CCPrep.com
A property of ARI

802.11n: A Feature Overview of The Latest WLAN Standard

By Steven Sowell

Introduction

I've been roaming the country lately doing a lot of WLAN consulting, design, diagnostics, and teaching. During this time, I've met many clients and students have questions about the new 802.11n draft standard. What follows is a brief, high-level overview of the key features and differences between the new 802.11n draft standard and the old 802.11a/b/g standards we've been using for years.

Standards Progress

In May 2008 in Jacksonville, Florida, the 802.11n Draft 4.0 passed recirculation ballot #124 by an 88% majority (75% required) with 253 votes to approve, 34 to not approve, 25 abstain.

The 802.11 Working Group is working through some technical issues on the 802.11n internal draft. This phase is expected to be complete by July 2008.
Next, the draft is sent for a "sponsor ballot." Based on comments received from the sponsor balloters, the 802.11 Working Group continues to update the draft until the sponsor balloters reach consensus that the draft is ready for final approval. The official 802.11 timeline for concluding this process and achieving IEEE Standards Board approval of 802.11n is March 2009.

Modulation Improvements

802.11a and 802.11g transmit using a technique called Orthogonal Frequency Division Multiplexing, or OFDM. This simply means that they divide their available spectrum into 48 sub-carriers, loosely analogous to 48 "mini radio stations" all transmitting at once. Each of these sub carriers transmits a portion of the data stream. This resulted in a maximum data rate of 54mbps.
802.11n increases the number of sub-carriers from 48 to 52, thereby bumping the data rate up to 65mbps for a single transmit radio. Since 802.11n defines up to four transmitters, we have possible data rates of 130mbps for two transmitters, 195mbps for three transmitters, and 260mbps for four transmitters.
Then, if we decide to bond to 20Mhz channels for 40Mhz of spectrum, we have data rates of 135, 270, 405, and 540mbps.

MIMO

Multiple Input Multiple Output, or MIMO exploits a radio-wave phenomenon called multipath: transmitted information bounces off walls, doors, and other objects, reaching the receiving antenna multiple times via different routes and at slightly different times. Uncontrolled, multipath distorts the original signal, making it more difficult to decipher and degrading Wi-Fi performance. MIMO harnesses multipath with a technique known as space-division multiplexing. The transmitting WLAN device actually splits a data stream into multiple parts, called spatial streams, and transmits each spatial stream through separate antennas to corresponding antennas on the receiving end. The current 802.11n draft provides for up to four spatial streams, even though compliant hardware is not required to support that many.

Doubling the number of spatial streams from one to two effectively doubles the raw data rate. There are trade-offs, however, such as increased power consumption and cost. The draft-n specification includes a MIMO power-save mode, which mitigates power consumption by using multiple paths only when communication would benefit from the additional performance. The MIMO power save mode is a required feature in the draft-n specification.

Beam Forming

Beam-forming is a technique that adjusts the phase relationship between to transmit antennae to optimize Signal-to-Noise Ratio (SNR) at the receiving host antenna, thereby improving range and performance. To understand how beam forming works, consider the popular noise-canceling headsets that many people like to wear. These headsets use a small microphone to detect the sound coming toward your ears. An equal but opposite sound is reproduced inside the headset, thus canceling out much of the original sound. This is an example of adding two *opposite* sound waves which cancel each other out.

Instead, we could add two *identical* sound waves together and they would build a very powerful sound. This is akin to beam forming, where two antennae transmit signals in just the right phase relationship that they add together as a stronger signal at the receiver. Beam forming requires that the receiver provide feedback to the transmitter about the signals it is receiving, and only 802.11n clients can provide this feedback.

This feedback is only relevant for the receiver’s current position and orientation relative to the transmitter. Further, this technique only works when transmitting to a single receiver. It doesn’t work when broadcasting or multicasting to multiple receivers. Since beacon broadcasts from the AP must be received by the client, beam forming does not significantly increase the coverage area for an AP. It can increase the throughput only to individual, relatively stationary clients. Finally, this technique is most effective in an area with minimal obstructions and reflective surfaces.

In my own testing in comparing 802.11n performance to that of 802.11a/g, I noticed beam forming in action. As I walked away from the AP, my performance would degrade about the same with 802.11n as with 802.11a/g. However, if I stopped walking and remained stationary for a short time, the performance of my client would improve slightly with 802.11n. At the time of my first 802.11n testing, I was not aware of this feature, and simply noted the phenomenon as something to investigate further….now I know! ☺

Spatial Diversity

Spatial diversity exploits the use of multiple antennas. Signal streams are transmitted by each antenna of the AP, which reflect off the various surfaces in the area, and hit the receiver at slightly different times. Using some heavy math, the receiver can combine these signals to achieve superior Signal-to-Noise ratios. The 802.11n specification supports up to four antennas, but most clients will have somewhere between one and three antennae.

Spatial Diversity could work in a so-called “two-by-one” (2x1) configuration, where the AP has two antennae, and the client has one (beam forming). Or you could have a 2x2 or 3x2 (three AP antennae, two client antennae), with ever-increasing performance benefits.

As with all things, the law of diminishing returns is applicable. There are significant gains as we progress up through a 3x2 configuration, but beyond that (3x3, 4x4) the gains are very small.

Channel Bonding

Another optional mode in the 802.11n draft effectively doubles data rates by doubling the width of a WLAN communications channel from 20 MHz to 40 MHz. We do this by bonding two adjacent channels together (1 & 2, or 5 & 6, or 10 & 11, etc).Since there is now twice the bandwidth available for transmissions, we can double effective throughputs.

Also, in standard 802.11b/g, the upper and lower portions of each channel are restricted to mitigate against co-channel interference. Since we are bonding the two channels to work together, we needn't worry about this interference, and we can use this previously reserved spectrum.

The trade-off is fewer channels available for other devices. In the case of the 2.4-GHz (802.11b/g) band, there is only enough room for three non-overlapping 20-MHz channels. Needless to say, a 40-MHz channel does not leave much room for other devices transmit in the same airspace. Thus, channel bonding is not recommended for 802.11b/g.

Since the 5Ghz, 802.11a band has at least 12 non-overlapping channels, channel bonding is a viable technique.

Aggregation

Aggregation improves efficiency in mixed-mode environments by allowing transmission burst of multiple data packets between overhead communications. Instead of a g client having to use slower "b-compatible" methods for gaining channel access for each frame, they can aggregate and transmit multiple frames with a single set of access/overhead frames. Instead of transmitting two separate data frames, each with their own Preamble, Radio header, MAC header, CRC, and acknowledgement overhead from the AP, we can transmit the data as a single frame, reducing overhead significantly. To help accommodate this feature maximum frame size is increased from 4k to 64k

Reduced Inter-frame spacing (RIFS)

A station must wait a minimum amount of time between frame transmissions. This is the Inter-Frame Spacing or IFS. 802.11n reduces this time, but only in a so-called "Greenfield" environment – one that does not require downward compatibility with 802.11a/b/g clients.

Reduced Guard Interval

Specifies a shorter "guard interval" delay between OFDM symbol transmissions than in 802.11a or g. Optionally, the standard 800 nanosecond delay can be reduced to 400 nanoseconds. This potentially increases the data rates to 72, 144, 216, and 288 Mbps for a 20- MHz channel and 150, 300, 450, and 600 Mbps for a 40-MHz channel.

Greenfield Mode

Improves efficiency by eliminating downward compatibility support for "legacy" 802.11a/b/g devices