The 802.11n standard: GROWN UP AT LAST

THE 802.11N WI-FI STANDARD PROMISES MUCH. AFTER AN ESPECIALLY ANGST-RIDDEN DEVELOPMENT PROCESS, IT'S READY TO GO TO WORK.

BY ANN R THRYFT • CONTRIBUTING TECHNICAL EDITOR

Even for a wireless-communications standard, the 802.11n specification has over the last several years been subject to an exceptionally messy development process. Battling vendors and standards proposals, predraft and postdraft 1.0 silicon and box products, and multiple interoperability issues have continued to plague this latest generation of Wi-Fi technology. The 802.11n spec is now mostly past the difficult stage and about to fulfill its promise: a range approximately twice as great and transmitting speeds five to 10 times as fast as those of legacy 802.11a/b/g products (**Table 1**). It's about time, too, because the spec has a big job to do in designs for consumer, enterprise, campus, and metroscale markets. Unfortunately, the spec is complex and has a huge number of possible variants and options.

Those complexities are at least one reason that the process has taken so long. According to some industry participants, vendor infighting has been another major reason that the approval process has taken so long. Even before the IEEE committee approved Draft 1.0, at least three vendors had offered different proposals for the standard's core technology. TGn (Task Group N), the IEEE group that handles the spec, eventually narrowed down these competing proposals to two, says Jagdish Rebello, director and principal analyst of wireless communications for iSuppli. "Last year, [the two proposals] merged, and TGn submitted a joint proposal to the full body in May 2006." But TGn failed to achieve the necessary 75% approval level for a Draft 1.0 version. In March 2007, the committee voted to approve Draft 2.0, which came with approximately 3000 technical and editorial comments. Industry participants expect TGn to issue Draft 3.0 for a recirculation ballot this month, when the group expects to have completed the comment resolutions. As is usual for an IEEE standard, the specification will likely go through some more tweaking for another six to 12 months or so. The unprecedented number of options in this 802.11x spec makes this fine-tuning especially necessary, but participants agree that the mandatory sections are unlikely to change now.

During this lengthy, drawn-out process, at least one vendor coalition arose to develop an alternative specification so that products could more quickly get to market than the standards committee's work would support. In frustration at the length of time involved, manufacturers began releasing products even before Draft 1.0, based on different vendors' silicon designs. As tested by various external labs and research companies, many of these designs did not interoperate with each other.

To help push forward the standards-development process and provide interoperability certification, the WFA (Wi-Fi Alliance) in August 2006 announced that, during the first half of 2007, it would launch an extensive certification process for products that included baseline features from the developing standard. In May 2007, after TGn approved Draft 2.0, the WFA unveiled the Certified for 802.11n Draft 2.0 program, which it based on the mandatory sections of Draft 2.0, and announced the

first certified chip, card, and box products, which form the testbed for certifying additional products. The formal certification program began in June 2007. Many vendors have stated that their WFA-certified 802.11n Draft 2.0 products will be firmware-upgradable to the final IEEE 802.11n standard.

Netgear, whose products have now passed WFA certification, in 2005 launched Draft n products that it based on chips from two suppliers. "At that point, interoperability across chip vendors was an issue, and we had to supply two lines of adapters," says Som Pal Choudhury, Netgear's product-line manager for advanced wireless. The company's latest router product automatically detects and self-installs the most recent firmware from the company's Web site.

802.11N ADVANTAGES

The 802.11n WLAN (wireless-local-area-network) technology is the only Wi-Fi technology today with the bandwidth to support multiple HDTV (high-definition-TV) streams at 20 Mbps each. This performance is sufficient for implementing some long-standing goals of Wi-Fi networks. One of these goals, wireless multimedia, comprises voice/VOIP (voice over Internet Protocol), data, video, and gaming in residential applications. Another goal is achieving throughput, QOS (quality-of-service), and security levels that compare favorably with those of Ethernet which are necessary for enterprise-grade, campus, and municipal networks. But the methods of achieving this performance are complex, leading to many of the options and variants in the standard. Another main reason for all the options is the large number of device types that users want to connect to Wi-Fi networks, each with its own distinct set of requirements. Because the market for Wi-Fi has become much more heterogeneous than it was in the early days of 802.11x, 802.11n networks must accommodate a much wider range of device types; many of the standard's optional-requirement portions reflect that range. New requirements from consumer-electronics companies, such as video applications, or from the handset market, in which manufacturers are interested in power savings and better coverage, have contributed to the long, drawn-out IEEE process, as well as to the standard's complications, says Frank Hanzlik, managing director of the WFA. "There are a lot more people to please, so the compromise process has been more complex," he says.

But now that WLAN capability has penetrated the consumer and communications-device markets and is increasingly embedded into DSL (digital-subscriber-line) and cable modems, as well as into Apple TV, the markets are bigger, and the stakes are much higher, increasing motivation and helping to drive resolution of those conflicts. As a result, 802.11n Draft 2.0 has turned into more of a framework than a standard, says Craig Mathias, principal of the Farpoint Group. How difficult it will be for engineers to navigate that framework remains to be seen.

AT A GLANCE

- The IEEE's Task Group N has approved 802.11n Draft 2.0, the latest Wi-Fi standard, which promises approximately twice the range and five to 10 times the transmitting speeds of legacy 802.11a/b/g. Mandatory features probably won't change when the group issues Draft 3.0 for recirculation this month.
- Because of the large number of options possible under the standard—including RF band, channel width, number of antennas, and modulation scheme—not to mention additional options for specific device types, designing 802.11n products and determining their interoperability can be tricky.
- The Wi-Fi Alliance has begun interoperability testing of products under its Certified for 802.11n Draft 2.0 program, based on the mandatory features of Draft 2.0.
- Many vendors have said that their WFA-certified 802.11n Draft 2.0 products will be firmware-upgradable to the final IEEE 802.11n standard.

WHAT'S IN THE STANDARD?

The fact that products based on 802.11n can operate in either the 2.4- or the 5-GHz bands or both makes them potentially backward-compatible with legacy products. When you use 802.11n in the 2.4-GHz band and in a 20-MHz channel,it is backward-compatible either with 802.11b, using CCK (complementary-code-keying) modulation or with 802.11g, using OFDM (orthogonal-frequency-division-muliplexing) modulation. When you use 802.11n at 20- or 40-MHz channel widths in the 5-GHz band with OFDM modulation, it is backward-compatible with 802.11a. Devices built to conform to more recent 802.11x standards tend to employ OFDM because it is more efficient than CCK, which older Wi-Fi networks use. With 802.11g, OFDM has become the Wi-Fi-modulation scheme of choice. The 802.11n standard introduces more efficient OFDM modulation to increase data rate. It uses 52 data subcarriers instead of the 48 in legacy networks, producing 65 Mbps per spatial stream instead of the 54 Mbps of legacy 802.11a or g. Another option in the standard shortens the guard interval from 800 to 400 nsec to increase the OFDM symbol rate, further boosting the data rate.

Previous 802.11x standards specified only one frequency band, one channel width, one spatial stream—transmitting or receiving—per direction, and one maximum data rate. Aside from OFDM improvements to enhance throughput, the 802.11n spec also doubles the channel width and introduces frame aggregation, block acknowledgment, and spatial multiplexing; spatial multiplexing is one of several possible MIMO (multiple-input/multiple-output) configurations (see **sidebar** "What MIMO does"). It allows one or two channel widths, 20, 40, or both 20 and 40 MHz; one to four spatial streams in either direction; and at least two other MIMO options (**Table 2**). Transmitting-data rate for 802.11n networks is therefore highly variable and is based primarily on modulation scheme, channel width, and the number of spatial streams.

Depending on the design, an 802.11n-compliant product can reach a typical throughput of 144 Mbps, assuming OFDM modulation, two transmitting and two receiving streams—known as a 2×2 configuration—a 20-MHz channel width, or a—currently theoretical—maximum throughput of 600 Mbps, assuming OFDM modulation, a 4×4 configuration, and a 40-MHz channel width. Most 802.11n products operating today achieve transmitting speeds between these two extremes: 300 Mbps with OFDM, a 2×2 configuration, and a

40-MHz channel width, or 450 Mbps by simply changing the MIMO configuration to 3×3.

As always, range also depends on several variables. But, with 802.11n, it's also more complex because of all the possibilities, including the transmitting power, number of receiving antennas, modulation scheme, and error-correction scheme. The spec also allows operation in 2.4- or 5-GHz frequency bands or a third dual-band 2.4/5-GHz option.

The differences between Draft 1.0 and Draft 2.0 are fairly minor, according to Bill McFarland, Atheros' chief technology officer and a member of the TGn. The biggest areas of change are associated with the coexistence mechanisms that govern how 802.11n devices behave with 802.11g devices, especially in the 2.4-GHz band. This issue is less important for 802.11a devices in the 5-GHz band: There, coexistence is simpler because the band is less crowded,

WHAT MIMO DOES

You can define MIMO (multiple-input/multiple-output) air-interface-technology methods in more than one way. The 802.11n Draft 2.0 spec includes spatial (space-division) multiplexing as the mandatory method for implementing MIMO and space-time block coding and beam forming as two optional methods. All of these methods use multiple antennas at both the transmitter and the receiver to send multiple data streams, increasing transmission speed, range, robustness, or some combination of these features without increasing bandwidth. These methods take advantage of the multipath effects, or reflections, inherent in RF communications. These reflections usually result in signal distortion at the receiver which, in turn, degrades performance.

In spatial multiplexing, which is mandatory in 802.11n, two to four antennas, each carrying a transmitting/receiving stream, simultaneously transmit multiple bits in the same frequency channel to increase throughput (Figure A). Because spatial-multiplexing techniques make receivers much more complex, designers usually combine them with OFDM (orthogonal-frequency-division-multiplexing)-modulation schemes, which are more efficient than other modulation schemes. Space-time block coding, an 802.11n option, uses different antennas for redundancy to increase robustness. The other optional method is beam forming, which uses multiple antennas as if they were parts of an array, forming a directional antenna that directs a beam to increase range.

The 802.11n spec allows one to four data streams for either the transmitting or receiving directions, but MIMO requires a minimum of two for each direction. In applications in which higher data rates, accurate data reception, and longer distances are critical, such as when incorporating 802.11n into a residential gateway, designers might want to increase data streams to three receiving and three transmitting for maximum data throughput. Further, they should use the best coding schemes, says Jagdish Rebello, director and principal analyst of wireless communications for iSuppli. "But if you're looking at a consumer product, such as a handset, you can fall back to the lowest configuration possible, such as single input, multiple outputs." Most applications use a two-transmitter/two-receiver (2×2) configuration, which is the minimum configuration that the Wi-Fi Alliance's Certified for 802.11n Draft 2.0 certification program is testing.

Access-point devices usually need symmetrical data streams. Asymmetrical arrangements, with more receivers than transmitters, improve receiver performance by way of better data correlation and range, which are especially useful in client devices, such as laptops and handsets. Because each antenna adds a radio chain, power consumption and cost of the silicon increase in both the RF and the baseband portions.

"This [cost and power-consumption issue] has historically been a challenge," says Broadcom's Kevin Mukai, senior product-line manager for 802.11n products. "That's where CMOS integration comes in. We've been pushing down the power budget to get 802.11n functionality onto smaller form factors." The advantages of a 2×2 configuration in silicon, such as Broadcom's chip set, are fewer antennas and therefore lower cost. The 802.11n spec incorporates a mandatory power-save mode for handsets. As single-chip silicon for 802.11n begins to appear, perhaps within the next year, greater opportunities for scaling down power will emerge.

"The cost, power consumption, and size of 802.11n silicon are issues that we, as an industry, still need to work through," says Craig Mathias, principal of the Farpoint Group. "But we're good at that."

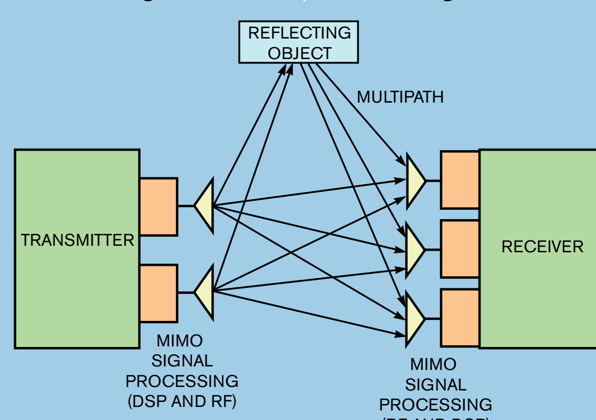

Figure A MIMO methods transmit multiple data streams to increase transmission speed, robustness, range, or some combination of these features (courtesy Farpoint Group).

TABLE 1 802.11N DRAFT 2.0 AND 802.11G COMPARISON

Feature	802.11g	802.11n Draft 2.0	Comments
RF band	2.4 GHz	2.4 GHz, 5 GHz	802.11n devices can be either single- or dual-band (2.4/5-GHz)-capable.
Channel width	20 MHz	20 MHz, 40 MHz	802.11n accommodates 20 MHz, 40 MHz, or both 20 and 40 MHz.
No. of transmitting or receiving spatial streams	One	One, two, three, or four	Common 802.11n transmitting and receiving configurations include 2×2, 2×3, 3×3, 3×4, and 4×4, but any combination of one to four streams per direction is possible.
Modulation schemes	Mostly OFDM; also backward-compatible with CCK and DSSS	Mostly OFDM; also backward-compatible with CCK and DSSS	OFDM encodes more bits per symbol than CCK based on the density of the QAM mode (maximum 64 points).
Typical transmitting data rate	25 Mbps with OFDM	144 Mbps with OFDM, 2×2, 20-MHz channel width	802.11n data rate depends on channel width, number of spatial streams, and modulation scheme.
Maximum transmitting data rate	54 Mbps with OFDM	600 Mbps with OFDM, 4×4, 40-MHz channel width	Current 802.11n equipment can transmit 300 Mbps with OFDM, 2×2, 40-MHz channel width, or 450 Mbps with OFDM, 3×3, 40-MHz channel width.
Typical indoor range	30 to 35m	50 to 70m	Range depends on multiple variables, including transmitter power, number of receiver antennas, modulation schemes, and error-correction schemes.
Typical outdoor range	110m	160m	Range depends on multiple variables, including transmitter power, number of receiver antennas, modulation schemes, and error-correction schemes.
802.11x backward compatibility	802.11b: 2.4-GHz band, CCK or DSSS, 20-MHz channel width	802.11b/g when you use 802.11n at a 20-MHz channel width in the 2.4-GHz band, CCK/OFDM (b/g); 802.11a when you use 802.11n at a 20- or 40-MHz channel width in the 5-GHz band, OFDM	802.11b maximum data rate is 11 Mbps. 802.11a maximum data rate is 54 Mbps.

Sources: Atheros, Broadcom, iSuppli, and Wi-Fi Alliance.

Notes: CCK=complementary-code keying. DSSS=direct-sequence spread spectrum. OFDM=orthogonal frequency-division multiplexing. QAM=quadrature-amplitude modulation.
2×2 denotes a two-transmitter/two-receiver configuration. 2×3 denotes a two-transmitter/three-receiver configuration. 3×3 denotes a three-transmitter/three-receiver configuration. 3×4 denotes a three-transmitter/four-receiver configuration. 4×4 denotes a four-transmitter/four-receiver configuration.

and most devices perform OFDM-style transmission and exceptions. Also, the layout and organization of the frequency channels work well with 802.11n. "In contrast, the 2.4-GHz band includes 802.11b devices using CCK modulation, and the frequency channels are laid out in a way that doesn't match up as well with how 802.11n likes to use the frequency," he says.

It's also difficult to get 802.11n devices to work well together with 802.11b and other devices in the 2.4-GHz band when using 40-MHz channels that double the data rate, because of potential interference issues. Techniques for dealing with this problem are still under discussion, and Draft 3.0 will probably resolve them, says McFarland. Some of these techniques are high-level-politeness algorithms based on measuring the amount of traffic in the environment: If there's a lot of traffic, channel width remains at 20 MHz. If traffic is low, channel width can expand to 40 MHz to take advantage of greater bandwidth. Another possible method calls for actively "listening" to detect any nearby access points or networks and moving to 20-MHz mode if any are detected, regardless of the amount of traffic. TGn is discussing modifications of this variant because it's potentially problematic.

Other, more fine-grained methods for preventing interference, such as CCA (clear-channel assessment), operate on a packet-by-packet basis. To prevent collision, before a 40-MHz-wide packet is transmitted, CCA checks to make sure that both channels are clear and that the packet can be transmitted on that whole 40-MHz frequency range. This option will probably remain. The spec now requires the use of both high-level-politeness and more fine-grained mechanisms. How the high-level algorithm will work is under discussion.

"At the MAC [media-access-control] layer, the concern is how devices share the airwaves on a packet-by-packet ba-

MORE AT EDN.COM

+ For more on the role of 802.11n in video distribution for home Wi-Fi networks, go to www.edn.com/article/CA6418210.

+ For another take on the 802.11n standards process and the problems that predraft and postdraft 1.0 products cause, go to www.edn.com/article/CA6405195.

TABLE 2 MAJOR 802.11N DRAFT 2.0 MANDATORY AND OPTIONAL FEATURES

802.11n Draft 2.0 specification feature type	802.11n Draft 2.0 specification feature	Tested by Wi-Fi Alliance Certified for 802.11n Draft 2.0 program	Explanation	Comments
Mandatory	MIMO spatial multiplexing	Support for two spatial streams in transmitting mode, two spatial streams in receiving mode (2×2); mandatory for Wi-Fi Alliance certification.	Transmitting mode required for access-point device, receiving mode required for access-point and client device, except for handheld devices	Simultaneously transmits multiple bits from different antennas, increases throughput
Optional	MIMO space-time block coding	Not tested	IEEE still finalizing details; candidate for future certification program	Uses different antennas for redundancy and increases robustness; targets handheld devices
Optional	MIMO beam forming	Not tested	IEEE still finalizing details; candidate for future certification program	Uses multiple antennas as an array to form a directional beam, increases range
Mandatory	Aggregated MAC Service/Protocol Data Units	Support for both; mandatory for Wi-Fi Alliance certification	Required for all devices	Improves MAC-layer efficiency by increasing maximum frame size (frame aggregation)
Mandatory	Block ACK	Support for block ACK; mandatory for Wi-Fi Alliance certification	Required for all devices	Improves protocol efficiency and throughput by reducing size of block-ACK frame and enabling back-to-back frame transmissions
Optional	2.4-GHz operation	Single-band device in 2.4-GHz band; tested if implemented	Devices must operate in one of three modes: single-band, 2.4-GHz; single-band, 5-GHz; or dual-band, 2.4- and 5-GHz.	802.11n devices can be either single- or dual-band (2.4/5-GHz)-capable. Most current 802.11n devices use 2.4-GHz band.
Optional	5-GHz operation	Single-band device in 5-GHz band; tested if implemented	Devices must operate in one of three modes: single-band, 2.4-GHz; single-band, 5-GHz; or dual-band, 2.4- and 5-GHz.	NA
Optional	Dual-band 2.4- and 5-GHz operation	Access-point or client-device capable of supporting both 2.4- and 5-GHz bands; tested if implemented	Devices must operate in one of three modes: single-band, 2.4-GHz; single-band, 5-GHz; or dual-band, 2.4- and 5-GHz.	NA
Mandatory	20-MHz channel width	Tested at 2.4 or 5 GHz for both one and two spatial streams	NA	802.11n accommodates 20 MHz, 40 MHz, or both 20 and 40 MHz.
Optional	40-MHz channel width	Tested only in 5-GHz band, if implemented, for one and two spatial streams; 40-MHz operation poses interference risks in crowded 2.4-GHz band.	IEEE has not yet decided protocol for allowing 40-MHz operation in the 2.4-GHz band.	Doubling legacy 20-MHz channel width can double throughput.
Mandatory	Mixed-mode operation: 802.11n with a, and 802.11n with b/g	Tested in both access-point and client devices	Mixed mode with b/g is mandatory for all 802.11n devices operating in the 2.4-GHz band; mixed mode with 802.11a is mandatory for all 802.11n devices operating in the 5-GHz band.	NA
Optional	Greenfield-mode operation: 802.11n only	Tested in both access-point and client devices if implemented	NA	Improves efficiency by eliminating support of legacy devices in an all-802.11n network
Mandatory	Protection protocols	Tested in both access-point and client devices	Mandatory when legacy networks and devices are within range. IEEE is discussing other optional methods.	Prevent negative impact on legacy-device traffic by 802.11n transmissions
Optional	Power-save multipoll	Not tested	NA	Saves handset power consumption in time-slotted network when 20 or more VOIP handsets connect to a single access point
Not in spec	NA	Optional untested features must not disrupt Wi-Fi Alliance-certified features and expected functions.	Required for all devices	NA

Sources: Atheros, Broadcom, iSuppli, and Wi-Fi Alliance.

Notes: ACK=acknowledgment. MAC=media-access control. MIMO=multiple input/multiple output. VOIP=voice over Internet Protocol.

sis," says McFarland. "The biggest enhancement here in Draft 2.0 versus 1.0 is packet aggregation." Traditionally, in 802.11x, a packet is transmitted and retried until an acknowledgment is returned. In 802.11n, mandatory packet aggregation combines a large number of packets into one superframe, sends it, and gets back a block acknowledgment specifying which packets were and were not received correctly. Only packets that failed must be retransmitted, resulting in a more efficient system and making better use of the high data rate. Additional differences in Draft 3.0 are likely to include changes to some of the spec's optional modes and features, including MIMO methods.

When a product implements one of the spec's optional features, the 802.11n spec provides for a negotiation associated with that feature so both devices can determine which options they include, says McFarland. But complications may arise in network configurations. For example, if the network requires beam forming, currently optional, both access points and clients must support it. Clients without that feature can still interoperate, but the network doesn't have the resulting performance enhancement.

For the first phase of its Certified for 802.11n Draft 2.0 program, the WFA defined a set of features that corresponds closely to most of the mandatory features of 802.11n Draft 2.0. The program also tests some of the spec's optional features if the device under test implements them.

In its second phase, currently targeting summer 2008, the program will test some optional features that the IEEE is considering for inclusion in the spec. The WFA has not yet decided how to label these capabilities. One possibility is a profile concept, or sets of features that correspond to the requirements of classes of devices, such as one set for PCs and data-centric devices; another for consumer devices that require video and audio streaming; and another that might include handheld devices requiring voice, VOIP, and telephony features. "This [arrangement] would give us better combinations of mandatory and optional features that make sense for that usage class," says the WFA's Hanzlik.

Available 802.11n chip sets differ in whether they target use in access points or client equipment and enterprise or consumer devices. They also have performance differences and varying MIMO configurations. Chip sets for 802.11n consist of CMOS baseband MAC engines, a separate CMOS radio, and an external discrete power amp.

Single-chip 802.11n silicon will likely become available within the next year or so. Legacy Wi-Fi a/b/g single-chip silicon already exists in CMOS SOCs (systems on chip), says Kevin Mukai, Broadcom's senior product-line manager for 802.11n products. Single-chip silicon will be more of a challenge for 802.11n products, because the a/b/g SOCs contain only one RF chain per chip, but 802.11n requires multiple RF chains. EDN

FOR MORE INFORMATION

Atheros
www.atheros.com

Broadcom
www.broadcom.com

Farpoint Group
craig@farpointgroup.com

iSuppli
www.isuppli.com

Netgear
www.netgear.com

Wi-Fi Alliance
wi-fi.org

REFERENCES

1 IEEE P802.11 Task Group N, Meeting Update, grouper.ieee.org/groups/802/11/Reports/tgn_update.htm.

2 "Wi-Fi Certified 802.11n draft 2.0: Longer-Range, Faster-Throughput, Multimedia-Grade Wi-Fi Networks," Wi-Fi Alliance, 2007, www.wi-fi.org/files/kc/WFA_802_11n_Industry_June07.pdf.

3 "Wi-Fi Certified makes it Wi-Fi: An Overview of the Wi-Fi Alliance Approach to Certification," Wi-Fi Alliance, September 2006, www.wi-fi.org/files/wp_16_Overview_WFA_Certification_09282006_2.pdf.

4 Wi-Fi Certified Products, Wi-Fi Alliance, http://certifications.wi-fi.org/wbcs_certified_products.php?lang=en.

5 "802.11n: Next-Generation Wireless LAN Technology," Broadcom, April 2006, www.broadcom.com/docs/WLAN/802_11n-WP100-R.pdf.

AUTHOR'S BIOGRAPHY

Contributing Technical Editor Ann R Thryft has been writing about technology, including wired and wireless networking, for more than 20 years. What she really wants is ESPnet: direct mind-to-screen/mind-to-mind communications, with LAN and WAN options for use at trade shows. You can reach her at athryft@earthlink.net.