



802.11n Demystified

Key considerations for n-abling the Wireless Enterprise



Introduction

IEEE 802.11n is a next generation wireless technology that delivers spectacular improvements in the reliability, speed, and range of 802.11 communications. Delivering about 6 times the data rate and 2-5 times the throughput of 802.11a/g with substantial improvements in network coverage and connection quality, 802.11n is expected to replace wired Ethernet as the dominant local area network technology of the future. This white paper examines the operation of current generation wireless LANs, explains key 802.11n technology components along with their benefits, and explores 802.11n's potential for delivering a totally wireless enterprise.

802.11n's Current State

The IEEE 802.11n working group, called *Task Group N* (TGN), is working on publishing the 802.11n standard. In March 2007 TGN approved the 802.11n Draft 2.0 which has become the early standard for several 802.11n products available in the market today. Availability and adoption of 802.11n Draft 2.0 products has been aided, to a large extent, by the Wi-Fi Alliance's announcement of a Draft 2.0 interoperability certification program. The Wi-Fi Alliance is an industry organization that certifies the interoperability of 802.11 devices from different vendors. Meanwhile TGN is continuing to work on refining the 802.11n standard based on comments received from various experts. Based on current time lines, 802.11n is expected to achieve final IEEE Standards Board approval by November 2009.

802.11n Overview

802.11n introduces several enhancements to the 802.11 PHY (radio) & MAC layers that significantly improve the throughput and reliability of wireless communication. These enhancements include:

- **Improved OFDM** — A new, more efficient OFDM (Orthogonal Frequency Division Multiplexing) modulation technique that provides wider bandwidth and higher data rates.
- **40 MHz Channels** — 802.11n doubles data rates by doubling the transmission channel width from 20 MHz (used in 802.11a/g) to 40 MHz.
- **Multiple-Input / Multiple-Output (MIMO)** — A radio system (transceiver) with multiple inputs into the receiver and multiple outputs from the transmitter capable of sending or receiving multiple spatial streams of data (current 802.11 radios can transmit and receive a single data stream on a single antenna used in a Primary/Secondary diversity configuration).

- **Frame Aggregation** — 802.11n enhances the MAC layer and reduces the transmission overhead by allowing multiple data frames to be sent as part of a single transmission. Further reducing the interframe spacing between frames allows transmissions to be completed in a shorter amount of time, making the wireless medium available for other transmissions, and increasing overall throughput.

802.11n Draft 2.0 systems combine all of the above techniques to deliver connection speeds of 300 Mbps with improved communication range and more consistent coverage. While many 802.11n benefits can only be availed in an "N-only" environment (comprising of 802.11n access points and 802.11n clients), Draft 2.0 does establish mechanisms for backwards compatibility with 802.11a/b/g devices and offers several benefits even for existing 802.11a/b/g devices.

802.11n operates in two frequency bands: 2.4 GHz for 802.11b/g/n and 5 GHz for 802.11a/n. A *Phased Co-existence Operation* (PCO) mode allows 802.11n to dynamically switch between 20 MHz and 40 MHz channels while communicating with 802.11a/b/g and 802.11n devices, allowing for backwards compatibility in both frequency bands.

Understanding Radio Frequency (RF) Multipath & MIMO

In indoor environments, RF signals typically encounter several barriers (walls, doors, partitions, etc.) in their communication path towards the receiver. These barriers either reflect or absorb the original RF signal, creating multiple reflected or "secondary" waveforms. Multipath results when multiple copies of the original RF signal travel different paths to arrive at the receiver.

Multipath traditionally has been considered the enemy of RF communication. Multiple reflected signals arriving at the receiver with varying delays

make it difficult for the receiver to separate a good signal from poor quality signals. Weaker signals may not be deciphered accurately, resulting in data corruption and retries; furthermore, coverage holes can occur if reflected signals are out-of-phase but are received simultaneously.

The higher the multipath in an environment, the more the likelihood of poor RF performance resulting from weaker received signal strength (RSSI), increased retries, and dead spots. Conventional RF system design has addressed the problem of multipath through the use of antenna diversity using two antennas for each radio in an access point. Antennas in a diversity configuration function almost like redundant antennas. A signal from only one antenna is used at any time. Diversity switching logic implemented on the access point decides when to switch between the primary and the secondary (diversity) antenna for receiving the best signal.

MIMO introduces a new paradigm in RF systems design. MIMO-capable radios actually perform better within a multipath-rich environment. A MIMO system has multiple radio chains each of which is a transceiver with its own antenna. A radio chain refers to the hardware necessary for transmit/

receive signal processing. A MIMO radio can then apply several techniques to enhance signal quality and deliver more throughput. It is this ability to add signal components from multiple antennas that differentiates MIMO access points from traditional access points that use antennas in a diversity configuration. An access point with antenna diversity selects signal components from the antenna that provides the best signal performance and ignores the other antenna.

A MIMO system has multiple Radio Frequency (RF) chains implemented in the radio allowing the processing of multiple RF signals from multiple antennas. Depending on the number of transmit/receive antennas and the number of spatial streams, a MIMO system is often classified as a TxR:S system. Under this nomenclature T refers to the number of transmit antennas, R refers to the number of receive antennas and S refers to the number of spatial streams (transmitted data streams) the system can process. For example a 3x3:2 system can transmit and receive 2 data streams on its 3 antennas.

A MIMO-enabled access point (a radio system supporting multiple radio chains) may employ one or more MIMO techniques:

Maximal Ratio Combining (MRC)

To understand MRC, consider a traditional access point that implements diversity antennas. Depending on the multipath caused by the environment, multiple RF signals will arrive on both antennas. The access point samples its antennas and selects the preferred signal, ignoring the other signal. Diversity, in a sense, wastes RF energy, using the signal from any one antenna for a transmission. MRC is a receive-side MIMO technique that takes RF signals from multiple receive antennas and combines them within the radio to effectively boost the signal strength. This MIMO technique is fully compatible with 802.11a/b/g devices and significantly improves receiver sensitivity and overall gain for the access point radio, especially in multipath environments.

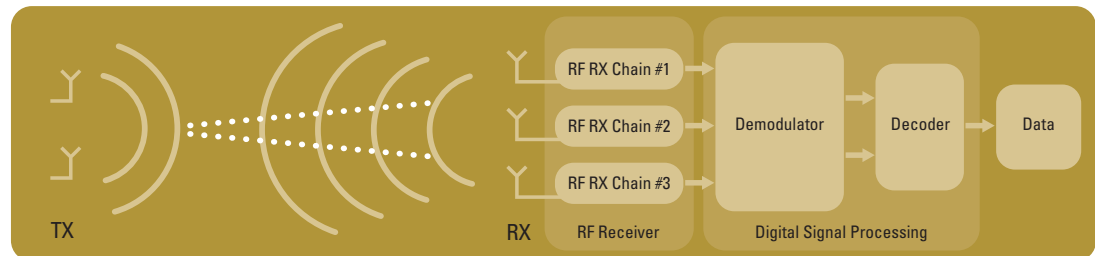


Figure 1: Signal Combining at Receiver

Transmit Beamforming (TxBF)

Transmit beamforming, often abbreviated as TxBF, uses multiple transceivers with antennas (radio chains) to focus RF energy toward the target receiver. An antenna array focuses energy toward the intended receiver in a way that less energy is wasted in other directions. As discussed previously, a single RF waveform arrives at the receiver as multiple waveforms and typically out-of-phase. The difference in phase can result in either an attenuated, amplified, or corrupted waveform at the receiver. TxBF adjusts the phase of the RF signal at the transmitter such that the receiver signal is an amplified waveform. Standards-based TxBF is only supported for 802.11n clients. It requires the client to participate in the beamforming process by sending special frequency characteristics and channel response information to the access point. This information is used by the AP in calculating the phase adjustment for its next transmission. However, any changes in the multipath environment (movement of device, changes in obstructions) will nullify the beamforming phase optimization.

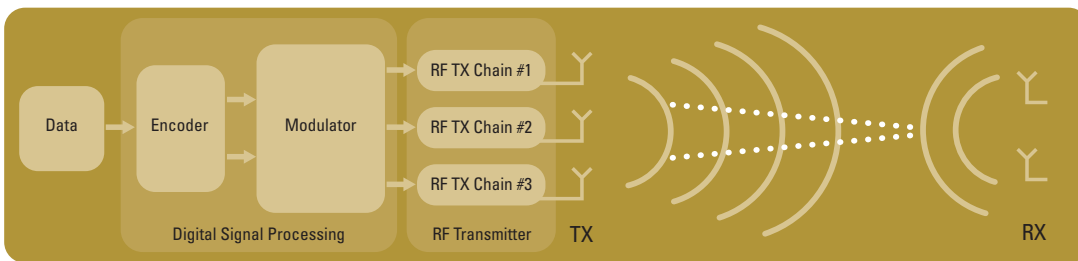


Figure 2: Transmit Beamforming

Spatial Multiplexing

Spatial multiplexing is the fullest and most powerful application of MIMO and involves the transmission of spatial streams using N (or more) antennas. Spatial multiplexing requires a compatible MIMO client capable of receiving and de-multiplexing N spatial streams. Each spatial stream carries a unique data stream allowing the system to dramatically improve data rates and range. Spatially multiplexing MIMO systems are represented as TxR:S, where T represents the number of transmit antennas, R the number of receive antennas and S the number of spatial data streams. As stated earlier, a 3x3:2 MIMO access point can transmit and receive 2 data streams on its 3 antennas. The majority of 802.11n clients available in the marketplace can transmit and receive 2 spatial streams.

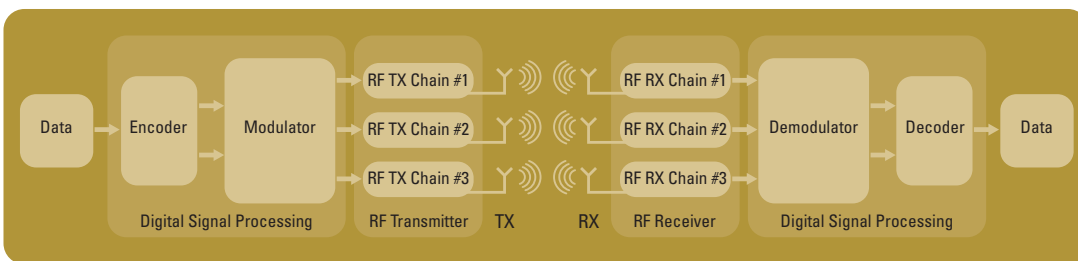


Figure 3: Multiple Transmit Chains and Multiple Receive Chains

802.11n PHY

Although MIMO is the core building block for 802.11n, MIMO alone cannot deliver high throughput. The 802.11n physical layer (radio) implements wider channel width along with improved modulation techniques that provide a high bandwidth medium for multiple spatial data streams. As an analogy, if MIMO's spatial multiplexing delivers 2 trains, the 802.11n PHY provides 2 separate tracks on which to run the 2 trains, thereby doubling the transport of packets and increasing throughput.

Improved OFDM and Channel Bonding

802.11b radios use DSSS modulation on a channel that is 22 MHz wide to deliver data rates between 1 Mbps and 802.11 Mbps. 802.11a and 802.11g both use 20 MHz channels, with OFDM modulation to deliver data rates between 6 Mbps and 54 Mbps. The table below shows the relationship between the spreading methods, modulation techniques and summarizes these data rates.

One of the reasons 802.11a and 802.11g offer higher data rates than 802.11b is they use OFDM modulation techniques.

802.11n uses a more efficient OFDM modulation and can use 40 MHz channels. This more than doubles the data rate for 802.11n when compared to 20 MHz channels. When operating within a traditional 20 MHz channel, OFDM further slices the channel into 52 subcarriers (48 of which are used for carrying data). However, when 802.11n applies OFDM on a 40 MHz channel, the number of data-carrying subcarriers

do not simply double to 96 sub-carriers. Instead, they actually more than double to 114 subcarriers, including pilots (which do not carry data). This allows 802.11n to deliver a 65 Mbps data rate (instead of 54 Mbps) per 20 MHz channel for a total of 135 Mbps on a 40 MHz channel when transmitting a single spatial data stream. When transmitting using 2 spatial streams on a 40 MHz channel, this data rate again doubles to 135 Mbps x 2 — 270 Mbps.

The 802.11-2007 standard dictates that every bit transmitted over the air is represented as a symbol. 802.11b transmits 1 symbol per microsecond for its 1 Mbps data, 2 symbols per microsecond for its 2 Mbps data rate and so on. An 802.11a/g symbol is transmitted for 4 microseconds and packs many more bits — up to 216 bits for the 54 Mbps data rate.

OFDM symbols are separated by a guard interval (GI) to reduce inter symbol interference. Since no bits or data is transmitted during this time, the GI is typically an overhead to the communication process. The longer the GI between symbols, the lesser the throughput. 802.11a and 802.11g symbols are separated by an 800ns GI. 802.11n introduces the concept of a short GI (400ns). In a high-multipath environment, a long GI setting will typically improve overall RF performance, whereas in low multipath environments, a short guard setting can be more advantageous.

As noted earlier, when transmitting 2 spatial streams in a 40 MHz channel, 802.11n supports data rates of up to 270 Mbps. However, when operating with a short GI 802.11n further improves data rates to deliver a top speed of 300 Mbps.

Modulation	DSSS		CCK		OFDM							
	BPSK	BPSK	QPSK		BPSK		QPSK		16-QAM		64-QAM	
Data Rate (Mbps)	1	2	5.5	11	6	9	12	18	24	36	48	54
802.11b (2.4 GHz)	*	*	*	*								
802.11g (2.4 GHz)	*	*	*	*	*	*	*	*	*	*	*	*
802.11a (5 GHz)					*	*	*	*	*	*	*	*

Table 1: Modulation Techniques and Data Rates

802.11a/g Rates	One Spatial Stream (MCS 0-7)			Two Spatial Streams (MCS 8-15)		
	802.11n Mandatory Rates (20 MHz channel)	802.11n Mandatory Rates (40 MHz channel)	Short Guard Interval Enabled	Two Spatial Streams	802.11n Mandatory Rates (40 MHz channel)	Short Guard Interval Enabled
6	6.5	13.5	15	13	27	30
9	13	27	30	26	54	60
12	19.5	40.5	45	39	81	90
18	26	54	60	52	108	120
24	39	81	90	78	162	180
36	52	108	120	104	216	240
48	58.5	121.5	135	117	243	270
54	65	135	150	130	270	300

Table 2: MCS Rate Index

802.11n Frame Aggregation Techniques

802.11 communication uses a shared medium that has a significant amount of overhead associated with it. Unlike wired Ethernet, there is no collision detection mechanism available in the wireless medium. Every wireless frame requires a positive *Acknowledgement* (ACK). This requirement for transmitting an ACK frame for every control/data frame has a huge performance penalty and significantly reduces throughput of 802.11 communication systems. 802.11 devices are also required to use a random wait interval (backoff period) after transmitting a frame and before getting access to the medium to transmit the next one. This further reduces aggregate system throughput.

802.11n introduces 3 key enhancements which address the inefficiencies of the traditional 802.11 MAC layer. Two of these techniques rely on frame aggregation; the third technique reduces interframe spacing between transmissions.

MSDU Aggregation

The term MSDU (MAC Service Data Unit) refers to the payload that is carried by the 802.11 MAC layer frame. An MSDU typically consists of an LLC header,

IP header and the IP packet payload from layers 4-7. When Mobile Units (MUs) communicate with different hosts in a network they still send all the 802.11 frames to the access point. Access point receives the 802.11 frames and forwards them to appropriate destinations. This may involve adding an 802.3 header if the destination is a wired host or a new 802.11 header if the destination is an 802.11 wireless host. The MSDU aggregation technique exploits this behavior and allows for a mechanism to aggregate multiple payloads in a single 802.11 frame. Since the MU has a single security association with the access point, this large 802.11 frame incurs the overhead of encryption (and decryption at the receiver) only once! This improvement is more pronounced for small size frames.

Since an A-MSDU frame is transmitted as a single 802.11 frame, receiver can acknowledge by sending a single ACK frame. There is also a significant increase in the maximum frame payload which reduces the acknowledgement overhead associated with 802.11 communications and improves overall throughput. The maximum A-MSDU size allowed by 802.11n is 8192 bytes. The disadvantage of aggregating A-MSDU is that each frame is only protected by a single checksum and an error in receiving an A-MSDU transmission incurs the overhead of having to retransmit the entire A-MSDU again.

No. -	Time	Source	Destination	Protocol	Info
1	0.000000	SymbolTe_ca:a6:f0	Broadcast	IEEE 802	Beacon frame, SN=3824, FN=0, Flags=
2	0.092643738	10.1.7.2	10.1.8.2	IP	IPv6 hop-by-hop option (0x00)
3	0.092645645	10.1.7.2	10.1.8.2	IP	IPv6 hop-by-hop option (0x00)
4	0.092651367	10.1.7.2	10.1.8.2	IP	IPv6 hop-by-hop option (0x00)
5	0.092655182	10.1.7.2	10.1.8.2	IP	IPv6 hop-by-hop option (0x00)
6	0.092657089	10.1.7.2	10.1.8.2	IP	IPv6 hop-by-hop option (0x00)
7	0.092658997	10.1.7.2	10.1.8.2	IP	IPv6 hop-by-hop option (0x00)
8	0.092672348	10.1.7.2	10.1.8.2	IP	IPv6 hop-by-hop option (0x00)
9	0.092678070	10.1.7.2	10.1.8.2	IP	IPv6 hop-by-hop option (0x00)
10	0.092683792	10.1.7.2	10.1.8.2	IP	IPv6 hop-by-hop option (0x00)
11	0.092689514	10.1.7.2	10.1.8.2	IP	IPv6 hop-by-hop option (0x00)
12	0.092693329	10.1.7.2	10.1.8.2	IP	IPv6 hop-by-hop option (0x00)
13	0.092697144	10.1.7.2	10.1.8.2	IP	IPv6 hop-by-hop option (0x00)
14	0.092704773	10.1.7.2	10.1.8.2	IP	IPv6 hop-by-hop option (0x00)
15	0.092708588	10.1.7.2	10.1.8.2	IP	IPv6 hop-by-hop option (0x00)
16	0.092714310	10.1.7.2	10.1.8.2	IP	IPv6 hop-by-hop option (0x00)
17	0.092720032	10.1.7.2	10.1.8.2	IP	IPv6 hop-by-hop option (0x00)
18	0.092723847	10.1.7.2	10.1.8.2	IP	IPv6 hop-by-hop option (0x00)
19	0.092729569	10.1.7.2	10.1.8.2	IP	IPv6 hop-by-hop option (0x00)
20	0.092735291	10.1.7.2	10.1.8.2	IP	IPv6 hop-by-hop option (0x00)
21	0.092741013	10.1.7.2	10.1.8.2	IP	IPv6 hop-by-hop option (0x00)
22	0.092744827	10.1.7.2	10.1.8.2	IP	IPv6 hop-by-hop option (0x00)
23	0.094577790	Intel_6a:b9:4a (TA)	SymbolTe_ca:a6:f0 (RA	IEEE 802	802.11 Block Ack, Flags=.....C

Figure 4: 802.11n Frame Aggregation with Block Ack

MPDU Aggregation with Block ACKs

MPDU (MAC Protocol Data Unit) aggregation gathers 802.11 frames, which each already have an 802.11 header for the same destination and transmits them as a single frame. Since this process involves transmitting multiple 802.11 frames as a single “grouped” frame, each frame requires its own ACK; however, instead of transmitting each ACK individually, 802.11n introduces a Block ACK frame which compiles all the individual acknowledgements into a single frame which gets transmitted from the receiver to the sender. The Block ACK frame is essentially a bitmap, or matrix of which frames are being acknowledged.

One of the disadvantages of MPDU aggregation is that each 802.11 frame needs to be encrypted separately, adding encryption overhead. On the other hand, MPDU aggregation allows for the selective retransmission of those frames not acknowledged within the Block ACK. This can be very useful in environments which have a high number of collision or transmission errors. The maximum A-MPDU size allowed by 802.11n is 64K bytes.

Reduced Interframe Spacing (RIFS)

Normal 802.11 transmitters are required to implement a random backoff between transmissions. DCF (*Distributed Coordinated Function*) is a contention-based service widely implemented in infrastructure networks that defines the backoff period for devices. The interframe spacing in DCF is referred to *DCF Interframe Spacing* (DIFS).

DIFS is the minimum idle time for transmissions if the medium is idle for longer than the DIFS interval. Wi-Fi Multimedia (WMM) based QoS allows frame bursting for certain devices without requiring a random backoff. These WMM devices typically separate their ACK receipt and subsequent transmissions with a shorter interframe spacing, referred to as *Short Interframe Spacing* (SIFS). 802.11n introduces an even shorter interframe spacing called *Reduced Interframe Spacing* (RIFS).

802.11n Beacons & Mixed Mode Operation

802.11n beacons are sent on a regular 20 MHz channel using a low rate modulation. Beacons in the 2.4 GHz frequency are fully understood by 802.11b/g devices and 5 GHz beacons are fully understood by 802.11a devices. An 802.11n beacon contains additional High Throughput (HT) operational information about the access point including channel width (20/40 MHz), guard interval, and number of spatial streams (usually 2). An 802.11n beacon also contains frame aggregation information, such as the maximum MSDU and MPDU sizes. Due to all this extra information carried in the 802.11n beacon, the size of the 802.11n beacon frame is much larger than the conventional 802.11a/b/g beacon. 802.11n is fully backwards compatible to support mixed mode operation with 802.11a/b/g devices.

The backwards compatibility 802.11n provides is similar to the protection mechanism 802.11g

provides for 802.11b devices. 802.11n provides this protection mechanism in both bands, 2.4 GHz (for 802.11b/g devices) and 5 GHz (for 802.11a devices). As with 802.11g, the protection mechanisms kick in as soon as an access point hears a legacy device transmitting on the same channel. The legacy device does not have to be associated to the 802.11n access point; the access point just needs to hear it on the same channel. However, mixed mode operation reduces the overall throughput for 802.11n cells. Keeping legacy clients on a different channel and dedicating separate channels for 802.11n devices can help resolve this problem.

Frequency Bands and Channel Availability

As noted previously, 802.11n can operate in both the 2.4 GHz and 5 GHz frequency bands and is fully backwards compatible with current generation devices in each band. 802.11g and 802.11a use channels that are 20 MHz wide, while 802.11n can use 20 or 40 MHz channels when communicating with other 802.11n devices. The 2.4 GHz frequency band has a total of 3 non-overlapping 20 MHz channels, while the 5 GHz band has a total of 23 non-overlapping 20 MHz channels (for most countries). For 40 MHz 802.11n communications, this means the 2.4 GHz band offers only a single 40 MHz non-overlapping channel, while the 5 GHz band offers up to 802.11 non-overlapping 40 MHz channels. At a connection speed of 300 Mbps, a 2.4 GHz deployment can deliver aggregate data rates up to 450 Mbps (a single 40 MHz 802.11n channel at 300 Mbps and a single 20 MHz 802.11n channel at 150 Mbps or three 20 MHz channels at 150 Mbps), whereas the 5 GHz deployment can deliver 3.45 Gbps (11 high throughput 40 MHz 802.11n channels at 300 Mbps each and a single 20 MHz 802.11n channel at 150 Mbps).

Clearly, the 5 GHz UNII band is better suited for multi-cell RF deployments; however, there are two issues that need to be considered with the 5 GHz frequency band. The 5 GHz band's range is a little less than that of 2.4 GHz, and there are radar detection requirements within this band. The range issue can be easily overcome with the proper antenna selection and transmit power adjustment. The second issue requires the 802.11n radio chip set provide radar detection and dynamic frequency selection (DFS) capabilities. Many early 802.11n access points (even those that are Wi-Fi Draft-N certified) use radio chipsets that cannot detect some

newly-introduced radar pulses. The operation of these access points is limited to only those bands that do not require radar detection, significantly reducing the number of available channels and overall system throughput.

Adopting 802.11n

802.11n RF Network Planning

RF network design typically involves determining the number of access points, their placement, channel assignment, and power settings to deliver the required coverage and throughput for a given physical environment. Traditionally, RF network design has evolved from a series of physical site surveys, in some cases, to combining physical site surveys with predictive planning tools.

The RSSI-based RF planning tools of the past are of little use in predicting and characterizing system performance for an 802.11n network. Planning tools must take into account the multipath characteristics of a wide variety of building construction material like concrete, glass, dry-wall, office cubicles, warehouse racks, etc. In addition accurately characterizing multipath is extremely important for optimal MIMO performance. Multipath characterization includes analyzing the geometrical structure of site drawings representing the actual physical dimensions of the intended radio coverage area.

Because 802.11n uses MIMO transmissions with multiple data streams of varying modulation, this analysis of the in-building multipath characteristics must be combined with knowledge of the predicted RSSI in order to arrive at an accurate prediction of the system performance. Thus, merely predicting RSSI at any given point is of little or no value. 802.11n requires a system capable of combining the effects of MIMO with multiple spatial streams and the underlying RSSI.

When supporting existing 802.11b/g clients and simultaneously deploying 802.11n in the 5 GHz band, there is no easy way to insert new 802.11n access points without causing harmful co-channel interference. Even with a rip and replace migration, a one-to-one replacement of legacy access points with an 802.11n AP (without transmit power adjustment) is likely to cause interference issues given the improved receiver sensitivity of the new radio technology.

802.11n planning can be simplified by using planning tools designed specifically for 802.11n and that offer legacy deployment support. Customers should look for tools that drive RF heat-map algorithms based on the detailed RF characterization of the specific access point radio deployed. Predictions and coverage maps generated for generic access point models can provide misleading results, delivering inconsistent coverage and poor performance when deployed.

802.11n Security Issues

The high throughput and reliability benefits of 802.11n will likely drive a greater proportion of new LAN deployments to be completely wireless. However, the need for security becomes even greater as businesses move mission critical applications to wireless and make it a primary network. A comprehensive WLAN security policy enforced by a dedicated Wireless Intrusion Prevention System (WIPS) system provides the foundation for wired-to-wireless network transitions.

Part-time WIPS can utilize the wireless infrastructure for limited detection of wireless attacks and rogue devices. These systems use access points that can service WLAN clients on one channel and perform off-channel scanning to detect rogue devices and to capture threats for analysis. The access point forwards information to a WIPS server that compares them against a database of well known attack signatures to identify patterns and generate events and/or alarms. Although the throughput of a typical 802.11n access point is expected to be about 6 times that of a traditional 802.11a/b/g access point, with improvements in hardware (faster CPUs, increased memory), most access points can handle both data and WIPS functions.

Part-time WIPS is sub-optimal for monitoring 802.11n deployments. Time-slicing radio resources for performing off-channel scanning is likely to have a serious impact on latency-sensitive and bandwidth-intensive applications like voice and video. Optimizing the radio's configuration to spend more time servicing WLAN clients, with fewer off channel scans, can improve WLAN performance but leave more channels

unmonitored and open for attack for longer periods of time. The time-slicing problem becomes more acute with 802.11n, because each WLAN channel must now be monitored for attacks at both the 20 MHz and 40 MHz channel widths. This almost doubles the number of channels the sensor has to monitor. 802.11n access points optimized for high throughput spend fewer cycles monitoring. Leaving several channels unmonitored for long periods of time exposes the wireless network to serious threats.

A dedicated sensor radio that does not serve WLAN clients (or perform wireless bridging or meshing) is able to spend all its time scanning for rogues and attacks, thus providing a much higher level of wireless security. A dedicated sensor can intelligently adjust its frequency scanning pattern based on wireless activity, spending more time on channels of interest, actively terminating rogues and unauthorized wireless sessions, without being constrained to a single operational channel.

Indoor 802.11n MESH

Mesh technology enables access points to wirelessly connect to each other and transmit data. Mesh provides a great way to extend network coverage, typically within hard-to-wire outdoor deployments. A dual-radio access point design allows deploying one radio for WLAN client access while leaving the second radio for a mesh backhaul. In 802.11a/b/g access points, the 802.11a radio was typically used for mesh backhaul and provided a 54 Mbps connection to the network. With 802.11n, the data-rate of this backhaul link increases to 300 Mbps, making mesh a higher speed network than Fast Ethernet for interconnecting access points.

It is important to note that bandwidth alone is not all that is involved in replacing wired access points with wireless mesh connections. Typically, local area networks are segmented using Virtual LANs (VLANs), each with its own quality of service (QoS) and security requirements. To replace a wired connection with a wireless link, the mesh access point should support intelligent backhauls that are VLAN and QOS-aware, supporting WPA2 security on the wireless link.

Mesh resiliency is another important consideration. Any changes in the physical environment can significantly change RF performance. For example, the movement of goods within a warehouse could temporarily downgrade the RF link between two access points to unusable levels. Consequently, mesh APs must maintain multiple redundant backhaul connections to other access points and dynamically self-heal based on changing RF conditions.

With the correct set of access point configurations, 802.11n-based mesh provides a great opportunity to reduce the number of wired access points in an indoor WLAN and replace them with a mesh design that saves cabling costs and delivers a higher performance network.

The Wireless Enterprise

802.11n offers many benefits to an Enterprise customer. In addition to high-speed connectivity and better range for 802.11n clients, the new standard also delivers significant improvements in WLAN reliability and coverage for existing 802.11a/b/g deployments. The following are some examples of high-bandwidth applications that can take advantage of 802.11n:

- **Education** — Wireless access in high density areas, like classrooms and auditoriums, with fewer access points
- **Healthcare** — Bedside access to medical imaging and patient records
- **Manufacturing** — Wireless sensor networks and video surveillance
- **Retail** — Store-of-the-future applications, like smart carts with location-based advertisement streaming, wireless digital signage, and video surveillance

In high-multipath environments (e.g. manufacturing, warehouses, etc.), 802.11n will also deliver a more reliable wireless network with better coverage and fewer RF dead spots. All things being equal, 802.11n is expected to deliver a significantly better wireless experience compared to the existing WLAN technologies and standards. With 802.11n, wireless networking catches up with wired networking in the area of predictability and reliability and exceeds Fast Ethernet in speed and throughput.

The high-bandwidth and improved reliability of 802.11n, when combined with the roaming benefits of wireless connectivity, will fundamentally transform the end-user mobility experience, impacting end-user behavior and business processes within the enterprise. With 802.11n, wireless will no longer be the network of convenience, and in many cases, will be the primary edge network running the most critical applications. For the enterprise CIO, 802.11n lowers the cost of ownership for deploying, supporting, and managing a WLAN. In addition, business applications now can be defined for a single mobile device platform without the need for developing and porting to multiple platforms. Employees will have reliable "anytime, anywhere" access to all business applications. Customers are better served by employees with seamless access to business information systems, regardless of location or device type. Traditional wireless networking enabled enterprise mobility in select verticals, allowing key applications to be wireless-enabled. 802.11n can unwire almost any business application used in the enterprise today and support seamless, location-independent information access for the workforce, reducing infrastructure/platform costs while enhancing productivity levels and improving customer satisfaction.

802.11n enables, "The Wireless Enterprise"; an IT-enabled business strategy that leverages wireless systems and applications to drive dramatic improvements in workforce productivity and customer satisfaction. From package delivery to warehouse and store operations, wireless technology has already improved business processes for verticals and retail. In the next 3-5 years, high-speed, reliable and user-friendly wireless technology will revolutionize the way businesses communicate, collaborate, and service customers. The future of the network is wireless. It's time to cut the wire and become a smarter, more efficient wireless enterprise.



MOTOROLA

motorola.com

Part number WP-802.11nDEM. Printed in USA 02/09. MOTOROLA and the Stylized M Logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners. ©2009 Motorola, Inc. All rights reserved. For system, product or services availability and specific information within your country, please contact your local Motorola office or Business Partner. Specifications are subject to change without notice.