Innovative Technology
for a Connected World

802.11n and Business-Critical Mobile Devices

Originally Published: March 2008

Updated: October 2012

A White Paper from Laird Technologies

The throughput of 802.11n is the result of enhancements that include packet aggregation, block acknowledgement, wider channels, decreased spacing between sent packets, and multiple input/multiple output (MIMO) technology. These enhancements not only boost throughput but also increase range, improve predictability of coverage, and improve quality of service.

Americas: +1-800-492-2320 Option 3
Europe: +44-1628-858-940
Hong Kong: +852-2268-6567 x026
www.lairdtech.com/wireless

CONTENTS

EXECUTIVE SUMMARY

With throughput much greater than that available with previous wireless local area networking (WLAN) standards, the IEEE 802.11n standard has had a significant impact on the WLAN, or Wi-Fi®, industry. Most of today's WLAN infrastructure products support 802.11n, and support on client devices is growing.

The throughput of 802.11n is the result of enhancements that include packet aggregation, block acknowledgement, wider channels, decreased spacing between sent packets, and multiple input/multiple output (MIMO) technology. These enhancements not only boost throughput but also increase range, improve predictability of coverage, and improve quality of service.

To achieve a significant throughput boost from 802.11n, both sides of the Wi-Fi link – the client and the infrastructure – must support 802.11n with more than one data stream. When the infrastructure supports 802.11n but clients do not, those clients still receive the non-throughput benefits of 802.11n. With 802.11n infrastructure devices rivaling pre-802.11n infrastructure devices in price, organizations should give 802.11n serious consideration when deploying a new infrastructure or refreshing an existing one.

Putting 802.11n on laptops and other general-purpose client devices makes sense if those devices need a throughput boost. Most business-critical mobile devices such as mobile computers don't need higher throughput. Because those devices may be used for five years or longer, consider dual-band 802.11n for them now if you can get it for a modest price premium.

INTRODUCTION

Since it was first proposed, the IEEE (Institute of Electrical and Electronics Engineers) 802.11n WLAN standard has generated a lot of excitement. Even before the standard was ratified in 2009, products that support 802.11n were in high demand. 802.11n unifies the preceding 802.11b, 802.11g, and 802.11a standards and provides significant enhancements to previous standards.

Before any 802.11 standard was defined, WLANs were deployed in "non-carpeted" vertical-market environments such as factories, warehouses, retail locations, and healthcare facilities. With standardization and performance improvements, WLANs moved into the mainstream. Today, WLAN radios are standard equipment on many client devices such as notebook computers and mobile phones, and WLANs are found in many "carpeted" environments such as homes and offices as well as public places such as coffee shops, hotels, and airports. The mainstream market for WLANs has become its dominant market.

The differences between mainstream areas and vertical-market areas go beyond floor covering. Applications, devices, usage scenarios, and requirements differ significantly between the areas. What may be a major benefit in a carpeted area may be of little or no use in a non-carpeted area.

The vast majority of business-critical mobile devices, such as mobile computers, operate in non-carpeted areas. The following sections of this paper analyze the value of 802.11n for users of such devices:

- Section 1: The main benefits of 802.11n
- Section 2: How 802.11n works
- Section 3: What to consider when deploying 802.11n in your infrastructure
- Section 4: The applicability of 802.11n features to business-critical mobile devices

While the authors have drawn on a number of sources while researching this topic, the opinions and conclusions provided in this document are the sole responsibility of Summit Data Communications. Sources are attributed at the end of this document.

SECTION 1: 802.11N BENEFITS

Performance

Since the first IEEE standard for WLANs was ratified in 1997, WLAN technology has evolved rapidly. Major advancements in the areas of security, management, and quality of service (QoS) have been matched by improvements in performance.

Note: For an explanation of data rate and throughput (two performance metrics), see Appendix 1.

The first IEEE 802.11 standard provided for a maximum data rate of two megabits per second (2 Mbps). 802.11b increased that data rate to 11 Mbps, and both 802.11a and 802.11g increased it further to 54 Mbps. 802.11n provides for a maximum data rate of up to 600 Mbps, which is faster than the 100 Mbps Fast Ethernet connections that today are common to the desktop.

802.11n supports bandwidth-intensive applications such as high-definition streaming video while still delivering the convenience, flexibility, and cost savings associated with wireless networking. As a result, experts within the industry are for the first time suggesting that WLANs can become the primary means of network access for typical users, freeing them from the costs and limitations that are inherent in wired LANs[1].

Coverage

When supporting one or more of the WLAN standards that preceded 802.11n, a Wi-Fi device has a single radio and a single antenna, and it transmits and receives a single data stream. With 802.11n, a Wi-Fi device with multiple antennas can transmit multiple data streams and receive multiple data streams.

When communicating with a client device that supports only one stream, an 802.11n infrastructure device such as an access point (AP) can transmit the same data multiple times and receive the same data multiple times. This technique improves the reliability of coverage for the client because it improves the likelihood that data will reach its destination without the need to retry the transmission.

By sending out multiple instances of a single transmission, an 802.11n AP "fills in" areas that may be missed with a single transmission. These otherwise un-covered areas, called *nulls*, can be common in environments with metal objects that reflects RF transmissions and other objects (such as wood) that absorb RF transmissions. By providing multiple receivers for a single transmitted signal, an 802.11n AP also reduces nulls and improves the likelihood that it will receive transmissions from a client device.

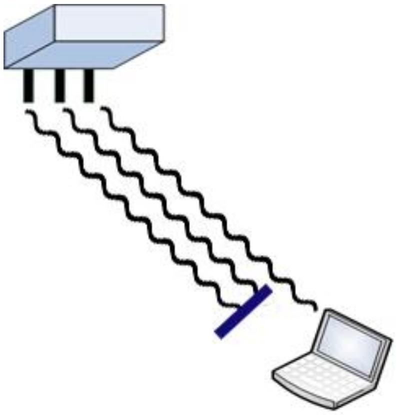

Figure 1: Multiple transmissions ensure that the client receives the data.

Quality of Service

All 802.11 standards include measures to allow operation on unlicensed frequencies that are often noisy and unpredictable. When a Wi-Fi radio sends a frame to another Wi-Fi radio, the receiver sends an acknowledgement to the sender that the frame was received. Until it receives that acknowledgement, the sender continues to resend the frame.

For multimedia applications such as voice and real-time video, a dropped packet usually is preferable to repeated transmissions while waiting for an acknowledgement. The “drop” may be imperceptible, whereas repeated attempts to resend may jeopardize the receipt of subsequent packets, leading to a further reduction in the quality of the transmission.

802.11n allows a recipient to provide a single acknowledgement that covers a block of frames and indicates which of the frames were received successfully. The sending device then can resend only those frames that merit a resend. The block acknowledgement capability of 802.11n made its debut in the IEEE 802.11e standard for QoS.

SECTION 2: HOW 802.11N WORKS

802.11n uses the same Orthogonal Frequency Division Multiplexing (OFDM) transmission method that is at the core of the 802.11a and 802.11g standards. OFDM is a technology that divides a single radio channel into multiple smaller channels referred to as “subcarriers”. Four modulation types are used with OFDM: Binary Phase Shift Keying (BPSK), Quaternary Phase Shift Keying (QPSK), 16 Quadrature Amplitude Modulation (QAM) and 64 QAM.

802.11a and 802.11g divide 20 MHz wide channels into 48 subcarriers for data. The use of these four modulation types with 48 subcarriers allows for eight data rates ranging from 6 to 54 Mbps. 802.11n divides a single channel into 52, not 48, subcarriers for data. When the same four modulation types are applied to the increased number of subcarriers on a 20 MHz channel, resulting data rates range from 6.5 to 65 Mbps.

As previously mentioned, 802.11n supports a top data rate of 600 Mbps (not 65 Mbps). The dramatic performance improvements of 802.11n result from major enhancements to the media access controller (MAC) and physical layer controller (PHY) of an 802.11-compliant radio. The remainder of this section provides a high-level explanation of these enhancements.

Note: For a more detailed technical explanation of 802.11n advancements, see Appendix 2.

The MAC of an 802.11 radio resides between the host device (such as a data terminal) and the PHY. The 802.11n standard provides two main MAC layer enhancements – block acknowledgement and packet aggregation. 802.11n also provides for changes to power save modes to accommodate the increased power consumption inherent in some 802.11n PHY layer enhancements.

The PHY (sometimes called the baseband processor or modem) resides between the MAC and the transceiver (2.4 GHz and 5 GHz to support both available unlicensed WLAN operating bands). Like a traditional modem, the PHY modulates and demodulates a signal between the MAC and the transceiver (converts digital to analog and analog to digital) to allow for send and receive operations. Most of the additional performance gains associated with 802.11n result from three major PHY enhancements: wider channels, reduced interframe space, and MIMO.

Table 1 provides a summary of 802.11n advancements and resulting benefits.

Table 1: Summary of 802.11n enhancements and benefits

Type	Name	Summary	Benefit(s)
MAC	Block acknowledgement	Acknowledgement of a block of frames	Increased throughput Improved quality of service
MAC	Packet aggregation	Multiple data packets with each header	Increased throughput
PHY	40 MHz channels	Twice the width = twice the data	Increased throughput
PHY	Reduced interframe space	Increased number of frames for a given period of media access	Increased throughput

Type	Name	Summary	Benefit(s)
PHY	MIMO: Transmit beam forming	Sending same signal multiple times	Increased throughput More predictable coverage
PHY	MIMO: Maximal ratio combining	Combining copies of a signal	Increased throughput Increased range
PHY	MIMO: Spatial multiplexing	Sending multiple streams of data	Increased throughput

MAC Enhancement: Block Acknowledgement

All WLAN standards – 802.11, 802.11a, 802.11b, 802.11g, and 802.11n – define a robust technology that can operate on unlicensed frequencies which are often noisy and unpredictable. Each time one radio sends a frame to another, the receiver sends an acknowledgement of receipt to the sender. Until it receives that acknowledgement, the sending station continues to resend the frame.

Although per-frame acknowledgements increase robustness, they decrease throughput. Resending data also reduces throughput but, for data such as voice or real-time video, can be unnecessary and detrimental to quality.

With block acknowledgement, a single acknowledgement can cover a range of frames, informing the sending station of which frames were successfully received and which were not. The sending device can then resend only the necessary frames. Block acknowledgement was first introduced in the IEEE 802.11e standard for QoS to support time-sensitive data such as voice or real-time video.

MAC Enhancement: Packet Aggregation

With packet aggregation, an 802.11n radio includes multiple data packets with a single header, which reduces the number of transmissions needed for a block of data. With fewer transmissions, there are fewer media accesses, collisions, and resends, which leads to higher throughput.

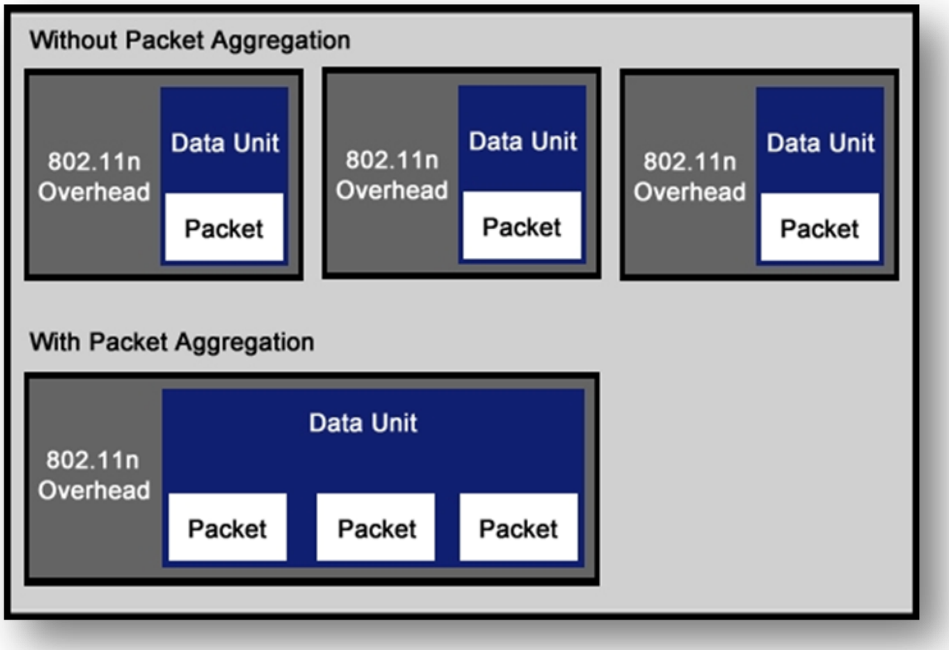

Figure 2: Packet aggregation

PHY Enhancement: Wider Channels

A channel can be thought of as a lane on a highway. If a two-lane highway can support a certain number of vehicles, switching to a four-lane highway doubles the number of vehicles supported.

Figure 3: Two-lane versus four-lane capacity

Pre-802.11n WLAN physical layer standards (802.11a, 802.11b, and 802.11g) provide for channels that are 20 MHz in width. At this width, and depending upon regulatory domain, a total of three non-overlapping channels are available in the 2.4 GHz band, and up to 23 non-overlapping channels are available in the 5 GHz band.

The 802.11n standard provides for channel widths of 20 MHz and 40 MHz. Wider 40 MHz channels allow about twice as much transmitted and received data as 20 MHz channels. Double-wide channels have been employed by fixed wireless systems for years, but 802.11n represents the first time that wider channels have been supported by an 802.11 standard.

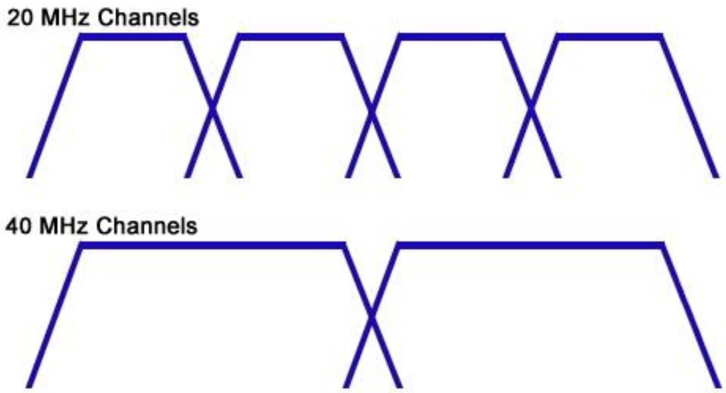

Figure 4: 20 MHz versus 40 MHz channels

PHY Enhancement: Reduced Interframe Space

A WLAN (802.11x) device first senses and then accesses the wireless medium prior to transmitting. The 802.11e standard is designed to enhance support for latency-sensitive traffic such as voice and real-time video. With 802.11e, a transmitting client may send multiple frames during a single media access, a period of time known as a *transmit opportunity*.

During a transmit opportunity, the sending station creates space between individual packet transmissions. Each space is referred to as an *interframe space* or a *guard interval*. The longer the interframe space, the less chance there is for self-interference where successive transmissions collide.

According to 802.11e, the short interframe space (SIFS) is 800 nanoseconds (800 ns). 802.11n introduces a reduced interframe space (RIFS) of 400 ns. If there is no self-interference, then RIFS improves throughput. 802.11n devices can switch between SIFS and RIFS as environmental conditions dictate.

PHY Enhancement: MIMO

Of all the 802.11n standard performance enhancements, MIMO (multiple input, multiple output) receives the most attention, has the most misunderstood concepts, and makes the greatest contribution to improved performance. MIMO, like double-wide channels, is employed for fixed wireless deployments. MIMO also is a foundational element of Wi-MAX or 802.16, a wireless standard for licensed band operation. 802.11n represents the first time that MIMO has been applied to an 802.11 standard.

MIMO can be thought of as a superset of systems supported by 802.11n:

- SISO – Single input, single output
- SIMO – Single input, multiple output
- MISO – Multiple input, single output
- MIMO – Multiple input, multiple output

There are three MIMO enhancements: transmit beam forming (TxBF), maximal ratio combining (MRC), and spatial multiplexing (SM). These enhancements take advantage of a phenomenon called *multipath propagation.* WLAN transmissions are radio waves. When a wave comes into contact with a reflective surface, it duplicates because of multipath propagation. When the wave is a sound wave, an echo results and the listener hears the same sound repeated in diminishing volumes over a period of time. When the wave is an 802.11 transmission, the recipient receives a series of duplicated transmissions that diminish in strength over time. A duplicate signal may arrive out of phase (out of sequence with other transmitted packets) and distort the original signal, which makes it more difficult for the receiver to decode.

Each of the three MIMO enhancements relies on support for more than one transmit or receive chain, where a chain is a set of WLAN radio components that operates on its own antenna. When a device supports at least two transmit chains and at least two receive chains, it can transmit or receive two or more spatial streams, or distinct sets of data, at a time. Table 2 on the next page details the capabilities of various systems of transmit and receive chains.

Table 2: Capabilities of SISO, SIMO, MISO, and MIMO systems

System Type	Notation	Transmit Chains	Receive Chains	Spatial Streams	Support
SISO	1 x 1	1	1	1	None
SIMO	2 x 1	2	1	1	TxBF
MISO	1 x 2	1	2	1	MRC
MIMO	2 x 2	2	2	2	TxBF, MRC, SM
MIMO	3 x 2	3	2	2	TxBF, MRC, SM
MIMO	3 x 2	2	3	2	TxBF, MRC, SM
MIMO	3 x 3	3	3	3	TxBF, MRC, SM
MIMO	4 x 4	4	4	4	TxBF, MRC, SM

Transmit Beam Forming (TxBF)

TxBF uses multiple transmitters to take advantage of multipath propagation. Working in harmony, transmitters send the same signal at slightly different times. When the sender and receiver know how long it takes for each signal to travel, then all signals arrive at the receiver simultaneously, in phase, and with the combined strength of all signals. Because a stronger signal is capable of supporting higher-order modulation schemes, the combined signal supports higher data rates; TxBF increases throughput.

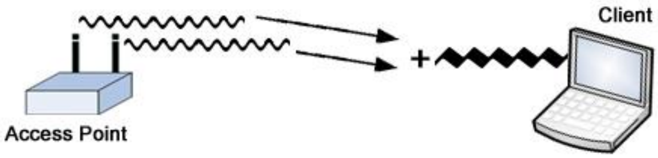

Figure 5: Transmit beam forming (TxBF)

When the sender and receiver do not know how long it takes for each signal to travel, the signals arrive at different times, so there is no throughput increase. Because the multiple signals fill in nulls, TxBF improves the reliability of coverage. Because TxBF applies to unicast traffic but not to multicast traffic (such as management frames), TxBF does not increase range.

Maximal Ratio Combining

MRC can be thought of as TxBF in reverse. With MRC, the recipient provides multiple receivers for a single transmitted signal, and these receivers receive and combine duplicate copies that arrive on multiple paths. The aggregated signal has the combined strength of all signals which boosts throughput. MRC does not require feedback from the transmitting device so it boosts the signal strength of unicast traffic and multicast traffic (including beacons). As a result, MRC improves throughput and range. It also reduces nulls which improves the reliability of coverage.

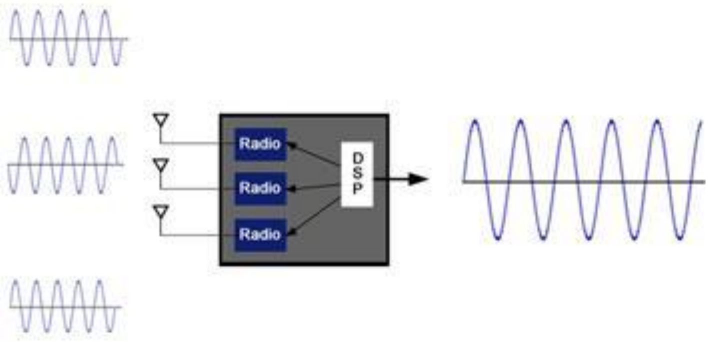

Figure 6: Maximal ratio combining (MRC)

Spatial Multiplexing

SM takes advantage of multipath propagation to send multiple streams of data (spatial streams) on a single channel. By enabling a transmitter to multiplex data on a given channel, SM can increase throughput by a factor equal to the number of spatial streams, although issues related to antenna spacing, cost, and power can result in diminishing marginal increases in throughput with additional spatial streams. The number of possible spatial streams is equal to the number of available transmit and receive chains, but an 802.11n system adjusts the number of spatial streams based upon the number of antennas available at any given time.

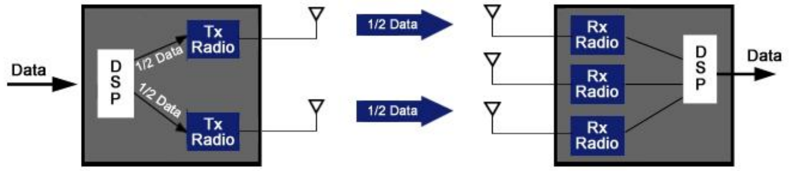

Figure 7: Spatial multiplexing

Summary: PHY Enhancements

Here is a summary of how PHY enhancements affect performance:

- Channel width: Using channels with a width of 40 MHz rather than 20 MHz results in an approximate doubling of the data rates, if all other factors remain constant. For example, with an 800 ns (nanosecond) interframe space and a single spatial stream, the maximum data rate with a 20 MHz channel is 65 Mbps. When the channel width is doubled to 40 MHz, the maximum data rate doubles to 130 Mbps.
- Interframe space: Decreasing the interframe space from 800 to 400 ns provides a relatively modest 11% data rate increase, if all other factors remain constant.
- Spatial streams: As spatial streams are added, the data rate increases as a factor of the number of spatial streams, if all other factors remain constant. With an 800 ns interframe space and a single spatial stream, the maximum data rate with a 20 MHz channel is 65 Mbps. When a second special stream is added, the maximum data rate doubles to 130 Mbps. When a third is added, the maximum triples to 195 Mbps.

A simplified table[4] of 802.11n data rates is provided in Table 3:

Table 3: 802.11n data rates

Channel Width (MHz)	Spatial Streams	Interframe Space (ns)	Minimum Data Rate (Mbps)	Maximum Data Rate (Mbps)
20	1	800	6.5	65
20	1	400	7.2	72
20	2	800	13	130
20	2	400	14.4	144
20	3	800	19.5	195
20	3	400	21.7	217
20	4	800	26	260
20	4	400	28.9	289
40	1	800	13.5	135
40	1	400	15	150
40	2	800	27	270
40	2	400	30	300
40	3	800	40.5	405
40	3	400	45	450
40	4	800	54	540
40	4	400	60	600

In terms of data rate, the biggest improvement results from the combination of 40 MHz channels and multiple spatial streams. When compared to a 20 MHz channel and a single spatial stream, a 40 MHz channel coupled with just one additional spatial stream quadruples data rate (a doubling of a doubling).

In summary, it is the multiplicative effect of 802.11n PHY enhancements that result in the considerable performance enhancements associated with 802.11n. Performance improvements are more modest when any one of the three enhancements is removed.

SECTION 3: DEPLOYING 802.11N IN YOUR INFRASTRUCTURE

When deploying 802.11n in a WLAN infrastructure, an organization must accommodate two aspects of the standard that do not exist with pre-802.11n standards:

- Throughput greater than 100 Mbps
- Power requirements greater than those supported by 802.3af Power over Ethernet (PoE)

In addition, the organization must determine which operational mode is best for its WLAN clients.

Throughput Greater than 100 Mbps

When deploying APs compliant with pre-802.11n standards, organizations have used 100 Mbps Fast Ethernet uplinks. Such an uplink could become a bottleneck with 802.11n APs, decreasing network performance. Most 802.11n AP vendors recommend upgrading the AP’s Ethernet link to Gigabit (1,000 Mbps) Ethernet to support even the highest data rates and throughput provided by 802.11n.

Increased Power Requirements

An AP that supports simultaneous operation in 802.11g and 802.11a has two radios, one for 802.11g (2.4 GHz band) and the other for 802.11a (5 GHz). Similarly, a dual-band 802.11n AP has two radios, one supporting 2.4 GHz operation and one supporting 5 GHz operation. A dual-radio AP requires more power than single-radio AP. While each of the radios in an 802.11a/g AP provides a single transmit and receive chain, each of the radios in a dual-band 802.11n AP provides multiple transmit and receive chains to support spatial multiplexing and associated higher throughput. Supporting more transmit and receive chains requires considerably more power.

The IEEE 802.3af specification for PoE provides for up to approximately 13 watts of power at a receiving device such as an AP. Most 802.11a/g APs consume nearly this much, and multiple transmit and receive chains in dual-radio 802.11n APs push the power requirements over this limit.

A successor standard to 802.3af, 802.3at, provides a minimum of 30 watts at the receiving device, which is more than sufficient to power dual-radio 802.11n APs. Prior to the ratification of 802.3at, vendors provided several means to power a dual-radio 802.11n AP over Ethernet, including:

- Powering the AP locally via AC power
- Introducing power to the Ethernet cable between the switch and the AP with an in-line power injector
- Reducing AP transmit power
- Limiting spatial streams
- Powering only a single radio

Operational Mode

The 802.11n standard provides three different operational modes. These modes address the frequent deployment of pre-802.11n and 802.11n clients and are designed for a spectrum of client device types:

- Legacy Mode – No 802.11n clients
- Mixed Mode – Combination of pre-802.11n clients and 802.11n clients
- Greenfield Mode – All 802.11n clients

Legacy Mode

Legacy Mode addresses a base of client devices compliant only with "legacy" or pre-802.11n standards such as 802.11a, 802.11b, and 802.11g. When an 802.11n AP is configured for Legacy Mode, it operates as if it were a legacy AP. In the 5 GHz band, it operates as an 802.11a AP; in the 2.4 GHz band, it operates as an 802.11g AP. In Legacy Mode, no 802.11n enhancements are enabled and no associated performance increase is expected.

An organization should consider using Legacy Mode when deploying 802.11n APs with non-802.11n clients. As 802.11n clients are added, the organization can move to Mixed Mode and enable 802.11n features in its infrastructure.

Mixed Mode

Mixed Mode is the best choice for most environments in which pre-802.11n client devices operate alongside 802.11n client devices. Mixed Mode operation provides some of the performance benefits of 802.11n while also providing normal operation for devices that lack 802.11n radios. This approach (backward compatibility at the expense of performance) is conceptually and technically similar to the approach taken with the 802.11g standard, which requires backward compatibility to 802.11b. Mixed Mode achieves backwards compatibility by making modifications to transmission frames and by changing the manner in which APs and clients access the wireless medium.

When in Mixed Mode, 802.11n APs add a legacy preamble and header to the front of every 802.11n frame to ensure that the frame can be decoded by clients that support pre-802.11n standards. Contained in the legacy header is an estimate of the duration of the transmission. With this information, pre-802.11n clients can compute the amount of time that they should avoid access attempts to the wireless medium. This is an effective means of decreasing or eliminating interference if the pre-802.11n clients actually receive the Mixed Mode frame. The issue of clients not receiving Mixed Mode frames from other clients in the same cell is addressed below.

While a client receives transmissions from the AP to which it is associated, that client may not be able to receive transmissions from another client associated to the same AP. The second client may be hidden from the first, leading to the "hidden node" problem.

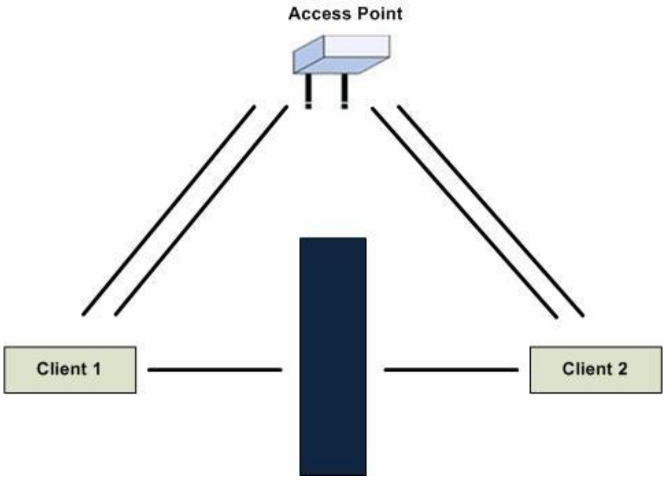

Figure 8: The "hidden node" problem

The use of RTS/CTS (Request to Send/Clear to Send), which is part of the original 802.11 standard, addresses the hidden node problem by introducing centralized media access control to 802.11's otherwise decentralized (and efficient) means of media access. When RTS/CTS is enabled, the sequence of events is:

- The client sends an RTS to the AP.
- The AP, which is capable of "hearing" all clients operating on its channel, determines if another client is currently accessing the wireless medium.
- If the medium is clear, the AP:
  - Issues a CTS, which frees the requesting client to send.
  - Instructs all other clients to refrain from media access for the estimated duration of the requesting client's transmission.

When an AP needs to transmit, it first sends a "CTS to Self", which notifies all clients that an AP transmission is about to occur and that they should refrain from a media access for the estimated duration of that transmission.

With 802.11n, the AP sends one type of CTS to pre-802.n clients and another type of CTS to 802.11n clients. When 40 MHz channels are employed, the AP sends one format for 40 MHz channels and another format for 20 MHz channels to support both 802.11n and pre-802.11n clients.

Clearly, additional network overhead is required to provide backward compatibility while in Mixed Mode. The additional overhead, marginally longer frames, RTSs, and multiple CTSs combine to decrease throughput. The centralized media access methods employed by Mixed Mode also are less efficient than decentralized media access; clients wait longer to transmit when in Mixed Mode which further decreases throughput.

Greenfield Mode

Greenfield Mode is provided for those environments in which all client devices support 802.11n. Greenfield Mode employs all of the enhancements found in the 802.11n standard and provides fully optimized performance.

Greenfield Mode does contain some risks. A device that lacks an 802.11n radio cannot decode 802.11n transmissions and perceives them as noise. When operating on the same channel as an 802.11n device, a pre-802.11n device has no specific indication that there are other devices contending for the same wireless medium beyond a noise assessment. Potentially chaotic results include un-received transmissions, continuing retries, and rising interference.

Pre-802.11n devices on different networks from the 802.11n device in Greenfield Mode may still cause issues; devices on neighboring networks but the same channel suffer the same issues as those on the same network.

SECTION 4: 802.11N AND MOBILE DEVICES

This final section examines the applicability of 802.11n features and benefits to mobile devices, especially those that run business-critical applications.

Types of Client Devices

Mobile Devices

Mobile devices are client devices that are in use while the user is in motion. Examples of mobile devices are mobile phones and mobile computers. When a mobile device runs mission-critical applications and requires a persistent network connection, it usually has stringent WLAN requirements in the areas of range and mobility. Range requirements result from the fact that a mobile device often operates at relatively long distances from the nearest AP, and the device's antenna orientation relative to the AP can change moment by moment. Mobility requirements stem from the fact that the device may move at high rates of speed, causing it to roam very quickly from one AP to the next.

Nomadic Devices

Most Wi-Fi client devices and their users that are commonly referred to as "mobile" are actually more "nomadic" in nature. A mobile device is used while it moves from one place the next. A nomadic device generally is not used while in motion. A laptop computer is a good example of a nomadic device because it generally is in use only when it is stationary.

Stationary Devices

A growing number of devices include Wi-Fi radios, not for mobile connectivity but to provide a cost-effective and convenient means of replacing cables. In consumer applications, desktop PCs and game consoles use Wi-Fi radios so that Ethernet cables are not required throughout the residence. In more commercial applications, WLAN connectivity is ideal for temporary workgroups and installation in difficult environments (such as factories). As the performance gap between wireless and mainstream wired technologies decreases, wireless cable replacement becomes increasingly viable for more and more applications.

802.11n Infrastructure: Benefits for Mobile Devices

Today, few business-critical mobile devices support 802.11n. When the infrastructure supports 802.11n but clients do not, those clients still receive the non-throughput benefits of 802.11n, namely improved predictability of coverage, increased range, and better quality of service.

MAC Layer Enhancements

Digitized voice packets are relatively small and require low latency. Packet aggregation and block acknowledgement, the two main MAC layer enhancements found in 802.11n, reduce latency for small packets, so they provide a significant benefit to voice-over-WLAN applications. Most mobile device data traffic uses small packets but is not latency-sensitive; the receipt of packets is more important than the order in which they were received. Because of this, 802.11n MAC layer enhancements do not provide significant benefits for most mobile device data traffic.

40 MHz Channels

When compared to 20 MHz channels, 40 MHz channels provide twice the data rate and throughput. Because they take up twice the bandwidth, however, they halve the number of available channels in a given operating band. If one employs 40 MHz channels in the 2.4 GHz spectrum, the number of non-overlapping channels is reduced from three to one, so using 40 MHz channels is not feasible. The 5 GHz operating band provides for approximately 10 double-wide channels, depending upon regulatory domain, so increasing throughput by implementing 40 MHz channels is a sensible approach. Pre-802.11n mobile devices, however, cannot operate on 40 MHz channels.

Transmit Beam Forming (TxBF)

Because TxBF sends out multiple instances of the same transmission, it reaches areas that would otherwise be "nulls" and therefore improves the predictability of coverage. While TxBF does not increase range, it results in more reliable connectivity in environments with a relatively high number of wall, shelves, and similar obstructions. Pre-802.11n mobile devices gain improved connectivity without the need for an 802.11n upgrade.

Maximal Ratio Combining (MRC)

For pre-802.11n mobile devices, MRC improves range and predictability of coverage without the limitations associated with TxBF. MRC is applied to the receive side of a link and does not require any distance calculations between the sending and receiving devices.

Spatial Multiplexing (SM)

SM provides no benefit to pre-802.11n devices because it requires that both sides of the Wi-Fi link – the client and the infrastructure – support 802.11n.

What about 802.11n on Mobile Devices?

To achieve a significant throughput boost from 802.11n, both sides of the Wi-Fi link must support 802.11n with more than one data stream. Putting 802.11n on laptops and other general-purpose client devices makes sense if those devices need a throughput boost. Most business-critical mobile devices don't need higher throughput.

While 802.11n provides for Legacy Mode and Mixed Mode to accommodate pre-802.11n client devices, both come at a significant cost to performance. When pre-802.11n devices operate only in the 2.4 GHz band, 802.11n devices can operate in the 5 GHz band in Greenfield Mode. Segregating mobile devices to the 2.4 GHz band is not the best long-term strategy, however, because those devices will be used for five years or longer. Consider dual-band 802.11n for mobile devices as soon as you can obtain it for a modest price premium.

CONCLUSION

802.11n enhancements such as double-wide channels and spatial multiplexing combine to bring significant performance improvements to client devices that support 802.11n with multiple streams. When 802.11n is deployed in the infrastructure, non-performance benefits such as improved coverage predictability accrue to pre-802.11n clients. With 802.11n infrastructure devices rivaling pre-802.11n infrastructure devices in price, organizations should give 802.11n serious consideration when deploying a new infrastructure or refreshing an existing one.

When considering 802.11n on client devices, organizations should focus first on stationary or nomadic devices that run media-rich applications and need a throughput boost. Pre-802.11n clients will limit the throughput of 802.11n clients, so organizations with heterogeneous client environments should plan for 802.11n upgrades on mobile devices. Once 802.11n radios are in all client devices, the network may be configured to operate in Greenfield Mode, maximizing throughput

APPENDIX 1: DATA RATE VS. THROUGHPUT

A notable benefit of 802.11n is greater speed of information transfer over the wireless medium. To understand the difference between data rate and throughput, it is helpful to understand how a WLAN radio transmits information.

To send data, the WLAN radio – in a client device or an infrastructure device such as an access point (AP) – must gain access to the wireless medium, i.e. the airwaves. Because this medium is shared with all other devices in a particular cell, the sending WLAN radio first must sense the medium (listen) and then send only when it has concluded that no other radios are currently accessing the medium (sending data). In the event that two or more radios send data simultaneously, a collision may occur and the radios must then restart the process and resend their data. The delays associated with resending data reduce performance. As only one device may access the medium at a given time, WLANs provide half-duplex communication: a device can either send or receive, but not both at a given time.

When sending data, a WLAN radio encapsulates each of these data packets inside a "frame" that includes a "header" followed by the actual data packet. A header contains information necessary for the frame to arrive successfully at the intended device. Depending upon the size of the data packet, this header can be as large as, if not larger than, the actual data. The header is discarded once the data reaches its destination. From the standpoint of performance, then, the header is considered "overhead" and the actual data is considered "payload".

Both data rate and throughput are measures of the speed at which information is sent, expressed in bits per second (bps). Data rate applies to everything that is sent, both overhead (headers) and payload (data). Throughput applies to just the data that is sent. As a result, throughput is a better measurement of the actual user experience. Unlike switched technologies, all client devices in a given cell (all clients associated to a given AP) share the same medium; an increase in the number of clients in a cell reduces per user throughput. It is important to understand that throughput typically is measured with a single client in a cell, and that scenario is inconsistent with real-world operation. Per user throughput with multiple clients in a cell will be lower.

APPENDIX 2: MIMO – TECHNICAL DETAILS

SISO: Pre-802.11n Systems

802.11a, 802.11b, and 802.11g represent SISO (single input/single output) systems. With these systems, each of the wireless devices receives a single data stream (single input) and transmits a single data stream (single output). A SISO device has a single transmitter and receiver and uses a single antenna. If a SISO device has two antennas instead of one, then only one is operational at any given time.

Multipath propagation may cause the recipient to receive multiple copies of a transmission. With a SISO system, duplicate transmissions of the same information arrive out of phase, interfere with each other, and potentially degrade performance. There two are ways to decrease the negative effects of multipath propagation on SISO systems:

1. Exploit capabilities within the radio to identify and ignore. This capability is quantified as *delay spread* (the interval during which a radio can discern a duplicate transmission from the primary transmission) which is typically expressed in nanoseconds. The greater the delay spread, the better the radio copes with multipath propagation.
2. Use diversity antenna systems, which provide two separate antennas that are typically spaced at least one-half of a wavelength from each other. In a diversity system, the radio dynamically chooses between the two antennas for transmitting and receiving and selects the antenna that provides the best path between the communicating devices. Although there is more than one antenna, only one is operational at a given time; the system supports only a single input and single output.

SIMO: Transmit Beam Forming (TxBF)

Fundamental to SIMO (single input/multiple output) is transmit beam forming (TxBF). TxBF is possible when the transmitting (output) device has multiple transmitters but the recipient has only a single receiver.

TxBF takes advantage of the multiple paths taken by a signal from a given transmitter to arrive at a given receiver. With a SISO system, the single output (transmission) multiplies and the duplicate transmissions arrive at the receiver at different times ("out of phase"), interfere with each other, and potentially degrade performance. With TxBF, the sending station uses multiple transmissions of the same signal with a goal of exploiting multipath propagation.

Suppose that, in a system with multiple transmitters, each transmitter is aware of the amount of time required for its signal to reach the single receiver. If each transmitter sends the same signal at a slightly different time and each signal follows a slightly different path, then each signal will arrive at the receiver at the same time ("in phase") and be perceived as a single signal with a strength that is essentially equal to the combined signal strength of the multiple signals sent from the transmitter. A stronger signal is capable of supporting higher-order modulation schemes, which are capable of supporting higher data rates and result in improved throughput.

When the sending and receiving devices cannot identify the amount of time required for a signal to arrive at the receiver, TxBF still provides a benefit. By sending out multiple instances of a single transmission, a TxBF-enabled device tends to "fill in" areas that are missed with a single transmission. These un-covered areas (nulls) are common in environments with a large number of walls, shelves, or similar obstructions. This reduction of nulls is sometimes referred to by device vendors as an improvement in coverage "predictability". This benefit applies to devices with 802.11n radios and devices with pre-802.11n radios.

Implicit in this explanation of TxBF are the following:

- TxBF is a SIMO technology (rather than a MIMO technology) because it provides multiple signals only on the transmit side.

- TxBF is not a "smart antenna" technology and does not employ "beam steering". TxBF does not rely upon specialized antennas or directional antennas. TxBF tends to work best with standard omnidirectional antennas.
- In order to improve throughput, TxBF requires close coordination between the transmitting device and the receiving device. In order for multiple signals to travel multiple paths and arrive simultaneously at the same destination, the transmitter must receive feedback from the receiver to calculate the phase offsets (different transmit times) for each of the multiple transmissions. While some devices implicitly claim to support TxBF without calculated feedback from the receiver, there is considerable debate within the industry as to the effectiveness of such TxBF. Without question, TxBF performs best when the receiver supports 802.11n and provides the transmitter with explicit feedback as to the relative time for receipt of the multiple streams.
- Due to the close coupling required between the transmitting and receiving devices to combine signals, TxBF applies only to unicast traffic (traffic between two points) and not to multicast traffic (traffic between one and many points). Wireless control and management packets (such as beacons and probes) are sent as multicast traffic. It is this traffic that dictates the maximum distance between devices such as an AP and a client device. While TxBF allows for relatively higher data rates at a given distance, *it does not allow for increased range*.
- To combine signals, path calculations made by the transmitter from information provided by the receiver must be location-specific to result in simultaneous arrival of the individual signals across different paths. Given that the wavelength of a signal at 2.4 GHz is 120 mm (about 4 ¾ inches) long and a signal at 5 GHz is 55 mm (just over two inches) long, even the smallest movement by the receiver requires a recalculation.
- Close coupling between the sending and receiving station is not required to reduce nulls. This applies to legacy devices and is not impacted by the dynamic nature of mobile clients. It also applies to 802.11n and legacy devices.
- TxBF is a smart transmitter technology that requires both multiple transmit antennas and multiple *transmit chains* inclusive of power amplifiers, filters, mixers, and additional components. Because of this, a radio that uses TxBF tends to draw significantly more power than an SISO radio. In addition, the calculations necessary for TxBF require more processing power than that used by a more traditional radio.

MISO: Maximal Ratio Combining (MRC)

Maximal ratio combining (MRC) can be thought of as TxBF in reverse. Whereas TxBF provides multiple transmit signals to a single receiver, MRC provides multiple receivers for a single transmitted signal. MRC represents a MISO (multiple input, single output) technology, the opposite of the SIMO technology represented by TxBF.

A device capable of TxBF also can be capable of MRC. As an example, when a TxBF-capable AP is transmitting to a client device, it is in SIMO mode. When that same MRC-capable AP is receiving from a client device, it is in MISO mode.

Similar to TxBF, the use of antennas with MRC differs from more traditional antenna diversity. With a traditional multi-antenna diversity system, only one antenna is active at a given time. MRC, like diversity, relies on multiple antennas but also requires multiple antennas and multiple *receive chains* inclusive of amplifiers, filters, mixers, and additional components. With these multiple receive chains, MRC can build upon the advantages of traditional diversity to provide additional performance and additional range.

Like diversity, the MRC algorithm identifies which of the multiple receivers is receiving the strongest signal at a given point in time; it primarily uses that signal when demodulating received traffic. MRC also moves beyond diversity; all available receivers remain active and receive duplicate signals that arrive on multiple paths. MRC is able to combine these additional signals with the primary signal which results in an aggregated signal of a strength that is essentially the sum of all received signals. Unlike TxBF, MRC does not require close coupling between the receiver and the transmitter; the receiving device processes all received signals as they arrive and combines them to the degree possible without any additional implicit or explicit information from the transmitter.

Implicit in the above explanation of MRC are the following:

- Because MRC does not require feedback from the transmitting device, it works well with 802.11n devices that support MRC and with 802.11a, 802.11b, and 802.11g devices that do not support MRC. MRC provides a quantifiable benefit for both new and legacy devices.
- Because MRC does not require feedback from the transmitting device, it works well with both unicast and multicast traffic. As such, it increases the signal strength of all transmitted traffic (including beacons), which *improves both throughput and range*.
- As with TxBF, MRC generally reduces the number of nulls in a particular environment, which improves the “predictability” of coverage. This benefit applies to 802.11n and legacy.
- MRC relies on multiple antennas, multiple receive chains, and the additional processing to combine received signals. As such, MRC-enabled devices generally draw more power than non-MRC devices.
- A radio consumes more power when transmitting than when receiving. While an MRC-enabled device draws more power than one that is not MRC-enabled, MRC consumes far less power than TxBF.

MIMO: Spatial Multiplexing

Spatial multiplexing (SM) is fundamental to MIMO (multiple input, multiple output). With SISO, MISO, and SIMO systems, a single stream of data is sent and received, but often in multiple transmissions. SM differs by taking advantage of multipath propagation to send multiple streams of data on a single channel.

By transmitting multiple spatial streams (such as separate and distinct sets of data), SM is able to multiplex data on a given channel and theoretically increase throughput by a factor equal to the number of spatial streams. The 802.11n standard allows for as few as one (for specialized devices) and as many as four spatial streams (two spatial streams anticipated to be the most practical implementation). The number of possible spatial streams is equal to the number of available transmit and receive chains. For example, two transmit/receive chains allow for two spatial streams. Due to environmental issues, the number of spatial streams supported is not necessarily equal to the number of transmit/receive chains. This is discussed in more detail later in this document. An 802.11n system dynamically adjusts the number of spatial streams (and performance) based upon the number of antennas available at any given time. If the number of transmit and receive chains is asymmetric, the additional transmit or receive chain may be dynamically employed to support TxBF or MRC, respectively.

Although the 802.11n specification defines as many as four spatial streams, the throughput gains associated with spatial multiplexing are subject to diminishing returns. This limitation on the number of spatial streams (and the resulting performance limitation) is due to three factors: antenna spacing, cost, and power.

Antenna Spacing

As the space between antennas is increased, their ability to support separate spatial streams is also increased. Ideally, antennas should be placed at least one wavelength apart. A wavelength in the 2.4 GHz band measures 120 mm (about 4 ¾ inches) and a wavelength in the 5 GHz band measures 55 mm (just over two inches). Therefore, a device with two optimally spaced 2.4 GHz antennas would be at least 120 mm wide, a device with three optimally placed 2.4 GHz antennas would be at least 240 mm (about 9 ½ inches) wide, and a device with four optimally placed 2.4 GHz antennas would be at least 360 mm (14 ¼ inches) wide. For operation in the 5 GHz band, the optimal antenna spacing is approximately half that of 2.4 GHz. Two conclusions can be drawn:

1. Given that APs are fixed (non-mobile) devices, it is far more practical to increase the size of an AP than to increase the size of a client device, particularly a highly mobile device. This may create common asymmetric systems such as 2 x 1, 1 x 2, 3 x 2, and 2 x 3, with the “extra” antenna residing on the AP and used to support TxBF and MRC.
2. Optimal antenna placement is more practical in the 5 GHz band than in the 2.4 GHz band because of the 5 GHz band’s shorter wavelength. Because the narrowness of the 2.4 GHz band also makes the use of 40 MHz wide channels impractical, 802.11n enhancements are exploited less in the 2.4 GHz band than in the 5 GHz band.

While it is generally ideal to space antennas a wavelength apart, it may not be practical for some client devices; therefore, some devices may have antennas that are less than one wavelength apart. The impact this may have on performance is explained below.

Cost

It costs more money to build a radio with multiple transmit and receive chains than it does to build a radio with single transmit and receive chains. Each chain requires a separate antenna, filters, mixers, and other components. This, coupled with the diminishing performance returns that accrue due to antenna correlation, results in decreasing value as additional transmit and receive chains are added.

Power

Transmit chains and, to a lesser extent, receive chains consume power. Multiple transmit and receive chains consume more power than single transmit and receive chains. Access points based on Draft 2 of the 802.11n standard consume more power than is provided by the 802.3af PoE standard. This standard is more than sufficient to power even dual radio 802.11a/g APs. This issue is especially acute on client devices that are often run on batteries. This also tends to promote asymmetric systems that have a greater number of transmit and receive chains on the infrastructure side than the client side.

APPENDIX 3: ADDITIONAL NOTES ON PERFORMANCE

Here are some other important points on the performance benefits of 802.11n:

- Within an 802.11a, 802.11b, or 802.11g packet, overhead is approximately equal to data, resulting in throughput that is approximately half that of data rate. With 802.11n packet aggregation, each packet contains more data and less overhead. Some 802.11n enhancements, however, add overhead that decreases throughput as a percentage of data rate.
- Data rate refers to the aggregate data rate for an entire cell, such as an AP and all the clients that are associated to the AP. Because 802.11n involves a shared medium, data rate (and throughput) on a per-client basis decrease as the number of clients in a cell increases.
- 802.11n allows for unequal modulation, which is the ability to support differing modulation types on different spatial streams. Different spatial streams on different paths present differing signal-to-noise ratios (SNRs), which allows for differing modulation types and resulting supported data rates. Unequal modulation optimizes performance on a stream-by-stream basis and thereby provides for optimized overall performance.

Two factors that can reduce the performance gains of 802.11n are an absence of multipath and correlated antennas:

- **Absence of multipath:** SIMO, MISO, and MIMO benefit from and rely upon the same multipath propagation that is a detriment to SISO systems. In the absence of multipath, a MIMO system operates essentially the same as a SISO system. Fortunately, the sorts of reflectors that cause multipath propagation exist in nearly every WLAN environment. That said, when there is line of sight between wireless devices, multipath propagation is significantly decreased or even eliminated.
- **Correlated antennas:** On a mobile device, multiple antennas can be placed in an orientation relative to the AP such that they become *correlated* (they are all receiving the same signal on the same path). The further the antennas are spaced apart (up the distance of the wavelength) the less likely it is that the antennas will become correlated. When antennas are correlated, the performance benefits from multiple paths are lost. The likelihood of the occurrence of correlation is far greater on a highly mobile device with relatively small spacing between antennas than on a larger, less mobile device with wider spacing between antennas.

RECOMMENDED READING AND END NOTES

Recommended Reading

The authors have drawn on a number of resources in researching this paper and are indebted to the authors of the excellent papers listed below. It should be noted that the conclusions drawn in this paper are those of the author and not necessarily in line with the conclusions drawn in the following documents.

1. *Wi-Fi CERTIFIED™ 802.11n draft 2.0: Longer-Range, Faster-Throughput, Multimedia-Grade Wi-Fi® Networks,* The Wi-Fi Alliance, 2007.
   http://www.wi-fi.org/whitepaper_80211n_draft2_technical.php

2. Dulaney, Ken and Ahlawat, Rachna, *Approach 'Draft 2.0' Wi-Fi-Certified 802.11n With Caution (ID Number: G00148909)*, Gartner, Inc. May 2007.
   http://www.gartner.com/DisplayDocument?id=506099&ref=g_fromdoc

3. DeBeasi, Paul, *802.11n: Beyond the Hype*, The Burton Group, July 2007.
   http://www.burtongroup.com/Research/PublicDocument.aspx?cid=1142

4. Sun, Winston, *MIMO Architecture: The Power of 3,* Atheros Communications, Inc., May 2006.
   http://www.atheros.com/pt/whitepapers/MIMO_Pwr3_whitepaper.pdf

5. *802.11n: The Next Generation of Wireless Performance*, Cisco Systems, Inc., 2007.
   http://www.cisco.com/en/US/netsol/ns767/networking_solutions_white_paper0900aecd806b8ce7.shtml

6. Thornycroft, Peter, *Designed for Speed: Network Infrastructure in an 802.11n World,* Aruba Networks, Inc., May 2007.
   http://www.arubanetworks.com/pdf/technology/whitepapers/wp_AWO.pdf

End Notes

1. Cox, John, “Does 802.11n spell the 'end of Ethernet'?”, *NetworkWorld Magazine*, September 11, 2007.
   http://www.networkworld.com/news/2007/091107-80211n-ethernet.html

2. Institute of Electrical and Electronics Engineers, *OFFICIAL IEEE 802.11 WORKING GROUP PROJECT TIMELINES*, September 2007.
   http://grouper.ieee.org/groups/802/11/Reports/802.11_Timelines.htm

3. Hamblen, Matt, “Update: Testing of Next-Gen Wi-Fi Standard to Start in June
   Final 802.11n Ratification Pushed Back to March 2009”, *Computerworld Magazine*, May 16, 2007.
   http://www.computerworld.com/action/article.do?command=viewArticleBasic&articleId=9019538

4. Thornycroft, Peter, *Designed for Speed: Network Infrastructure in an 802.11n World,* Aruba Networks, Inc., May 2007.
   http://www.arubanetworks.com/pdf/technology/whitepapers/wp_AWO.pdf