# ON Semiconductor

## Is Now

# onsemi™

## To learn more about onsemi™, please visit our website at
## www.onsemi.com

**ON Semiconductor®**

# An Introduction to 802.11ac

Originally published in Sep. 2011 under the company name Quantenna® Communications, Inc.

# An Introduction to 802.11ac

## Introduction

802.11ac is the name of a proposed amendment to the IEEE 802.11 specification for Wireless Local Area Networks (WLANs). IEEE task group TGac was started in November 2008 to consider technical proposals and to draft a text proposal for the amendment. This work is currently still in draft stage, with final approval targeted for December 2013. Once approved, it will become an official 802.11 amendment under the name 802.11ac. Even though the work has not been officially approved yet, we will refer to it as 802.11ac in this document.

The main goal of the new 802.11ac amendment was to significantly increase the throughput within the Basic Service Set (BSS). The official target rates, as defined at the start of the project, are a maximum Multi−Station (Multi−STA) throughput of at least 1 Gbps and a maximum single link throughput of at least 500 Mbps. These higher rates are motivated by the continuing trend to transition devices and applications from fixed links to wireless links and by the emergence of new applications with ever higher throughput requirements.

Existing 802.11 technologies operate in the 2.4 GHz band (802.11b, 802.11g), the 5 GHz band (802.11a), or both (802.11n). 802.11ac operates strictly in the 5 GHz band, but supports backwards compatibility with other 802.11 technologies operating in the same band (most notably 802.11n).

To achieve its goals, 802.11ac relies on a number of improvements in both the MAC and Physical Layer (PHY).

The PHY improvements include:

- Increased bandwidth per channel
- Increased number of spatial streams
- Higher−order modulation − 256 Quadrature Amplitude Modulation (QAM)
- Multi−User Multiple Input Multiple Output (MU−MIMO)

In addition to these new PHY features, 802.11ac also supports a number of advanced digital communication concepts that were first introduced in 802.11n, such as space division

multiplicity, Low−Density Parity  Check (LDPC) coding, shortened guard interval (short GI), Space−Time Block Coding (STBC), and explicit−feedback transmit beamforming (Tx BF).

The Media Access Control (MAC) layer includes many of the improvements that were first introduced with 802.11n. One notable enhancement is the larger maximum size of aggregate MAC Protocol Data Units (MPDUs). Also, the Request to Send/Clear to Send (RTS/CTS) mechanism has been enhanced to allow more efficient implementation of dynamic bandwidth operation.

Table 1 summarizes some of the advanced features used in 802.11n, 802.11ac, or both. We'll discuss them in more detail in the next section.

**Table 1. Features of 802.11n and 802.11ac**

|  | 802.11n | 802.11ac |
|---|---|---|
| MIMO | Supported | Supported |
| Channel bandwidth | 20 MHz and 40 MHz | 20, 40 and 80 MHz (Mandatory) 160 MHz and 80 + 80 MHz optional |
| LDPC | Supported (optional) | Supported (optional) |
| STBC | Supported (optional) | Supported (optional) |
| Implicit feedback Transmit beamforming | Supported (optional) | Not supported |
| Explicit feedback Transmit beamforming | Supported (optional) | Supported (optional) |
| Short GI | Supported (optional) | Supported (optional) |
| MU−MIMO | – | Optional |
| Spatial streams | 1 to 4 spatial streams (2 SS mandatory for AP) | 1 to 8 spatial streams (> 1 SS optional) |
| Modulation | BPSK, QPSK, 16 QAM, 64 QAM | BPSK, QPSK, 16 QAM, 64 QAM, 256 QAM (optional) |
| Unequal Modulation MCS | Supported (optional) | Not supported |
| Frequency | 5 GHz and 2.4 GHz band | 5 GHz band |
| Greenfield preamble | Supported (optional) | Not supported |
| MAC improvements: Reduced inter−frame spacing | Supported | Not supported |
| MAC improvements: block ACK, aggregation (A−MSDU and A−MPDU) | Supported | Supported |
| MAX A−MPDU size | Maximum size 65,535 octets | Maximum size 1,048,575 octets |

While 802.11ac promises big increases in PHY rates, it is important to see this performance improvement in context. Specifically, one has to be aware of the impact that interference will have on the actual throughput of wider−bandwidth systems. This aspect will be covered in some length later in this paper.

# Overview of 802.11ac Features

### Introduction

802.11ac shares many features with 802.11n. Advanced coding (LDPC) is used for increased coding gain. STBC technology can be used for transmitter diversity. A number of other features provide various levels of performance improvement, such as explicit feedback Tx BF, and short GI. These features are essentially identical in 802.11ac and 802.11n. Below we will discuss and evaluate features that truly distinguish 802.11ac from 802.11n.

## Modulation Enhancements

Like most recent wireless specification, 802.11ac uses Orthogonal Frequency–Division Multiplexing (OFDM) to modulate bits for transmission over the wireless medium. While the modulation approach is identical to that used in 802.11n, 802.11ac optionally allows the use of 256 QAM in addition to the mandatory Quadrature Phase Shift Keying (QPSK), Binary PSK (BPSK), 16 QAM and 64 QAM modulations. 256 QAM increases the number of bits per sub–carrier from 6 to 8, resulting in a 33% increase in PHY rate under the right conditions. It should be noted however that 256 QAM can only be used in high signal–to–noise ratio (SNR) scenarios (across the used spectrum and desired streams); i.e. for very favorable channel conditions. The support of 256 QAM will increase the maximum PHY rate that can be supported by the system, but will have no effect in typical scenarios and will not lead to any reach increase for the service. Also, supporting 256 QAM requires transmitter and receiver to be designed such that the inherent SNR (transmit and receive Error Vector Magnitude, or EVM) of the system is able to accommodate the higher constellation.

This will make the RF design of a system that supports 256 QAM more challenging. Unlike 802.11n, 802.11ac does not support the use of unequal modulation (UEQM). This means that all streams in a multi–stream transmission have to be modulated with the same constellation size. UEQM, by contrast, enables the system to modulate weaker streams with lower modulations, which allows for more fine–grained optimization of the data rate to a particular channel environment. This may be important for higher numbers of streams, especially in combination with beamforming.

### Increased Bandwidth

The most notable feature of 802.11ac is the extended bandwidth of the wireless channels. 802.11ac mandates support of 20, 40 and 80 MHz channels (versus 20 and 40 MHz in 802.11n). Optionally, the use of contiguous 160 MHz channels or non–contiguous 80 + 80 MHz channels is also allowed. The doubling of the channel bandwidth (from 40 to 80 MHz) is a very efficient way to increase performance in a cost–efficient way. Alternatively, an 80 MHz system can use a lower number of antennas to provide the same performance as a 40 MHz system. However,

this approach should be weighed against other spectrally efficient techniques that provide performance increase. In addition, in most realistic scenarios the performance is not only a function of the PHY rate, but will also be affected by interference from other networks in close proximity. Different bandwidth levels will be affected differently in an interference scenario. Also, reducing the number of antennas eliminates diversity and reduces the robustness of the transmission. These aspects will be discussed further below.

**Cost of Increased−Bandwidth Solution**

By doubling the bandwidth (from 40 to 80 MHz), each spatial stream can roughly support twice the number of bits per symbol. As such, an 80 MHz single−stream transmission can provide the same performance as a two−stream 40 MHz transmission. To support two streams, both transmitter and receiver should have at least two antennas, while the single stream transmission can be sent (or received) with a single antenna. This means that an 80 MHz systems designed with only a single RF and baseband transmit/receive chain requires less hardware for the same performance than a 40 MHz system designed with two RF and base−band transmit/receive chains. It is true that the 80 MHz RF and baseband design will be more demanding than the 40 MHz design, but it is likely that some cost advantage remains for the wider−bandwidth system. However, single-antenna systems will not provide the reliability that is required for certain Quality of Service (QoS) scenarios. Even 80 MHz systems will need multiple chains to reliably transport video. Any remaining cost advantage is expected to become marginal over time as silicon costs decrease.

**Power Consumption of Increased−Bandwidth Solution**

When enhanced bandwidth is used to deliver the same data rate with fewer RF chains, the power consumption of the device will be lower by virtue of the lower number of RF components. This gives an advantage to the 80 MHz system over a 40 MHz system with two streams from this perspective. However, one has to consider the fact that a single antenna will not suffice for certain services.

**Required Antenna Diversity**

Increasing the bandwidth enhances the performance of a single stream. If the target is to improve the PHY rate or the maximum throughput of a system regardless of QoS considerations, this may be all that is needed. One has to recognize, however, that transmission of high−quality content such as video has more requirements than just increasing the maximum ideal PHY rate. To ensure stable delivery of video, the number of antennas should be higher than the number of spatial streams. Diversity is a critical part of stable data delivery with QoS (see e.g. [1]). Therefore, even 80 MHz systems will have to be built using multiple antennas if they are going to be used in applications that require stable and reliable transmission of data

(such as video). This narrows the cost and power advantage between a (single−stream) 80 MHz bandwidth system and a (two−stream) 40 MHz system.

**Performance in Interference Environment**

The use of the 5 GHz band has significantly increased the amount of bandwidth available for wireless transmission. However, even this band is ultimately a limited resource, and ever−increasing competition for bandwidth share will be a reality for any 802.11 system operating in this band.

802.11ac specifies that 80 MHz channels consist of two adjacent 40 MHz channels, without any overlap between the 80 MHz channels. This results in channel allocations as illustrated in Figure 1 and Figure 2 for the U.S. and Europe, respectively. The number of 80 MHz channels for the U.S. is 5, while in Europe and Japan the number is 4.
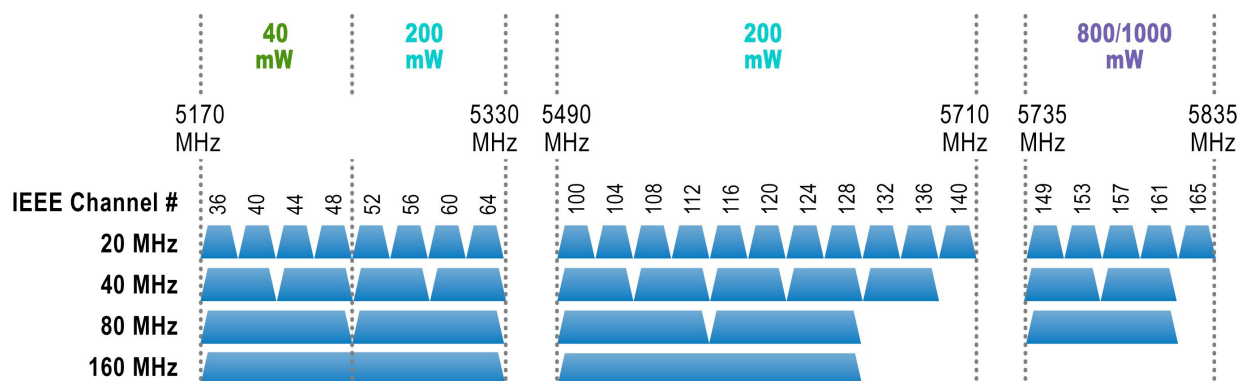


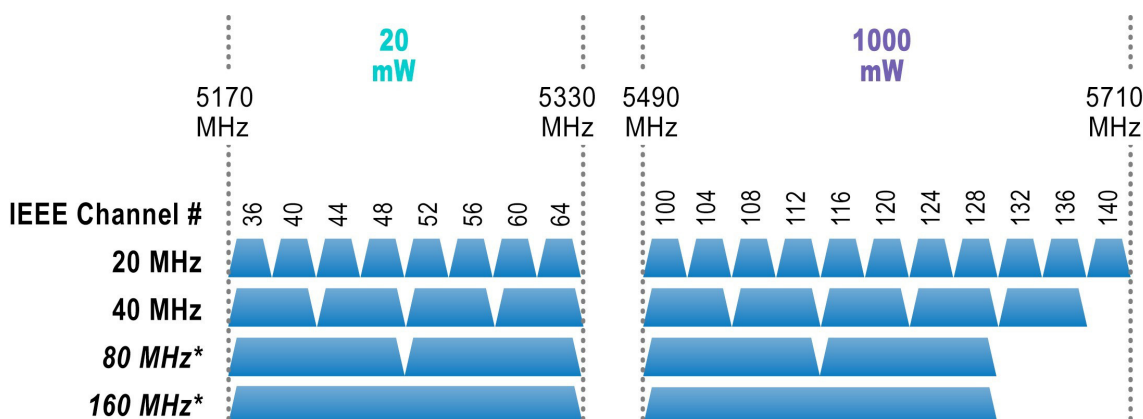**Figure 1. US Channel Allocations for 20/40/80/160 MHz**



**Figure 2. European Channel Allocations for 20/40/80/160 MHz**

For 20 and 40 MHz systems, one of the primary mechanisms for interference mitigation is a judicious channel selection algorithm. This allows neighboring networks to essentially avoid

each other by using different channels when they are within each others range. Channels can be reused by networks that are sufficiently far removed.

With only four or five available 80 MHz channels, it becomes much harder for an 80 MHz system to avoid interference from neighboring networks (which could be either 80 MHz networks or 20/40 MHz networks). Overlapping Base Station Subsystem ((OBSS) problems will be more prevalent for 80 MHz systems. This means that neighboring 80 MHz networks will have to share the same channel with a neighboring BSS, with each getting access only part of the time. In addition, an 80 MHz system is disadvantaged by the fact that all four 20 MHz (or two 40 MHz) channels with which it overlaps have to be clear before the 80 MHz transmission is allowed to start. Even a 20 or 40 MHz transmission in part of the 80 MHz channel will preempt the complete 80 MHz channel from sending. This reduces the probability that an 80 MHz system will gain access to the channel in a dense environment. It should be noted that 802.11 ac provides a mechanism for 80 MHz to fall back to lower bandwidth modes, but only under some conditions (for instance, 40 MHz transmission on the secondary channel, only, is not allowed). Also, this mechanism is optional and need not be implemented to be 802.11ac−compliant.

Figure 3 illustrates a possible difference in interference scenarios between an 80 MHz system and a 40 MHz system. Figure 3(a) shows an 80 MHz system occupying four 20 MHz channels, while one of the 20 MHz channels is also used by a legacy 20 MHz system. In this scenario, the 80 MHz system has no way to avoid the occupied channel and has to share access to the medium with the 20 MHz system. If access is equally shared between the two systems, the capacity of the 80 MHz system is cut in half. Note that the 80 MHz system cannot fall back to 40 MHz transmission in this case, since the overlap happens in the primary 40 MHz channel. In comparison, Figure 3(b) shows how a 40 MHz system can avoid the occupied 20 MHz channel by channel selection. In this scenario, the 40 MHz system has full unshared access to the medium. A single−stream 40 MHz system would have the same capacity as the single−stream 80 MHz system shown in Figure 3(a). If the same 40 MHz system were to support two streams, however, its capacity would be double that of the single−stream 80 MHz system shown in Figure 3(a).
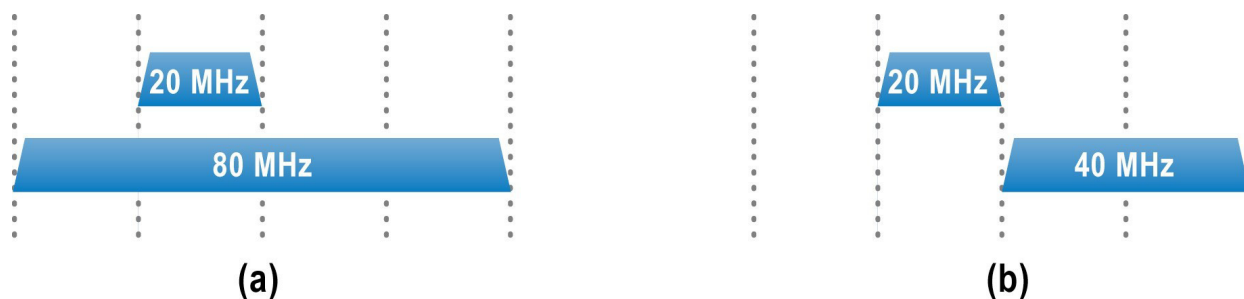


**Figure 3. Two Interference Scenarios**

There are many different interference scenarios, and the above example only aims to illustrate the importance of fully understanding interference dynamics and implications when it comes

to assessing the true capacity of a system. It is also important to realize that one is unlikely to encounter "greenfield" scenarios without legacy or other 802.11ac systems, where no pre−existing interference−related constraints exist. Interference will virtually always be an issue, to one degree or another.

In the appendix, we include simulation results that compare 80 MHz and 40 MHz performance in a specific simulation scenario. This further emphasizes the fact that performance is, to a large extent, dependent on the interference environment. The interference issue is very important to service providers, where it is essential that the wireless infrastructure be future−proofed.

To be competitive in an interference environment, it is important for an 80 MHz system to have a fallback mode to lower bandwidth levels that provide similar performance. This means the equipment has to support a sufficient number of antennas to fully and effectively exploit the benefits of channel diversity.

## Increased Number of Streams

802.11ac allows support for up to 8 spatial streams − up from a maximum of 4 streams in 802.11n. Support for more than one spatial stream is optional, however. It is not clear whether a real−world, single−user MIMO channel can realistically support that many streams. The increased number of streams may be most useful in combination with MU−MIMO.

### MU−MIMO

MU−MIMO was added to 802.11ac to address the multi−STA throughput requirement. In MU−MIMO, the Access Point (AP) − or possibly another STA − transmits independent data streams to several STAs at the same time. Through preprocessing of the data streams at the transmitter (similar to what happens in beamforming), the interference from streams that are not intended for a particular STA is eliminated at the receiver of each STA. Therefore, in theory, each STA receives its data free of interference from the transmissions that are simultaneously directed towards other STAs. In MU−MIMO, the spatial degrees of freedom are used to create independent transmissions to different STAs, while in single−user MIMO, these spatial degrees of freedom are used to increase the throughput from AP to STA.

The complexity of MU−MIMO falls mostly on the AP (or transmitting STA), where the preprocessing happens. The receiving STAs only need the capability to report channel information to the AP so it can calculate the preprocessing matrices. The required channel information from the receiving STA is very similar to what is required for explicit feedback beamforming. As such, the complexity for the STA is no more than the complexity already involved in supporting explicit feedback beamforming as a receiver.

One drawback of MU−MIMO is that the amount of time that the medium is occupied is determined by the slowest link among all AP−STA pairs (or, more generally, the link that requires

the most time to finalize its transmission). No new data can be sent to any of the STAs until all transmissions to STAs in the MU−group have ended. If there is too much difference in either the amount of data or throughput going to various STAs, this may lead to inefficient use of the wireless medium.

At this point, MU−MIMO is a well−studied concept, but practical considerations will likely defer implementation of this feature to later generations of 802.11ac products. Additional work may be needed to guarantee the efficient use of MU−MIMO.

## MAC Improvements

### Increased Aggregated MPDU (A−MPDU) size

The maximum size of an A−MPDU can optionally be increased to a maximum of 1,048,575 octets (compared to a maximum of 65,535 octets in 802.11n).

### RTS/CTS Operation for Wider Bandwidth

Because of the wider bandwidth used in 802.11ac and the limited number of 80 MHz channels, hidden nodes on the secondary channels are an important problem to address. The RTS/CTS mechanism has been updated to better detect whether any of the non−primary channels are occupied by a different transmission. To this end, both RTS and CTS (optionally) support a "dynamic bandwidth" mode. In this mode, CTS may be sent only on the primary channels that are available in case part of the bandwidth is occupied. The STA that sent the RTS can than fall back to a lower bandwidth mode. This helps to mitigate the effect of a hidden node. Note however that the final transmission bandwidth always has to include the primary channel.

### Other

Reduced Inter−Frame Spacing (RIFS) is a deprecated feature of the 802.11n specification whose purpose was to increase MAC efficiency by reducing the gap between successive transmissions. RIFS can be applied between transmissions within the same burst. This mechanism was removed from 802.11ac, except for what is needed to maintain backward compatibility with 802.11n. It was felt that aggregation provided a more efficient way to increase MAC efficiency, and that the complexity of RIFS implementation did not outweigh its gains as a stand−alone mechanism.

### Conclusion

802.11ac has the potential to provide the next generation in high−throughput wireless systems. To fully realize this potential, 802.11ac systems will have to go beyond a minimal implementation that simply exploits the wider bandwidth channels available to this technology. Any new system will be measured against currently available 802.11n systems that already

implement MIMO processing with space–division multiplexing, LDPC, STBC, beamforming, multiple streams and a variety of other PHY, MAC and coexistence enhancements. First–generation 802.11ac systems must be evaluated in light of this comparison. As a minimum, such systems would have to match the feature set that is already provided by current–generation 802.11n. Preferably, any next–generation system would include some truly next–generation features (such as MU–MIMO) in addition to the channel bandwidth increases that are readily available in this new technology. The bandwidth increase of 11ac is currently a concern in situations with limited bandwidth resources. Frequency is a scarce resource that needs to be used as efficiently as possible. Exploiting channel diversity by using a higher number of spatial streams allows more efficient spectrum use than simply doubling the bandwidth of the transmission. Channel and antenna diversity, therefore, remain important requirements, even for systems that are capable of wider bandwidth. It is believed that a 4x4 system with a maximum number of spatial streams and MU–MIMO will be required, at a minimum, in order for 11ac to fully realize its potential. Such a system would provide higher bandwidth in sparsely populated networks, while providing QoS, good performance and coexistence in denser network environments.

## Appendix: 80 MHz vs. 40 MHz Performance

There are many possible interference scenarios that can be considered when evaluating the relative performance of 80 and 40 MHz systems (or other bandwidth scenarios). We'll consider one method that has been applied in a number of IEEE 802.11 contributions to the TGac Task Group [4]–[7].

The method that follows evaluates the effective bandwidth of a system that operates in an environment characterized by a certain number of busy 20 MHz channels. The effective bandwidth is a measure that is ultimately proportional to the throughput. Each busy channel is characterized by a load between 0 and 100%, which represents the likelihood that the 20 MHz channel is occupied at any given time.

For a given number Nb of busy channels out of a total of Nt channels, we consider all possible ways that  the Nb busy channels can be allocated within the available Nt channels. For each allocation of busy channels, we find the best channel allocation for an 80 MHz or 40 MHz system. The best channel allocation for the 80 MHz or 40 MHz system is the one that provides the lowest probability that the 80 or 40 MHz system will have to defer its transmission because of interference from the busy channels.

The simulation approach is illustrated in Figure 4.

**Figure 4. Illustration of the Optimal Location of an 80 MHz System in an Environment with Five Busy Channels**

Figure 4 shows an environment with five busy 20 MHz channels, which are shaded in the diagram to identify them. Within such an environment, an 80 MHz channel can be assigned to four different locations. Only one of these locations (shown as un–shaded) is free of interference for this specific choice of five busy channels. In this location, the 80 MHz system has the best probability (in this case, 100%) of gaining access to the medium. Figure 4 shows just one possible allocation of five busy channels. In the simulation, the process is repeated for many different allocations, and the location of the 80 MHz system that gives it the optimal probability of gaining channel access is found for each of them. The probability is averaged over all possible environments with the same number of busy channels. The effective bandwidth is the "native" bandwidth of the system (e.g. 80 MHz) multiplied by the probability that the system has access to the medium.

Figure 5 shows an example that illustrates the performance of an 80 MHz and a 40 MHz system, as calculated using this methodology. The horizontal axis shows the environment, as characterized by the number of busy (20 MHz) channels. The vertical axis shows the effective bandwidth (which is a measure that is proportional to the throughput of the system).
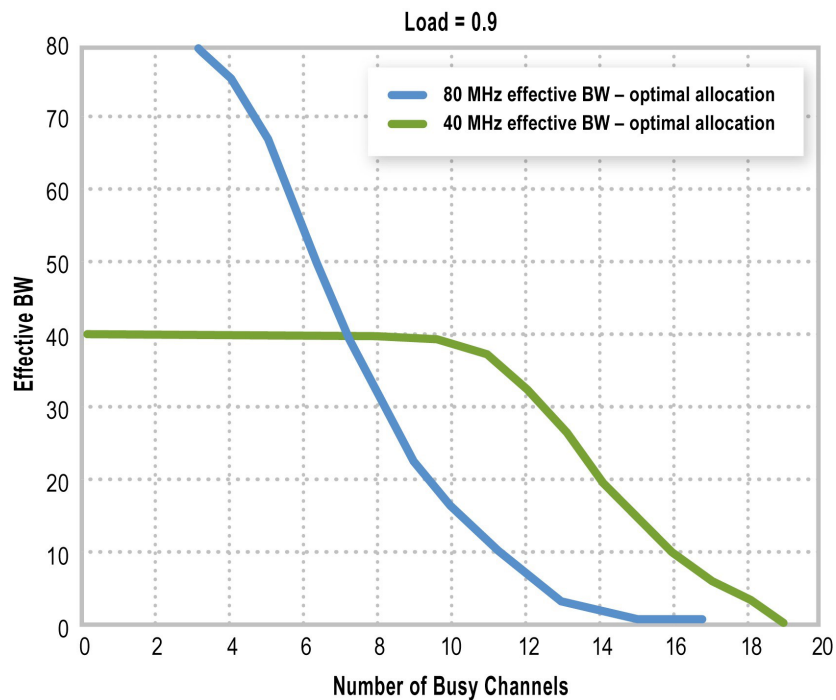


**Figure 5. Performance Comparison of 80 MHz and 40 MHz System**

In an environment with no or low interference, the 80 MHz system will achieve about double the performance of a 40 MHz systems (assuming the same number of streams for both). Not surprisingly, the 80 MHz systems begins to exhibit the impact of interference sooner (i.e. for a lower number of busy channels) than a 40 MHz system would. The performance of the 80 MHz system degrades rapidly as the number of busy 20 MHz channels increases. For the results shown in Figure 5, the throughput of the 80 MHz systems drops below the throughput of the 40 MHz channel when the number of busy channels exceeds seven. In a high−interference environment, the 40 MHz system performs significantly better than the 80 MHz system.

## References

[1]  Designing reliable Wi−Fi® for HD delivery throughout the home, Quantenna white paper
[2]  Harmonized European Standard ETSI EN 301 893, Broadband Radio Access Networks; 5 GHz high performance RLAN
[3]  MIC Equipment Ordinance for Regulating Radio Equipment
[4]  P. Christin et al, Non contiguous 40 + 40 MHz mode for Europe, Japan and global, IEEE 802.11−10/127
[5]  L. Cariou, J. Benko, Gains provided by multichannel transmissions, IEEE 802.11−10/0130
[6]  L. Cariou, P. Christin, 80 MHz and 160 MHz channel access modes, IEEE 802.44−10/0385
[7]  L. Cariou et al, Evaluation of saturation of the 5 GHz band, IEEE 802.11−10/0846

◊