, for instance—Wi-Fi operates in the haze of the public airwaves, where the equivalents of snow squalls and thunderstorms often wreak havoc on reception.

Wireless vendors have tried to hide the vagaries of Wi-Fi physical-layer signaling with promises of easy installation and seamless integration. But you need RF know-how to manage a large wireless network, just as you must know structured wiring to efficiently manage your Ethernet LAN.

ABCs of RF Communication

Like dial-up and cable modems, Wi-Fi networks use a technique called modulation to convert digital signals from your computer into analog RF signals. The speed at which data can be transmitted over a modulated carrier depends on a number of factors, including available bandwidth and

the specific type of modulation used. Complex modulation schemes, such as 64-Quadrature Amplitude Modulation on 54-Mbps 802.11 WLANs, carry more bits per unit of time than simpler schemes, like the Differential Binary Phase Shift Keying used on 1-Mbps WLANs. If a complex modulation scheme isn’t supported by high-quality RF signals, errors will occur.

Because signal quality decreases over the RF medium, there’s always a trade-off between speed and distance. Radio waves that travel through the air attenuate faster than RF signals carried by cable modems running over a hybrid fiber-coax cabling system.

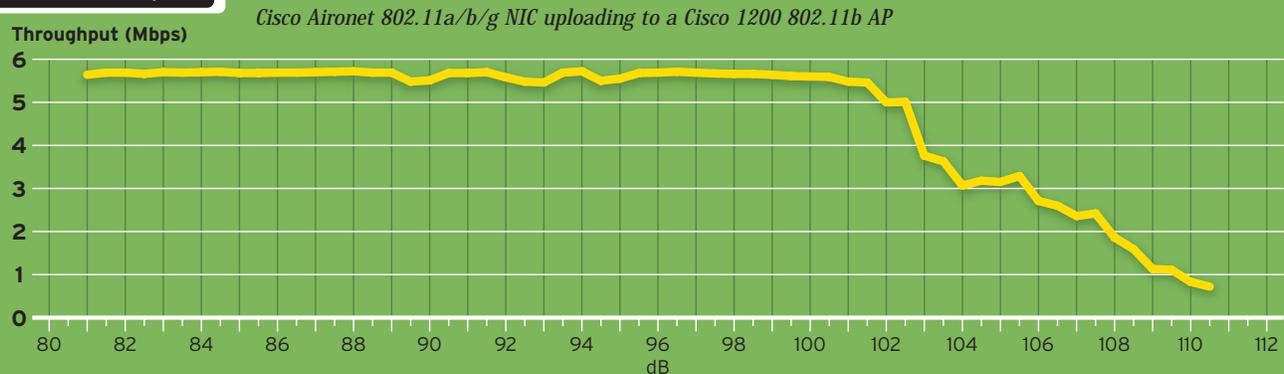
Wi-Fi networks use radio spectrum designated by the FCC and other regulatory bodies for unlicensed operators. This “buy, install and run” approach is one thing that makes Wi-Fi so appealing. Although you don’t need a license to operate a Wi-Fi system, vendors’ products must be certified to ensure they adhere to FCC rules.

FCC regulations govern the use of 83.5 MHz of spectrum between 2.4 and 2.4835 GHz, known as the 2.4 GHz ISM (Industrial, Scientific and Medical) band, as well as the 300 MHz of spectrum in the 5 GHz UNII (Unlicensed National Information Infrastructure) bands. The 5 GHz UNII bands include 200 MHz between 5.15 and 5.35 GHz and 100 MHz between 5.725 and 5.825 GHz.

The spectrum bands are further divided into channels, typically 22 MHz wide for Wi-Fi systems.

The FCC certification process restricts the freedom of radio-system designers, and for good reason: Because shared, unlicensed bands are at issue, products must be designed to minimize interference. Wi-Fi systems operating in unlicensed bands must radiate signals at very low

Rate vs. Range



WORKSHOP

WIRELESS

power levels, yet be robust enough to function amid modest levels of interference from other unlicensed systems.

The underlying spread-spectrum signaling method is the radio magic that allows systems to co-exist. Spread-spectrum systems are robust, but beware: It's almost impossible to design a multicell enterprise spread-spectrum Wi-Fi system without encountering interference-related performance problems. You are well-advised to understand these issues, which will help you design and support your wireless LAN.

Each Wi-Fi device, whether it's a PC card, a NIC or an AP (access point), acts as a transceiver that transmits and receives radio signals. Since all Wi-Fi systems use high-frequency microwave signals, the signals attenuate rapidly. Higher-frequency 5-GHz 802.11a signals experience somewhat greater attenuation than 2.4-GHz systems, especially when solid objects like floors and walls are within the coverage area.

Aside from contending with the normal signal attenuation, receivers must deal with ambient noise sources. That high-speed CPU in your notebook computer, for example, is a source of RF noise, as is the microwave oven down the hall. But modern radios can keep functioning even at very low signal-to-noise levels.

What's Watt

RF output power is often measured in watts. While a surround-sound stereo system could have total output power of 500 watts, for instance, microwave-based RF data systems operate at much lower signal levels. Even the most powerful Wi-Fi radios come with output power

of only 200 milliwatts, or one-fifth of a watt. Because wireless systems operate at low power, engineers use a logarithmic transformation, known as a decibel, to express signal levels. When decibels are used to measure power in relation to a milliwatt, we use the abbreviation "dBm." A signal level of 0 dBm equals 1 mW.

Signal levels that fall below 1 mW are measured in negative numbers. For example, a cal 802.11b Wi-Fi NIC might have a minimum receive sensitivity of -90 dBm to support a data rate of 2 Mbps.

There are two rules of thumb for decibels. First, when you increase a signal by 3 dB, you're actually doubling the signal. Likewise, decreasing a signal by 3 dB cuts the power output in half. Second, when you increase the signal level by 10 dB, you increase the power output tenfold. So if 0 dBm is equivalent to 1 mW, then 10 dBm equals 10 mW, 20 dBm equals 100 mW and 30 dBm equals 1,000 mW, or 1 watt. By combining the two rules, you also can see that 23 dBm is equal to 200 mW. Got it?

Your Gain, My Loss

Radio transmitters include amplifiers for generating stronger signals, which result in RF gain. RF technology designers can improve gain by increasing the power input within regulatory parameters. But that means greater energy consumption (read: shorter battery life) and heat dissipation (hot laptop).

Antennas can extend the range of coverage and improve performance by focusing the signal in a specific direction. On the receiving end, high-gain antennas may enhance range and performance.

How RF Modulation Rates

The speed of a digital data network—its bit rate—is measured in bits per second, or bps. Another measure with analog signaling is the baud rate, the number of distinct "symbols" transmitted each second.

The simplest symbols used in RF systems have two possible states. Thus, one bit of information (1 or 0) can be transmitted with each symbol, and the bit rate is the same as the baud rate.

RF modulation techniques can use three unique radio-signal properties—amplitude, phase and frequency—to encode digital information. For example, a 300-baud BPSK (Binary Phase Shift Keying) modulation scheme measures the phase of a signal 300 times per second. A normal radio signal could be equated with a digital 1, while a signal

that is 180 degrees out of phase could be equated with a 0.

Complex modulation schemes combine multiple characteristics of a radio signal to increase the bit rate. For example, QPSK (Quadrature Phase Shift Keying) supports symbols with four different states and can therefore carry 2 bits of information (00, 01, 10, 11). QAM (Quadrature Amplitude Modulation) increases the bit rate even further by combining unique combinations of amplitude and phase. For example, QAM 16 supports 16 unique states per symbol. Translated to binary, a QAM 16 signal carries 4 bits per symbol. Each time you double the number of unique symbol states, you increase the number of bits that can be carried by 1.

Although 256-QAM modulation

is in use on wired RF systems today, the best today's Wi-Fi systems (802.11a and 802.11g) achieve is 64-QAM, which supports 6 bits per symbol.

Amplifiers and receivers must be well-designed to decode the digital information contained in this relatively complex symbol architecture. And the signal-to-noise ratio must be quite high. Because of these requirements and the tendency of RF signals to attenuate over distance, the highest supported data rates usually require the transmitter and receiver to be in close physical proximity.

On 11a and 11g systems, QAM-16 is used at data rates of 24 and 36 Mbps, QPSK is used at 12 and 18 Mbps, and BPSK is used at 6 and 9 Mbps.

For a wireless network to function properly, it needs enough total-system gain between two nodes to overcome the effects of path loss. Otherwise, you'll experience the data equivalent of your favorite radio station breaking up as you drive farther away from home. Attenuation in the open air is called free-space loss—a slow fading as signals disperse in the atmosphere. Within buildings, WLANs encounter other types of loss, including absorption (walls, floors and doors), scattering (chaotic reflections off irregular surfaces), diffraction (bending around objects like a turn in a hallway) and refraction (bending of a wave as it passes through an object, like a glass wall). These characteristics differ based on frequency. For example, 5-GHz systems experience slightly greater attenuation due to absorption by walls and floors than do 2.4-GHz systems.

Although gain and loss often determine whether you can connect to the WLAN or not, another factor influencing

A wireless network needs enough **total-system gain** to overcome **path loss**.

radio waves inside buildings is multipath, a form of interference that occurs when radio waves bounce off solid surfaces. The result: multiple instances of a signal arriving at a receiver at slightly different times.

Engineers are constantly looking for ways to design systems that overcome multipath interference. Most of today's Wi-Fi systems, for example, come with dual-diversity antennas that can sometimes help. In most cases, it's the design of the radio receiver that determines its robustness in dealing with multipath. That's why a wireless NIC with high output power can experience poorer rate-versus-range performance than a different vendor's NIC that has lower output power but enhanced multipath capabilities.

Put to the Test

The **mathematical difference** between the output power of the transmitting device and the sensitivity of the receiving device is known as the path loss, or link budget. A Cisco Aironet 802.11b NIC, for example, has a maximum output power of 20 dBm, whereas a Cisco 1200 AP has a receive sensitivity of -85 dBm (remember, negative dBm numbers represent fractions of a milliwatt) at the 11-Mbps data rate. That means the total link budget is 105 dB.

As RF signal quality degrades, Wi-Fi systems experience diminished performance due to errors and retries. To compensate, Wi-Fi systems automatically drop their data rates as signal levels decrease. More precisely, they change to a less efficient modulation scheme, which in turn reduces the data rate. So for 802.11b systems, data rates drop from 11 Mbps to 5.5 Mbps to 2 Mbps and, finally, to 1 Mbps. When the signal level is insufficient to support the 1-Mbps data rate, the connection is lost.

In our Syracuse University Real-World Labs®, we use test gear from Azimuth Systems to perform rate-versus-range testing by programmatically adjusting the attenuation level in an RF isolation environment. The graph on page 1 shows rate versus range for a Cisco Aironet 802.11a/b/g NIC configured to constantly upload data to a Cisco 1200 802.11b AP. Overhead limits the throughput of Wi-Fi systems to a little more than 50 percent of the data rate, so the Cisco NIC starts at about 5.7 Mbps of throughput with 80 dB of attenuation.

The throughput remains relatively steady to about 100 dB of attenuation. Then, signal degradation begins to cause errors and retries at the 11-Mbps data rate, decreasing throughput. At around 103 dB of attenuation, the NIC shifts its data rate to 5.5 Mbps. It doesn't take long before performance starts to degrade further, and another rate shift occurs at around 109 dB. At 111 dB, the signal level is insufficient to support even the 1-Mbps data rate, and the network connection (802.11 association) between NIC and AP gets dropped.

The Azimuth test system provides the most systematic analysis of client rate-versus-range performance. We expect test labs and equipment manufacturers to deploy it widely. In the few months we've worked with this system, we've seen significant differences in rate-versus-range performance of various network cards.

However, while the Azimuth Systems test environment is effective for getting accurate, controlled and reproducible performance test data on the most important RF behavior—attenuation—it doesn't simulate the effects of multipath interference.

With Ethernet, you can use cable scanners to certify UTP and fiber-cabling systems. If your systems are installed and terminated properly, these certifications are usually routine, and the system rarely changes once it's certified. Wireless, however, is more fluid, with the medium constantly changing as doors open and people move around. Understanding how RF systems work will help you get more out of popular RF site-survey and troubleshooting tools, enhancing both your planning and support of Wi-Fi networks. [NWC](#)

Dave Molta is a NETWORK COMPUTING senior technology editor. He is also assistant dean for technology at the School of Information Studies and director of the Center for Emerging Network Technologies at Syracuse University. Write to him at dmolta@nwc.com.

Reprinted for



www.azimuthsystems.com