Meraki White Paper:
802.11n Technology

February 2011

This document describes the benefit of adopting 802.11n technology in a Meraki wireless network.

660 Alabama St.
San Francisco, California 94110

Phone: +1 415 632 5800
Fax: +1 415 632 5899

Table of Contents

1 Introduction to Meraki 802.11n

This chapter summarizes the benefits of a wireless LAN (WLAN) network for an introduction to Multiple-In Multiple-Out (MIMO) technology

1.1 Technology Overview

The IEEE 802.11n standard (formally known as IEEE 802.11n-2009) offers several technical benefits over previous technology generations, which result in improved throughput to 802.11n-based clients, as well as greater reliability for legacy 802.11a/b/g clients.

802.11n is much more than just a new radio for 802.11. In addition to providing higher bit rates (as was done in 802.11a, b, and g), 802.11n makes dramatic changes to the basic frame format that is used by 802.11 devices to communicate with each other. This section will describe the changes incorporated in 802.11n, including MIMO, radio enhancements, and MAC enhancements.

Environmental characteristics and network density plays a significant role in the ultimate performance of a network. In well-designed networks, each access point can serve well over 150 Mbps of TCP throughput to clients using 802.11n technology, and multiple radios can operate simultaneously to provide several gigabits of throughput.

Technique	Effect	Clients Affected
Packet Aggregation	Multiple TCP packets are combined together in a single MAC layer frame, to reduce overhead from headers	802.11n
Block Acknowledgements	Multiple packets can be acknowledged as a block at the link layer, reducing the amount of airtime spent on low-speed ACK frame	802.11n
Channel Bonding and Coding Schemes	Utilize 40 MHz wide channel bandwidths and high density modulation to improve line rates from 54 Mbit/sec in 802.11g to 600 Mbit/sec in 802.11n	802.11n
Spatial Multiplexing	Simultaneously send multiple streams of data and decode with multiple receivers, to increase channel capacity	802.11n

Technique	Effect	Clients Affected
Multi-ratio Combining (MRC) and Improve Receive Sensitivity	Combine data from subcarriers on each receive antenna, to improve client receive sensitivity for non-11n clients	802.11a/b/g
Transmit Beamforming	Use physical layer phase delays between antennas to electronically steer transmissions when sending a single spatial stream to non-11n clients.	802.11a/b/g

1.2 Understanding MIMO

Multiple-input multiple-output is the heart of 802.11n. This technical discussion of MIMO provides a basis for understanding how 802.11n can reach data rates of 600 Mbps.

1.2.1 Radio Operation Basics

To understand the improvement brought by MIMO technology, it is important to understand some of the basics that determine how well a traditional radio operates. In a traditional, single-input single-output radio, the amount of information that can be carried by a received radio signal depends on the amount by which the received signal strength exceeds the noise at the receiver, called the signal-to-noise ratio, or SNR. SNR is typically expressed in decibels (dB). The greater the SNR, the more information that can be carried on the signal and be recovered by the receiver.

To understand this situation, imagine the analogy of an eye as the receiver. Is the eye able to detect whether a table lamp is on or off in the house next door? In this analogy, ambient light is the noise. At night, detecting that the lamp is on or off is quite easy. However, in full daylight, it is much more difficult to make the same determination, because the ambient light is much brighter, and the tiny amount of additional light from the lamp can be undetectable.

Light, like a radio wave, disperses uniformly from its source. The farther the receiver is from the source, the less power is received from the source. In fact, the amount of power received decreases more rapidly than the square of the distance from the source. Noise, unfortunately, is often constant in the environment, due to both natural and man-made causes.

So, returning to the table lamp example, when it is too bright to determine if the lamp next door is on or off, it might be possible to make that determination from just outside the neighbor’s window. Alternatively, it might be possible to make the determination if the neighbor changed out the 40 watt bulb with a 150 watt bulb. In both cases, the SNR increases, in the first case, because the distance to the source is reduced, and in the second, because the power of the transmitter is increased.

Once the minimum SNR is achieved to allow information to be exchanged at the desired rate, any additional SNR can be likened to “money in the bank.” That additional SNR can be spent on increasing the information rate, increasing the distance, or a little bit of both. However, the same dB cannot be spent more than once, just as the same dollar cannot be spent more than once (at least, not without encountering some unpleasant consequences).

All this is background to understand the improvements that MIMO technology brings to 802.11.

1.2.2 MIMO Technology: Beamforming

MIMO technology takes advantage of other techniques to improve the SNR at the receiver. One technique is *transmit beamforming*. When there is more than one transmit antenna, it is possible to coordinate the signal sent from each antenna so that the signal at the receiver is dramatically improved. This technique is generally used when the receiver has only a single antenna and when there are few obstructions or radio-reflective surfaces, e.g., open storage yards.

To understand transmit beamforming, consider a radio signal as a wave shape, with a wave length that is specific to the frequency of the signal. When two radio signals are sent from different antennae, these signals are added together at the receiver’s antenna (see Figure 1). Depending on the distance that each radio signal travels, they are very likely to arrive at the receiver out of phase with each other. This difference in phase at the receiver affects the overall signal strength of the received signal. By carefully adjusting the phase of the radio signals at the transmitter, the received signal can be maximized at the receiver, increasing SNR. This is what transmit beamforming does: It effectively focuses the transmitters on a single receiver, as shown in Figure 2.

Figure 1: Destructive Interference

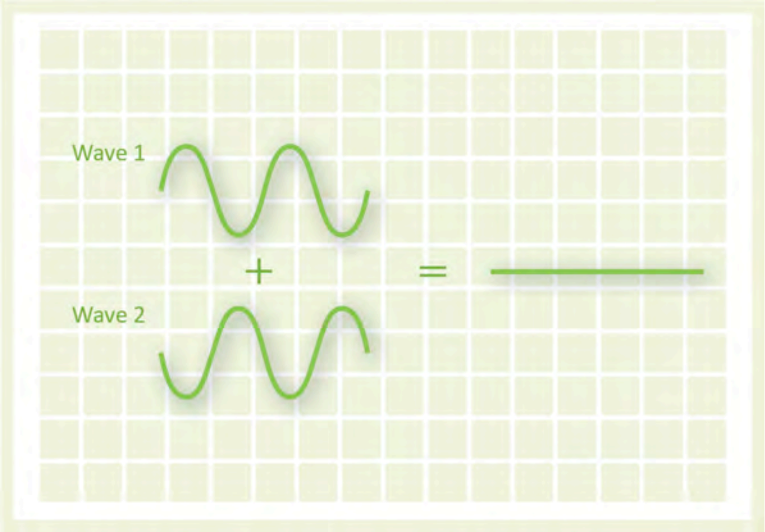

Figure 2: Constructive Interference with Transmit Beamforming

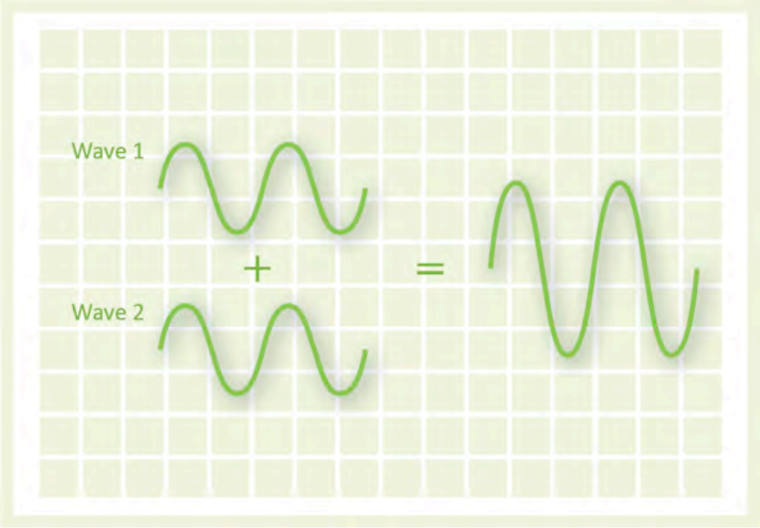

Transmit beamforming cannot easily be done at the transmitter without information from the receiver about the received signal. This feedback is available only from 802.11n devices, and not from 802.11a, b, or g devices. To maximize the signal at the receiver, feedback from the receiver must be sent to the transmitter so that the transmitter can tune each signal it sends. This feedback is not immediate and is only valid for a short time. Any physical movement by the transmitter, receiver, or elements in the environment will quickly invalidate the parameters used for beamforming. The wavelength for a 2.4 GHz radio is only 120mm, and only 55mm for a 5 GHz radio. Thus, a normal walking pace of 1 meter per second will rapidly move the receiver out of the spot where the transmitter's beamforming efforts are most effective.

Transmit beamforming is useful only when transmitting to a single receiver. It is not possible to optimize the phase of the transmitted signals when sending broadcast or multicast transmissions. For this reason, in general networking applications, the utility of transmit beamforming is somewhat limited, providing improved SNR at the receiver for only those transmissions that are sent to that receiver alone. Transmit beamforming can increase the data rate available at greater distances from the AP. But, it does not increase the coverage area of an access point, since that is determined, in large part, by the ability to receive the beacons from the access point. Beacons are a broadcast transmission that does not benefit from transmit beamforming.

1.2.3 MIMO Technology: Multipath or Spatial Diversity

In typical indoor WLAN deployments, (e.g., offices, hospitals, and warehouses) the radio signal rarely takes the shortest and most direct path from the transmitter to the receiver. This is because there is rarely “line of sight” between the transmitter and the receiver. Often there is a cube wall, door, or other structure that obscures the line of sight. All of these obstructions reduce the strength of the radio signal as it passes through them. Luckily, most of these environments are full of surfaces that reflect a radio signal as well as a mirror reflects light.

Imagine that all of the metallic surfaces, large and small, that are in an environment were actually mirrors. Nails and screws, door frames, ceiling suspension grids, and structural beams are all reflectors of radio signals. It would be possible to see the same WLAN access point in many of these mirrors simultaneously. Some of the images of the access point would be a direct reflection through a single mirror. Some images would be a reflection of a reflection. Still others would involve an even greater number of reflections. This phenomenon is called *multipath* (see Figure 3).

Figure 3: Multipath Interference

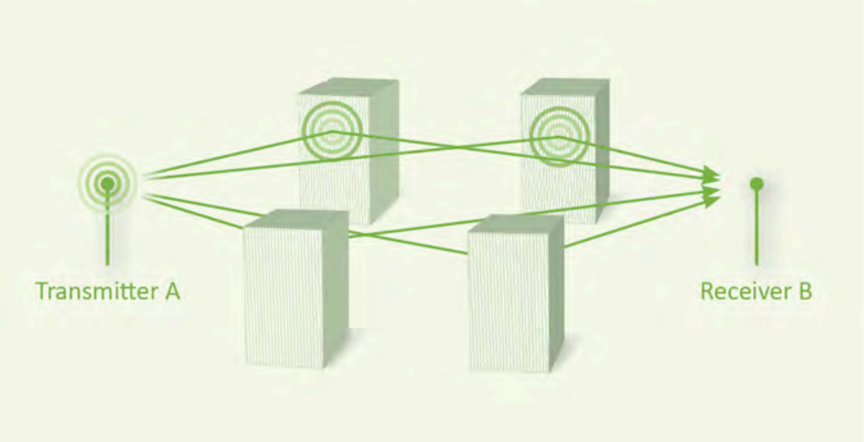

When a signal travels over different paths to a single receiver, the time that the signal arrives at the receiver depends on the length of the path it traveled. The signal traveling the shortest path will arrive first, followed by copies or echoes of the signal slightly delayed by each of the longer paths that the copies traveled. When traveling at the speed of light, as radio signals do, the delay between the first signal to arrive and its copies is very small, only nanoseconds. (A rule of thumb for the distance covered at the speed of light is roughly one foot per one nanosecond.) This delay is enough to be able to cause significant degradation of the signal at a single antenna because all the copies interfere with the first signal to arrive.

A MIMO radio sends multiple radio signals at the same time and takes advantage of multipath. Each of these signals is called a *spatial stream*. Each spatial stream is sent from its own antenna, using its own transmitter. Because there is some space between each of these antennae, each signal follows a slightly different path to the receiver. This is called *spatial diversity*. Each radio can also send a different data stream from the other radios. The receiver has multiple antennas as well, each with its own radio. Each of the receive radios independently decode the arriving signals (see Figure 4). Then, each radio's received signal is combined with the signals from the other receive radios. With a lot of complex math, the result is a much better receive signal than can be achieved with either a single antenna or even with transmit beamforming. One of the two significant benefits of MIMO is that it dramatically improves the SNR, providing more flexibility for the WLAN system designer.

Figure 4: Spatial Multiplexing

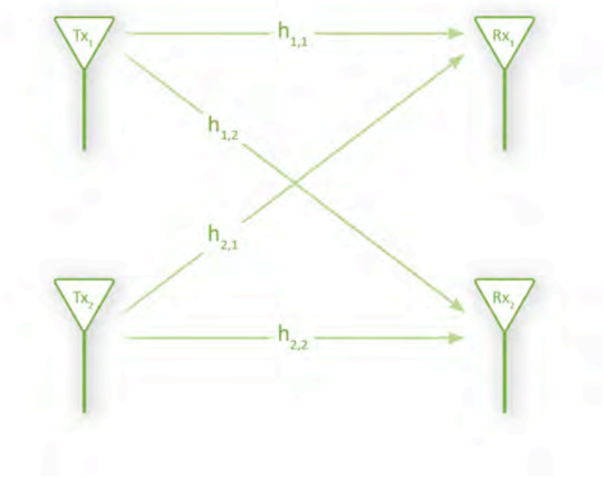

MIMO systems are described using the number of transmitters and receivers in the system—for example, 2×1 is "two by one," meaning two transmitters and one receiver. 802.11n defines a number of different combinations for the number of transmitters and the number of receivers, from 2×1, equivalent to transmit beamforming, up to 4×4. Each additional transmitter or receiver in the system increases the SNR. However, the incremental gains from each additional transmitter or receiver diminish rapidly. The gain in SNR is large for each step from 1×1 to 2×1 to 2×2, the improvement with 3×3 is not quite as large, and beyond 3x3 is more moderate. The use of multiple transmitters provides the second significant benefit of MIMO: the ability to use each spatial stream to carry its own information, providing dramatically increased data rates.

2 802.11n Radio Enhancements

In addition to MIMO technology, 802.11n makes a number of additional changes to the radio to increase the effective throughput of the WLAN. The most important of these changes are increased channel size, higher modulation rates, and reduced overhead. This section will describe each of these changes and the effect they have on WLAN throughput.

2.1 Physical Layer Enhancements

2.1.1 20 MHz and 40 MHz Channels

The original 802.11 direct sequence radio and the 802.11b extension to the base standard use a radio channel spacing that is 22 MHz wide. 802.11a and 802.11g use 20 MHz wide radio channel spacing. Because 802.11g is an extension to 802.11b, 802.11g spaces its channels just as 802.11b does, every 22 MHz. The size, or bandwidth, of the radio channel is an important measure of the efficiency of the radio. This is called the *spectral efficiency* and is measured in bits per Hertz. The spectral efficiency of 802.11b is one-half the bits per Hertz (for example, 11 Mbps in 22 MHz). 802.11a and 802.11g have higher spectral efficiency, as much as 2.7 bits per Hertz at 54 Mbps.

Using exactly the same technology as 802.11a and 802.11g, some proprietary WLAN systems are available that provide up to 108 Mbps. These proprietary systems use a simple technique to double the data rate of 802.11a and 802.11g. They use two channels at the same time. This is called *channel bonding*. With channel bonding, the spectral efficiency is the same as 802.11a and 802.11g, but the channel bandwidth is twice as great. This provides a simple way of doubling the data rate.

Figure 5: 20MHz and 40MHz channels

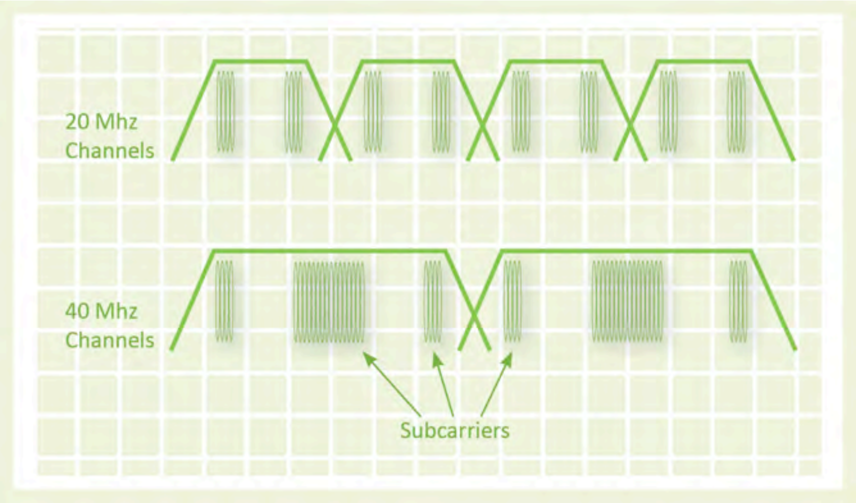

802.11n uses both 20 MHz and 40 MHz channels. Like the proprietary products, the 40 MHz channels in 802.11n are two adjacent 20 MHz channels, bonded together. When using the 40 MHz bonded channel, 802.11n takes advantage of the fact that each 20 MHz channel has a small amount of the channel that is reserved at the top and bottom, to reduce interference in those adjacent channels. When using 40 MHz channels, the top of the lower channel and the bottom of the upper channel do not have to be reserved to avoid interference. These small parts of the channel can now be used to carry information. By using the two 20 MHz channels more efficiently in this way, 802.11n achieves slightly more than doubling the data rate when moving from 20 MHz to 40 MHz channels (see Figure 5).

2.1.2 Higher Modulation Rates

The original 802.11 direct sequence radio transmitted a *symbol* to represent each bit (or set of bits) sent from a transmitter to a receiver. Each symbol lasted one microsecond. A symbol consisted of a fixed series of 11 *chips*. Each chip was modulated on the radio signal using a phase shift key (PSK) technique. For the 1 Mbps data rate, a single symbol was sent using binary PSK every microsecond. The 2 Mbps rate sent two symbols each microsecond, using quaternary (4-phase) PSK (QPSK). 802.11b extended the direct sequence radio by coding more bits into each symbol, while continuing the use of the QPSK modulation method. This allowed the extension of data rates to 11 Mbps.

802.11a and 802.11g changed the way information is transmitted on the radio signal. These standards adopted a method called *orthogonal frequency division multiplexing* (OFDM). OFDM divides a radio channel into a large number of smaller channels, each with its own *subcarrier* signal (see Figure 5 above). Each of these carrier signals is able to convey information independent of all the other carrier signals. It is roughly the same as having a group of independent radios bunched together.

For 802.11a and 802.11g, a symbol lasts 4 microseconds, including an 800 nanosecond *guard interval*. For the highest data rate, 54 Mbps, each symbol carries 216 data bits. These data bits are spread out over 48 subcarriers. In addition, there are 72 error-correction bits sent in each symbol at 54 Mbps, for a total of 288 bits in the symbol. To pack this many bits on each subcarrier, the subcarrier is modulated using 64 QAM (Quadrature Amplitude Modulation), 16 times the highest modulation rate of 802.11b. This means that each subcarrier is able to carry 6 bits (a combination of data and error correction bits).

802.11n continues to use OFDM and a 4-microsecond symbol, similar to 802.11a and 802.11g. However, 802.11n increases the number of subcarriers in each 20 MHz channel from 48 to 52. This marginally increases the data rate to a maximum of 65 Mbps, for a single-transmit radio. 802.11n provides a selection of eight data rates for a transmitter to use and also increases the number of transmitters allowable to four. For two transmitters, the maximum data rate is 130 Mbps. Three transmitters provide a maximum data rate of 195 Mbps. The maximum four transmitters can deliver 260 Mbps. In total, 802.11n provides up to 32 data rates for use in a 20 MHz channel.

When using 40 MHz channels, 802.11n increases the number of subcarriers available to 108. This provides a maximum data rate of 135 Mbps, 270 Mbps, 405 Mbps, and 540 Mbps for one through four transmitters, respectively. Similarly, there are eight data rates provided for each transmitter, 32 in total, for the 40 MHz channel.

802.11n can also use a short guard interval that is 400 nanoseconds long, instead of 800 nanoseconds. This slightly increases the maximum data rates, for example in 40 MHz channels, to 150 Mbps per transmitter. A four-transmitter 802.11n radio operating with 40 MHz channels and using the short guard interval can therefore deliver a maximum of 600 Mbps.

The rates described so far use the same modulation (*equal* modulation) on all of the subcarriers, for example, all the subcarriers use QPSK or 64 QAM. This is the same as 802.11a and 802.11g. 802.11n adds the ability to modulate different spatial streams using different modulation methods—that is, some spatial streams use QPSK, some other spatial streams use 16 QAM, and yet other spatial streams use 64 QAM. This dramatically increases the number of data rates available to be used. In fact, there are dozens more possible data rates, using this *unequal modulation* method. However, it is unlikely that many practical implementations would be able to take advantage of this method, as it requires a significant amount of feedback from the receiver to the transmitter to identify the individual spatial streams that must use each of the different modulation methods.

2.1.3 Maximal-ratio Combining (MRC)

The 802.11n standard includes the ability for the receiver to combine the received signals from multiple antennas to reassemble a single spatial stream. Multipath echoes in an environment can lead to frequency selective fading, in which certain subcarriers within a 20 MHz or 40 MHz signal are stronger than others. Maximal-ratio combining (MRC) enables the receiver to correlate the signal reception from multiple antennas and select the strongest of each antenna before decoding a particular subcarrier.

The result is illustrated in Figure 6: the aggregate signal of several antennas offers a stronger, more consistent result than any individual antenna. This effectively increases the receive sensitivity, solely through the use of digital signal processing.

Figure 6: Maximal Ratio Combining (MRC)

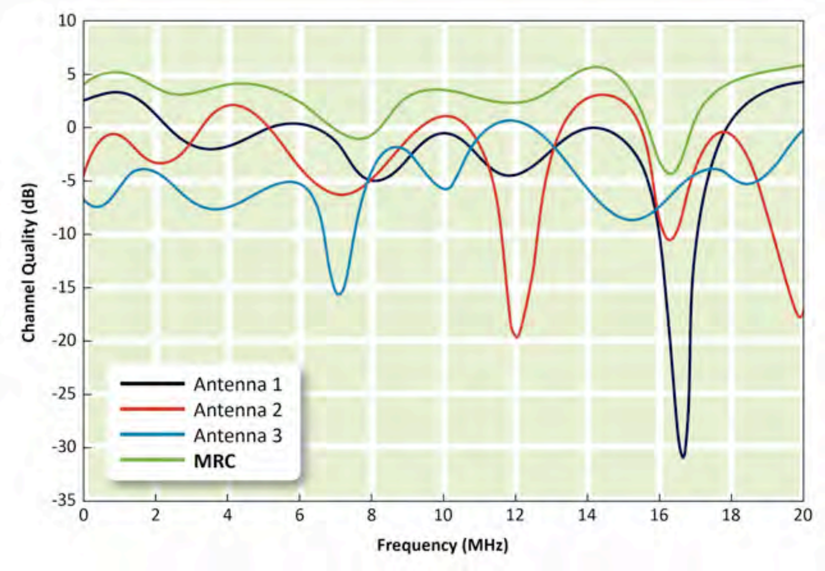

2.2 MAC Enhancements

There is only so much improvement that can be obtained in 802.11 by increasing the data rate of the radio. There is a significant amount of fixed overhead in the MAC layer protocol, and in the interframe spaces and acknowledgements of each frame transmitted, in particular. At the highest of data rates, this overhead alone can be longer than the entire data frame. In addition, contention for the air and collisions also reduce the maximum effective throughput of 802.11. 802.11n addresses these issues by making changes in the MAC layer to improve on the inefficiencies imposed by this fixed overhead and by contention losses.

2.2.1 Aggregation

Every frame transmitted by an 802.11 device has fixed overhead associated with the radio preamble and MAC frame fields that limit the effective throughput, even if the actual data rate was infinite (see Figure 7).

Figure 7: Frame Overhead

To reduce this overhead, 802.11n introduces *frame aggregation*. Frame aggregation is essentially putting two or more frames together into a single transmission. 802.11n introduces two methods for frame aggregation: Mac Service Data Units (MSDU) aggregation and Message Protocol Data Unit (MPDU) aggregation. Both aggregation methods reduce the overhead to only a single radio preamble for each frame transmission (see Figure 8).

Figure 8: Aggregated Frame

Because multiple frames are now sent in a single transmission, the number of potential collisions and the time lost to backoff is significantly reduced. The maximum frame size is increased in 802.11n, as well, in order to accommodate these large, aggregated frames. The maximum frame size is increased from 4 KB to 64 KB. One limitation of frame aggregation is that all the frames that are aggregated into a transmission must be sent to the same destination; that is, all the frames in the aggregated frame must be addressed to the same mobile client or access point. Another limitation is that all the frames to be aggregated have to be ready to transmit from the client or access point at the same time, potentially delaying some frames to wait for additional frames, in order to attempt to send a single aggregate frame. A third limitation of aggregation is that the maximum frame size that can be successfully sent is affected by a factor called *channel coherence time*. Channel coherence time depends on how quickly the transmitter, receiver, and other items in the environment are moving. The faster things are moving the smaller the maximum frame size can be as the data rate is reduced, i.e., the time for the transmission must be less than the channel coherence time.

There are slight differences in the two aggregation methods that result in differences in the efficiency gained. These two methods are described here.

2.2.2 MAC Service Data Units Aggregation

MSDU aggregation is the more efficient of the two aggregation methods. It relies on the fact that an access point receives frames from its Ethernet interface, to be translated to 802.11 frames and then transmitted to a mobile client. Similarly, most mobile client protocol stacks create an Ethernet frame, which the 802.11 driver must translate to an 802.11 frame before transmission. In both these cases, the “native” format of the frame is Ethernet, and it is then translated to 802.11 format for transmission.

Theoretically, MSDU aggregation allows frames for many destinations to be collected into a single aggregated frame for transmission. Practically, however, MSDU aggregation collects Ethernet frames for a common destination, wraps the collection in a single 802.11 frame, and then transmits that 802.11-wrapped collection of Ethernet frames (see Figure 9). This method is more efficient than MPDU aggregation, because the Ethernet header is much shorter than the 802.11 header.

Figure 9: MSDU Frame Aggregation

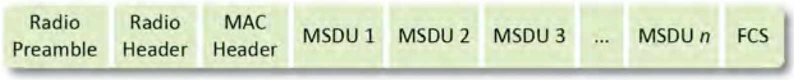

For a mobile device, the aggregated frame is sent to the access point, where the constituent Ethernet frames are forwarded to their ultimate destinations. For an access point, all of the constituent frames in the aggregated frame must be destined to a single mobile client, since there is only a single destination in each mobile client.

With MSDU aggregation, the entire, aggregated frame is encrypted once using the security association of the destination of the outer 802.11 frame wrapper. A restriction of MSDU aggregation is that all of the constituent frames must be of the same quality-of-service (QoS) level. It is not permitted to mix voice frames with best-effort frames, for example.

2.2.3 MAC Protocol Data Units Aggregation

MPDU aggregation is slightly different from MSDU aggregation. Instead of collecting Ethernet frames, MPDU aggregation translates each Ethernet frame to 802.11 format and then collects the 802.11 frames for a common destination. The collection does not require a wrapping of another 802.11 frame, since the collected frames already begin with an 802.11 MAC header (see Figure 10).

Figure 10: MPDU Aggregation

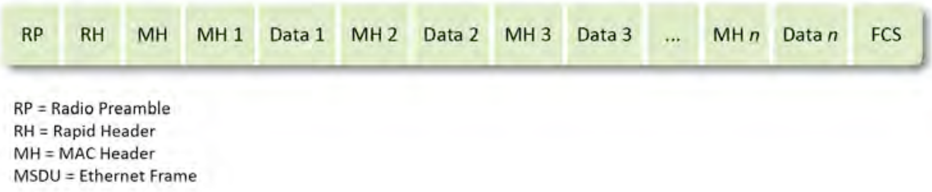

MPDU aggregation does require that all the 802.11 frames that constitute the aggregated frame have the same destination address. However, this results in the same behavior as MSDU aggregation, since the destination of all frames sent by a mobile client is that client's access point, where the 802.11 frames are translated to Ethernet and forwarded to the ultimate destination. Similarly, the destination of any frame sent by the access point is a single mobile client.

With MPDU aggregation, it is possible to encrypt each constituent frame independently, using the security association for each individual 802.11 destination address. This does not have any effective difference from the encryption done in MSDU aggregation, as all frames sent by a mobile client are encrypted using the security association for the access point, and all frames sent by the access point are encrypted using the security association for the single mobile client that is the intended recipient of the frame.

Similar to MSDU aggregation, MPDU aggregation requires that all of the constituent frames be of the same QoS level.

The efficiency of the MPDU aggregation method is lower than that of the MSDU aggregation method, because of the extra overhead of the individual 802.11 frame headers for each constituent frame. The efficiency is further reduced when the encryption is used. Encryption adds overhead to each of the constituent frame in MPDU aggregation, where MSDU aggregation incurs overhead for a single encryption of the outer 802.11 wrapper.

2.3 Block Acknowledgement

For the 802.11 MAC protocol to operate reliably, each of the frames transmitted an individual address, i.e., not multicast or broadcast frames, is immediately acknowledged by the recipient. MSDU aggregation requires no changes to this operation. The aggregated frame is acknowledged, just as any 802.11 frame is acknowledged. This is not the case for MPDU aggregation. For MPDU aggregation, each of the individual constituent 802.11 frames must be acknowledged. The mechanism to deal with this requirement that 802.11n introduces is called *block acknowledgement*.

Block acknowledgement compiles all the acknowledgements of the individual constituent frames produced by MPDU aggregation into a single frame returned by the recipient to the sender. This allows a compact and rapid mechanism to implement selective retransmission of only those constituent frames that are not acknowledged. In environments with high error rates, this selective retransmission mechanism can provide some improvement in the effective throughput of a WLAN using MPDU aggregation over that of one using MSDU aggregation, because much less is retransmitted when an error affects some of the constituent frames of an MPDU aggregated frame as compared to an MSDU aggregated frame.

2.4 Lower Overhead: Reduced Interframe Space

When aggregation of frames is not possible, 802.11n provides a mechanism to reduce the overhead involved with transmitting a stream of frames to different destinations. This mechanism reduces the interframe space between receiving a frame, typically an acknowledgement frame, and sending a subsequent frame. The 802.11e extension for quality of service added the ability for a single transmitter to send a burst of frames during a single, timed *transmit opportunity*. During the transmit opportunity, the sender does not need to perform any random backoff between transmissions, separating its frames by the smallest allowable interframe space, the short interframe space (SIFS).

802.11n improves on this mechanism, reducing the overhead between frames, by specifying an even smaller interframe space, called the *reduced interframe space* (RIFS). RIFS cuts down further on the dead time between frames, increasing the amount of time in the transmit opportunity that is occupied by sending frames. The one unfortunate aspect of using RIFS is that it is restricted to being used only in *greenfield deployments*-that is, only deployments where there are no legacy 802.11a, b, or g devices in the area.

2.5 Backward Compatibility

Compatibility with existing 802.11a, b, and g devices is a critical issue addressed in 802.11n. Just as 802.11g provides a protection mode for operation with 802.11b devices, 802.11n has a number of mechanisms to provide backward compatibility with 802.11 a, b, and g devices, allowing these devices to understand the information necessary to allow 802.11n devices to operate in the same area.

For quite a long time, 802.11n will need to operate in the presence of legacy 802.11a, b, and g devices. This mixed-mode operation will continue until all the devices in an area have been upgraded or replaced with 802.11n devices. The mixed-mode protection mechanism for 802.11n is quite similar to the protection mechanism of 802.11g.

Like 802.11g, 802.11n transmits a signal that cannot be decoded by devices built to an earlier standard. To avoid descending into absolute chaos in the presence of massive interference and collisions, 802.11n operating in mixed mode transmits a radio preamble and signal field that can be decoded by 802.11a and 802.11g radios (see Figure 11). This provides enough information to the legacy radios to allow them to indicate that there is another transmission on the air and how long that transmission will last. Following the legacy preamble and signal field, the 802.11n device sends the remaining information using 802.11n rates and its multiple spatial streams, including an 802.11n preamble and signal field.

Figure 11: Legacy Frame Format

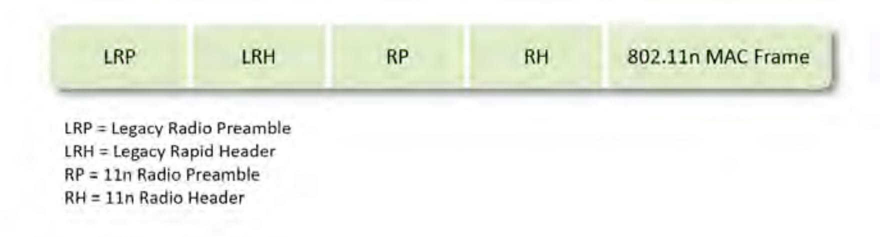

In addition to the legacy preamble and signal field, it can also be necessary to use additional protection mechanisms provided by 802.11g to allow the MAC in legacy devices to correctly determine when it is allowed to transmit and when it must perform backoff before transmission. The mechanism provided by 802.11g and utilized by 802.11n when either 802.11g or 802.11a devices are present is the CTS-to-self mechanism. CTS-to-self allows the 802.11n device to transmit a short CTS frame, addressed to itself, that includes the timing information necessary to be communicated to the neighboring legacy MACs that will protect the 802.11n transmission that will follow. The CTS frame must be transmitted using one of the legacy data rates that a legacy device will be able to receive and decode.

The cost of this additional legacy preamble and signal field, as well as the CTS-to-self, is more overhead on every 802.11n transmission. This reduces the benefits of all the 802.11n improvements, resulting in significantly lower effective throughput by 802.11n devices in mixed environments. Similar to 802.11g, legacy devices do not need to be associated to the same access point as an 802.11n device to require the use of protection mechanisms. If there are legacy devices on the same channel on any nearby access points, this will cause protection mechanisms to be invoked as well.

It can be expected that protection mechanisms will be in use in the 2.4 GHz band (802.11b and 802.11g) until nearly every legacy device has disappeared. This is because there are too few channels available in that band to effectively overlay pure 802.11n WLANs in the same areas as legacy 2.4 GHz WLANs. Given the larger number of channels available in the 5 GHz band in many countries, it is possible that two completely separate WLANs could be operating in the same area, with 802.11a operating on one set of channels and 802.11n operating on a different, nonintersecting set of channels. This would result in 802.11n operating in pure high-throughput (greenfield) mode, achieving the highest effective throughput offered by this new standard.

2.6 Summary of 802.11n Technology

To summarize the benefits of 802.11n technology, it is simplest to say that there are two major areas of improvement over previous 802.11 devices. The first area of improvement is in the use of MIMO technology to achieve greater SNR on the radio link. The second area of improvement is in the greater efficiencies in both radio transmissions and the MAC protocol. These improvements translate into benefits in three areas: reliability, predictable coverage, and throughput.

Reliability: Greater SNR on the radio link translates directly to more reliable communication, often at higher data rates. Higher SNR means that more interference is needed to corrupt a transmission. This means greater client densities can be supported.

Predictable Coverage: The use of multiple spatial streams provided by MIMO technology means that there will be fewer dead spots in a coverage area. Areas that previously suffered from destructive multipath interference now make use of that same multipath effect to provide robust communication.

Throughput: The efficiency improvement in 802.11n provides a greater transfer of the high bit rates of the 802.11n radio to effective throughput seen by actual applications, at least in greenfield deployments. Even in mixed-mode deployments with legacy 802.11 devices, 802.11n will provide greater effective throughput, although significantly less than the greenfield mode.

3 Migration to 802.11n

The migration to 802.11n has already begun. 802.11n client devices, beginning with laptops, are already available, and new client devices continue to appear. In larger devices, 802.11n has already become the default WLAN adapter. These client devices are completely compatible with existing 802.11a, b, and g access points and will operate just as existing devices do today, but can take advantage of 802.11n infrastructure where it has been deployed.

3.1 Planning

There are several areas to consider when planning the migration of a network to support 802.11n. Because of the higher speeds and greater power requirements of 802.11n access points, planning the migration needs to take into account more than just the access point.

3.1.1 Radio Bands

802.11n operates in both the 2.4 GHz (802.11b and g) and 5 GHz (802.11a) radio bands. Planning for each of the radio bands should be done independently, because of the constraints that are sometimes very different for each band.

The 2.4 GHz band is no more than 100mm wide, and often much less than that in many countries. The same channelization that is used for 802.11b and 802.11g can be used for 802.11n operating in this band. However, the use of the 40 MHz mode of operation of 802.11n is not recommended in this band, because a significant portion of the band will suffer from interference from a single 40 MHz transmitter. In addition, it is required that the second 20 MHz channel, concatenated with the original 20 MHz channel to form the 40 MHz channel, must be free of any legacy transmissions. This drastically reduces the chance that any 40 MHz operations will be feasible in this band.

In much of the world where it is typical to utilize three non-overlapping channels in this band, a single 40 MHz access point will present a significant challenge to developing a channel plan that will provide adequate capacity in most enterprises. Even when all legacy 802.11b and g devices are removed from the band, it will be difficult to deploy access points utilizing the 40 MHz channels in this band. There is just not enough bandwidth available to even begin to duplicate the three non-overlapping channels of the legacy layout.

The 5 GHz band has been opened up significantly in much of the world, due to recent changes by many regulatory agencies. There are significantly more channels available in the 5 GHz band than in the 2.4 GHz band. The larger number of channels in this band makes planning the deployment of an 802.11n network much simpler, even while allowing for 40 MHz operation.

There are at least two possible ways to migrate to 802.11n in the 5 GHz band. The first way is to replace individual legacy access points with 802.11n access points as budget allows and user demand for additional capacity dictates. This gradual migration can be accomplished over a planned period of time or as the need arises. This migration method would have the new 802.11n access point operating on the channel of the legacy access point it replaced. The new 802.11n access point would support 802.11n clients, as well as legacy 802.11a clients. It would operate in mixed mode, providing protection for the legacy 802.11a clients. Eventually, as the last legacy access point is replaced and the last legacy client is retired, the entire set of new 802.11n access points could be switched to operate in greenfield mode.

The second way to migrate to 802.11n would be to reassign the channels on some of the legacy access points to free a set of channels that could be used for 802.11n exclusively. Then as budget allows and demand dictates, new 802.11n access points would be *added* to the existing WLAN deployment, operating in parallel in overlapping areas with the legacy access points. The new 802.11n access points, however, would support only 802.11n devices and be able to operate in the greenfield mode, providing the greatest benefits of the new standard. Eventually, as 802.11n access points cover an entire area served by legacy access points, 802.11n clients would have the ability to operate in greenfield mode everywhere, while the legacy access points still provide service to the legacy clients. Once the last legacy client is retired, the legacy access points can also be retired.

3.1.2 Wired Infrastructure Stresses

Today’s dual-band access points can theoretically put a load on their Ethernet connections of as much as 54 Mbps. Practically, however, due to the inefficiencies of the 802.11 protocol, they top out at a peak load of 20 Mbps.

802.11n access points can demand much more of their Ethernet connection. With the higher bit rates on the air and the improved efficiency of the protocol, it is possible that a single dual-band 802.11n access point supporting a 20 MHz channel in the 2.4 GHz band and a 40 MHz channel in the 5 GHz band can place a peak demand on its Ethernet connection of more than 200 Mbps. Obviously, this is greater than a single 100 Mbps Ethernet connection can support.

For this reason, planning for a migration to support 802.11n should also include planning to upgrade the edge Ethernet switching capabilities to support a 1-Gbps connection to each 802.11n access point. This will eliminate any bottlenecks that might occur in areas of high-capacity demand by the 802.11n clients.

3.1.3 Access Point Deployment

Planning the positioning of the 802.11n access points should also be considered during migration planning. If the migration plan is to gradually replace the existing legacy access points, there is no further deployment planning necessary. However, if the new 802.11n access points are to be deployed in a new installation or alongside an existing deployment, it is possible to use the increased SNR provided by 802.11n to cover greater areas per access point, although at a cost of reducing the overall capacity of the resulting 802.11 WLAN. Earlier, SNR was discussed using the analogy to “money in the bank.” SNR can be used for either increased data rate, increased range, or a little of both, but cannot be used for the maximum of both at the same time.

3.2 Conclusions

There is no reason to avoid installing a new 802.11n WLAN or migrating an existing WLAN to support 802.11n. Some care, though, should be taken when selecting the 802.11n equipment to install.

802.11n has the ability to dramatically increase the capacity of a WLAN and the effective throughput of every client. The time to begin moving to this new standard is as soon as it is necessary to add a new access point, in order to address the demand for additional capacity in the WLAN, and bring Ethernet-level speeds to the wireless client.