WHITE PAPER

Key Considerations for Wi-Fi 6E Testing

Introduction

Wi-Fi 6/6E refers to the IEEE 802.11 standard based Wireless Local Area Network (WLAN) technology that introduces several new enhancements and features compared to previous generations of Wi-Fi solutions. Wi-Fi 6E is based on the IEEE 802.11ax standard. As an extension of Wi-Fi 6, it enables the operation of features in the unlicensed 6 GHz band, in addition to the currently supported 2.4 GHz and 5 GHz bands . Wi-Fi 6/6E provides networks the capacity, efficiency, and performance required to support next-generation applications. In addition to 6 GHz operation, Wi-Fi 6/6E will soon encompass the "Release 2" set of features identified by the Wi-Fi Alliance and referred to as Wi-Fi 6/6E R2. In this paper we will include the R2 set of features in our discussion.

The IEEE 802.11ax standard on which Wi-Fi 6/6E is based includes the following enhancements and improvements compared to previous generations of Wi-Fi:

- Enhanced operation in 2.4GHz, 5 GHz, and new 6 GHz bands
- Increased average throughput per station by at least 4x in dense deployment scenarios
- Environments including indoor AND outdoor operation
- Scenarios include dense public area, corporate office, outdoor hotspot, dense residential apartments, and stadiums
- Improved power efficiency of the stations

The major improvements introduced in this new WLAN standard are summarized but not limited to those in Figure 1, including:

- Higher order modulation with 1024 QAM support for ~20% higher rates
- Downlink and uplink Orthogonal Frequency-Division Multiple Access (OFDMA),
- Downlink and uplink Multiuser, Multi-Input, Multiple-Output (MU-MIMO)
- Longer OFDM symbol with symbol duration increased by 4x
- Target Wake Time (TWT)
- Tri-band support
- Basic Service Set (BSS) color for spatial re-use
- Support for up to 8 spatial streams

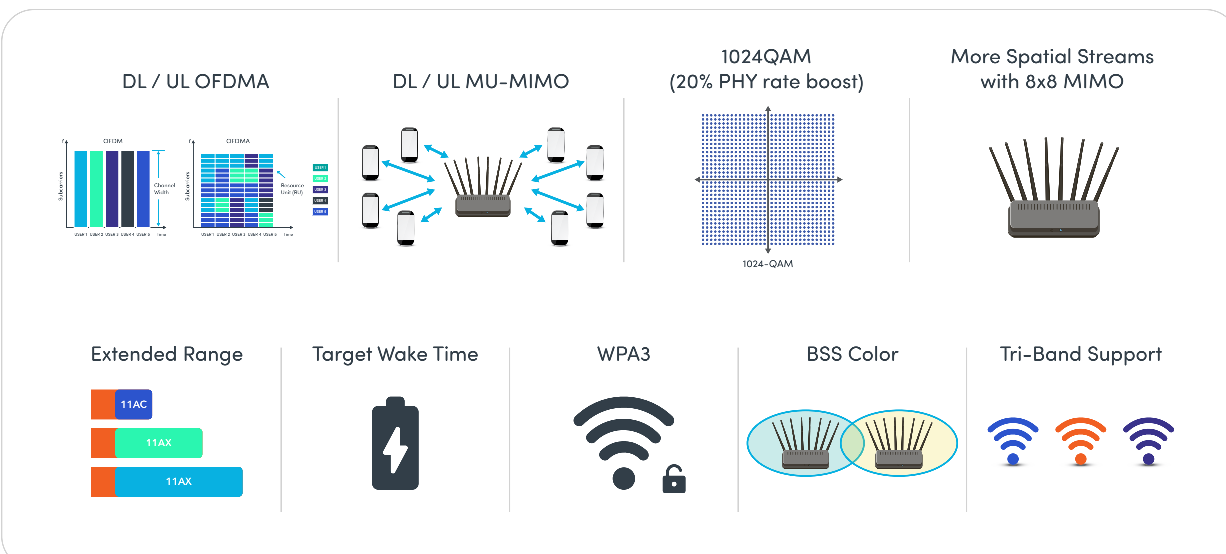

Figure 1

A note on terminology. The fact that Wi-Fi 6/6E has overlapped in time with the availability of the 6 GHz spectrum band can cause some confusion in understanding the terminology. The "6" in Wi-Fi 6/6E corresponds to the technology generation. Wi-Fi 5, for example, refers to technology that is part of the IEEE 802.11ac specification, whereas Wi-Fi 6/6E refers to 802.11ax. It is the "E" in Wi-Fi 6E that refers to the 6 GHz band. The "E" in this case means "extended", that is, it refers to Wi-Fi 6/6E technology designed to operate in an extended range of frequencies that include the 6 GHz bands. To add to the confusion, the Wi-Fi Alliance has dictated that, from the point of view of Wi-Fi certified products, only Wi-Fi 6/6E products will be permitted to operate in the 6 GHz bands. But, again, those two "6s" are not the same. Wi-Fi 6/6E refers to the technology generation, and 6 GHz refers to the frequency band.

Key Elements of Wi-Fi 6/6E (IEEE 802.11AX) and Testing Implications

The key new features that are part of Wi-Fi 6/6E and hence, the main focus on this white paper are:

- Multi-user operation using a variety of techniques.
  - OFDMA for spectral efficiency and improved latency.
  - MU-MIMO for increased capacity using spatial diversity. Note that while MU-MIMO was part of Wi-Fi 5, it has been enhanced in Wi-Fi 6/6E, including the addition of UL MU-MIMO in R2.
- Enhanced power-save operation using Target Wake Time (TWT).
- Improved high-density deployment handling using BSS Coloring.

As will be discussed in more detail below, Wi-Fi 6/6E testing is fundamentally different from previous Wi-Fi testing for a few reasons:

1. Possibly the most discussed new feature in Wi-Fi 6/6E is OFDMA, and OFDMA is, by definition, a multi-station technology. There is no meaning to OFDMA without at least two (and usually more) stations communicating, simultaneously, with the AP. Whereas in previous Wi-Fi generations testing could be configured primarily around a single AP / single STA model, that is no longer possible with Wi-Fi 6/6E. Testbeds must be configured to support multiple station operation and testing.

2. Wi-Fi 6/6E changes the balance between Access Points (APs) and Stations (STAs) that existed in previous Wi-Fi generations. Although there have always been APs and STAs, prior to Wi-Fi 6/6E they were largely similar devices which used the Wi-Fi channel in similar ways. Under Wi-Fi 6/6E, however, the AP has become a controlling entity, so tests need to be able to probe this new functionality.

New Wi-Fi 6/6E test tools and methodologies are needed to meet the industry need for comprehensive testing strategies covering the expanded range of test cases without significantly increasing investment and test time.

Wi-Fi 6/6E breaks the Wi-Fi symmetry between APs and STAs

Before 802.11ax, AP and STAs were treated equally in the Physical (PHY) layer; therefore, Wi-Fi used the same type of PHY Protocol Data Unit (PPDU) format for APs and STAs. Since 802.11ax has adopted the concept of multi-user Downlink (DL) (AP to non-AP STA) and Uplink (UL) (non-AP STA to AP) communications, the standard now introduces four types of PPDU formats for various transmissions. (Note: the designation "HE" in the below, and throughout the remainder of this whitepaper, refers to functionality specific to the 802.11ax standard. This standard, or amendment to the 802.11 standard, is known as "Enhancements for High Efficiency WLAN", and "HE", therefore, refers to "High Efficiency.")

- **HE-SU PPDU:** [DL/UL] Normal Single User PPDU between an AP and an STA
- **HE-MU PPDU:** [DL] Multi-user PPDU from AP to multiple STAs to enable simultaneous transmission via OFDMA and/or MU-MIMO. The new HE-SIG-B introduced indicates the Resource Unit (RU) and spatial stream info for MU operation.
- **HE-TB PPDU:** [UL] Trigger-based PPDU for UL MU (OFDMA, MU-MIMO) operation from STAs. Compared with HE MU PPDU, this format does not have a HE-SIG-B field because the AP indicates the RU allocation information in the Trigger frames. (Sometimes referred to as HE-TRIG.)
- **HE-ER SU PPDU:** [DL/DU] Enhanced Range SU PPDU. Not a multi-user format, this format has training fields boosted by 3dB to extend range. The packet structure is same as HE SU PPDU, except for the HE-SIG-A field which is repeated. Compared with the HE SU PPDU, this format has 8us HE-STF vs. only 4us in HE SU PPDUs.

This broken symmetry is an indication of how Wi-Fi 6/6E has changed the Wi-Fi testing environment. In previous generations of Wi-Fi, it may have been sufficient to look at end-user metrics like throughput or packet latency in order to quantify the performance of a Wi-Fi setup. In Wi-Fi 6/6E, on the other hand, there are more complicated effects to examine. For example, since the AP has taken on a significant scheduler responsibility in OFDMA operation, new tests that probe that scheduler functionality are needed. In addition, whereas it used to be possible to look at all of the packets on the air using a sniffer in promiscuous or monitor mode, that is no longer possible in Wi-Fi 6/6E. As we will discuss in more detail below, OFDMA and MU-MIMO packets are sent using specific IDs (association IDs, or AIDs) that allow receivers to identify traffic even when multiple packets are being sent simultaneously. Without using these IDs, it is no longer possible to sniff the packets. For these reasons, among others, testing a Wi-Fi network has become much more complicated with Wi-Fi 6/6E than it was with previous Wi-Fi generations.

IEEE 802.11AX OFDMA

Orthogonal Frequency-Division Multiple Access (OFDMA) introduced in Wi-Fi 6/6E is a large step beyond Orthogonal Frequency Division Multiplexing (OFDM), which is used by some of the previous generations of Wi-Fi. OFDMA is a feature that has been used in wireless communication systems such as LTE for many years to improve spectral efficiency.

With OFDMA in Wi-Fi 6/6E, multiple users can simultaneously share the bandwidth of a full radio frequency (RF) channel for sending data more efficiently. Imagine the following simple scenario. There are four users connected to a single access point using an 80 MHz channel. Each user has traffic to send to the AP, but not very much; the data could easily be sent using only a 20 MHz channel. Prior to Wi-Fi 6/6E, each user would gain access to the channel (the whole channel) and send its data, thus wasting the unused capacity. In Wi-Fi 6/6E, however, the AP can arrange it so that each of the 4 users sends its information to the AP at the same time, using only one-fourth of the channel capacity. In this way, the channel is 100% used, but this is accomplished by scheduling four users to send data at the same time. This is the essence of Wi-Fi 6/6E OFDMA.

The difference between OFDM and OFDMA is illustrated in Figure 2. In OFDM mode, one of 4 users always occupies the entire frequency channel bandwidth at a time. In OFDMA, all 4 users can share the same channel at the same time, whenever it makes sense (is more efficient) for the network to operate that way.

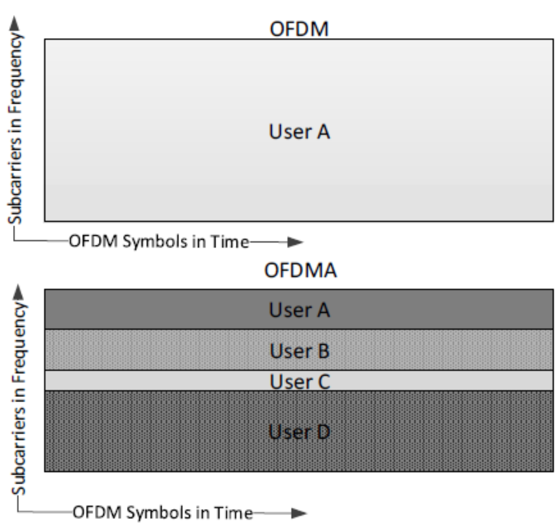

Figure 2: OFDM vs OFDMA in 802.11ax
(from IEEE 802.11 Draft Standard)

Each user is told when to transmit, and what part of the channel to use. This combination of the frequency to use and the time during which to use it for each user is an assignment called a Resource Unit (RU). (For those familiar with LTE, this is equivalent to the concept of a Resource Block (RB).) As discussed above, this means that a wide range of test coverage is now required to validate the OFDMA operation. APs must now decide how many users to schedule simultaneously, what RU assignments to give them, as well as many other details. And those assignments are not fixed; they can change with each transmission opportunity.

One fundamental change to testing procedures should be evident from the above description. Whereas in previous generations of Wi-Fi many tests were point-to-point (one AP and one STA), test cases which can probe this OFDMA operation require a testbed with many independent client station (STA) radio interfaces in order to initiate OFDMA operation by the AP. In addition, in order to probe different types of scheduler behavior, the testbed also needs to be able to emulate various realistic use cases with different types of RUs assignments managed by an AP for possible OFDMA operation.

Downlink OFDMA Transmission

In downlink operation from an AP to many clients, OFDMA allows a single transmission to be split into different frequencies within a single RF channel so that multiple concurrent users can get access to the information in each frequency allocation.

HE DL MU operation allows an AP to transmit simultaneously to one or more STAs. For a downlink OFDMA transmission opportunity from an AP to multiple STAs, the AP schedules RUs in different sections of the RF channel to each STA, and then sends packets to each STA in the proper RU over the air interface. Since the AP is the controller in the case of OFDMA, there is no coordination with the STAs required; the STAs simply need to follow whatever directions are sent by the AP. The AP will use the HE MU PPDU packet type, which will contain all the encoded information required for the STAs to recognize data addressed to them, and to find and decode that data.

In the previous (OFDM) case, the AP would only be able to communicate with a single STA at a time. Packets would be sent to that STA, and that STA would then be able to respond with a block ACK. This is shown in Figure 3 below. Especially when the STAs do not have enough data to "fill the pipe", this could be a very inefficient way to send traffic.

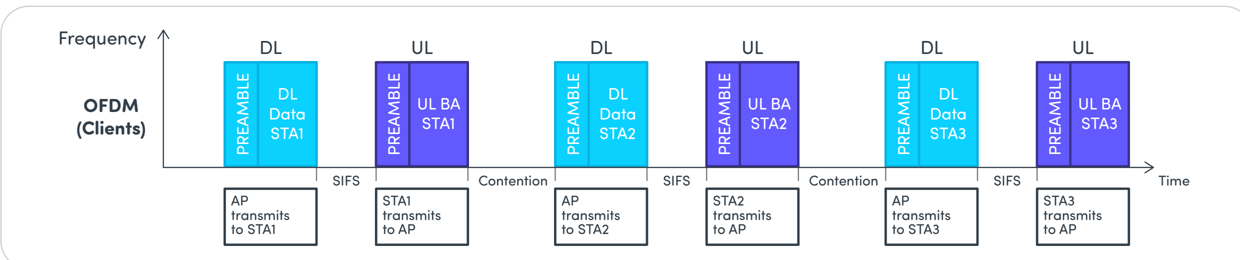

Figure 3

Using OFDMA, however, the AP can send data to more than one STA at a time. Since the packet being sent contains all of the information required, the radio interface of each of the STAs allocated for the downlink OFDMA operation will tune to its own sub-channel to independently receive and decode its own data from the assigned RU. The preamble field of each MU frame specifies the frame duration and subcarrier assignment for each RU for OFDMA. When the STAs successfully receive the OFDMA data, the AP sends a multi-user (MU) Block Acknowledgment Request (BAR) frame called MU-BAR to the STAs so that they can use uplink OFDMA to simultaneously send back Block Acknowledgments (BA) to the AP.

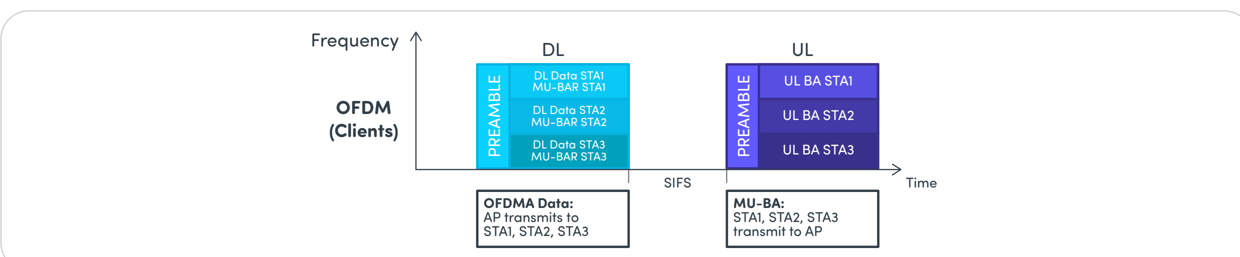

Figure 4

The DL MU-BAR (multi-user block ACK request) is usually a Trigger frame type sent as an independent transmission, but it can also be a part of the MU PPDU to each STA for DL OFDMA as shown in Figure 4. More popularly, the new MU-BAR frame introduced in 802.11ax is a multicast frame for all STAs as shown in Figure 5.

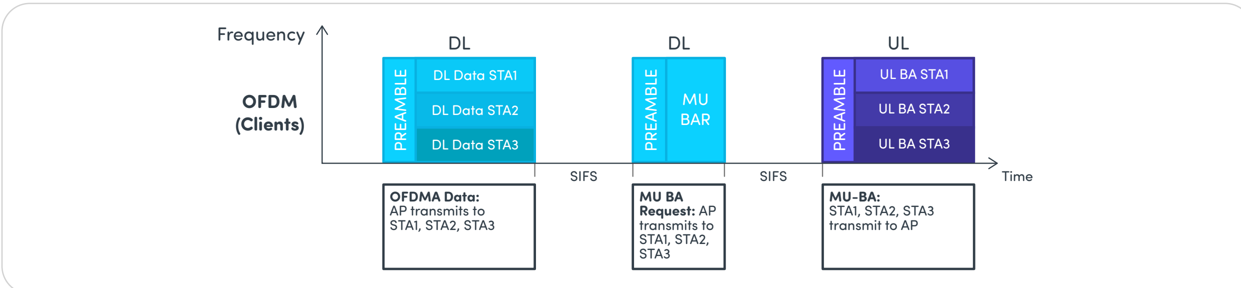

Figure 5

The STAs may also use sequential OFDM transmissions to send BAs back to the AP, in case uplink OFDMA is not available. However, this may not be an efficient ACK mechanism since multiple BAs from the STAs consume more airtime and bandwidth than would simultaneous block ACKs. This is illustrated in Figure 6.

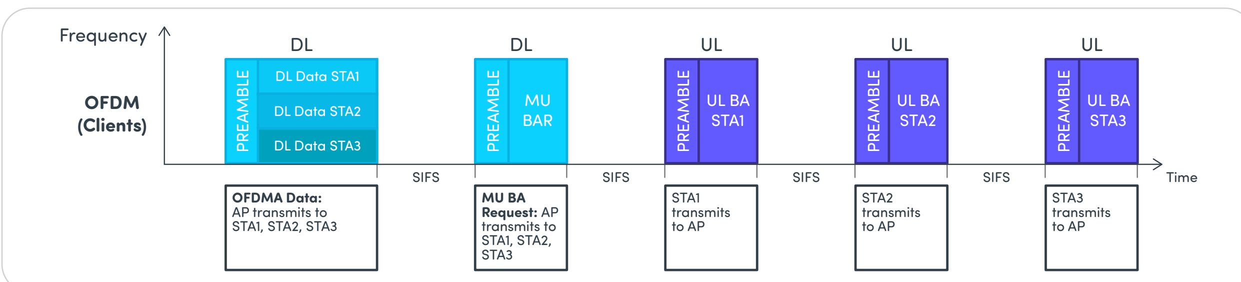

Figure 6

Uplink OFDMA Transmission

UL MU operation allows an AP to solicit simultaneous immediate response frames from one or more HE STAs. UL OFDMA transmission is more challenging than downlink since the AP now has to orchestrate multiple STAs to operate together. For uplink OFDMA operation, the AP uses information from the STAs, including its data buffer status which it obtains via a query or from the STA's unprompted reporting. The AP is the center of the scheduling hub that sets a synchronized transmission timing for all STAs being allocated an uplink OFDMA transmission opportunity along with assigning RUs for those transmissions. IEEE 802.11ax introduces a new type of control frame called the Trigger frame that specifies the common parameters of the upcoming UL MU transmission (duration, guard interval (GI), etc.), allocates RUs to the STAs, and defines transmission parameters such as MCS type, coding, etc.

The mechanism for uplink OFDMA is more complicated than for downlink OFDMA due to the fact that multiple individual STA transmitters, in different physical locations, must be synchronized to complete the operation. But Wi-Fi does not require tight time alignment of transceivers in a network. So, to achieve synchronization, STAs are required to start MU transmission immediately after one predefined time unit (known as a SIFS time) after the Trigger frame is received. When the AP receives the data from all of the scheduled STAs, it will send back a BA using a downlink OFDMA transmission as shown in Figure 7.

The STAs assigned for UL OFDMA must meet the requirements for the transmission timing window, transmission power, and allowed frequency offset error. For UL OFDMA transmission, the AP receives signals from multiple STAs at almost the same transmit power level. 802.11ax defines a power pre-correction mechanism in order to enable this. The AP indicates, in the Trigger frame, its current transmit power and the desired target signal strength that the AP is expected to receive from an STA in the following UL transmission. Therefore, knowing the AP's transmit power and the signal strength of the received Trigger frame, an STA can estimate the path loss to the AP, and it can calculate an appropriate transmit power to use. This implies that the STAs that are very close to an AP and STAs that are very far from the AP may not be able to perform UL OFDMA together due to the large difference in RSSI that the AP senses.

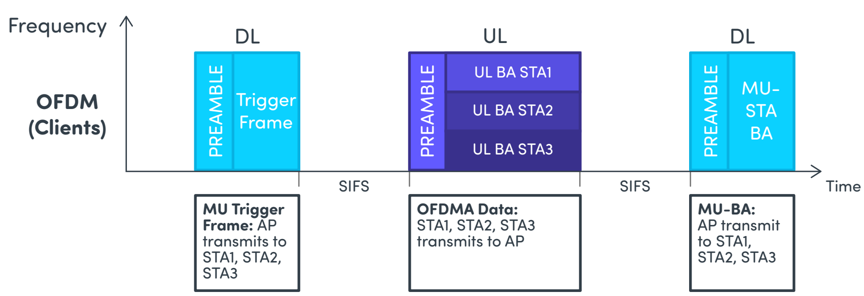

Figure 7

Maximum number of simultaneous users supported by Wi-Fi 6/6E OFDMA

As described above, OFDMA handles simultaneous users on a single RF channel by breaking that physical channel into separate frequency allocations. The number of simultaneous allocations is based on the number of resource units (RUs) that are defined for each channel bandwidth. The IEEE 802.11ax standard specifies a few possible RUs per PPDU for different channel bandwidths as shown in Table 1; each of the RU types can have any one of 26, 52, 106, 242, 484, 996, or 2x 996 subcarriers. (These subcarriers are often referred to as “tones”; hence, the specification will refer to, for example, a “26 tone” RU.) These RU types are represented as RU26, RU52, RU106, RU242, RU484, andRU996 in this white paper. The IEEE 802.11ax standard does not allow multiple RU allocations for a STA. Table 1 shows the maximum number of RU allocation scenarios in terms of the possible RU numbers that can be scheduled for each specific channel bandwidth of 20MHz, 40MHz, 80MHz, and 160MHz. It only shows a single type of RU used in each allocation, so it does not cover other possible mixed RU allocations cases that are also allowed in 802.11ax. It is also clear from Table 1 that RU sizes can be larger when the channel bandwidth in use is larger. The possible RU allocation scenarios with mixtures of different RU types under the maximum allowed RUs per PPDU conditions are numerous and complicated.

Table 1.

RU type	Channel Bandwidth (CBW20) 20MHz	Channel Bandwidth (CBW40) 40MHz	Channel Bandwidth (CBW80) 80MHz	Channel Bandwidth CBW80+80 and CBW160 80+80MHz and 160MHz
RU26: 26-subcarrier RU	9	18	37	74
RU52: 52-subcarrier RU	4^{+1}	8^{+1}	16^{+5}	32^{+10}
RU106: 106-subcarrier RU	2^{+1}	4^{+2}	8^{+5}	16^{+10}
RU242: 242-subcarrier RU	1	2	4^{+1}	8^{+2}
RU484: 484-subcarrier RU	N/A	1	2^{+1}	4^{+2}
RU996: 996-subcarrier RU	N/A	N/A	1	2
2x996 subcarrier RU	N/A	N/A	N/A	1

Notes:
1: x+n means + n 26-subcarrier RUs
2: The darker colored rows show the subcarrier RUs that can be used with MU-MIMO simultaneously

Testing OFDMA

What to test

Testing OFDMA from either the AP or STA side can be a challenging task. OFDMA operation is a complex process, and OFDMA scheduling is mainly handled by the AP and is based on many real-time parameters such as the number of STAs in the network, RF conditions, current serviced traffic loading, traffic characteristics (packet size, type, QoS, etc.). The exact use of OFDMA is not specified, so each AP will have chipset vendor-specific implementations in terms of OFDMA scheduling. An AP is a black box in the way it handles OFDMA scheduling and implements all of the OFDMA rules imposed by the 802.11ax standard. As the AP is in control, there is much more to learn about the OFDMA operation on the AP side than with the STA. A primary focus of OFDMA testing, then, is on understanding how APs have implemented OFDMA.

Since the goal of ODFMA is to enable efficient use of the Wi-Fi channel, the ideal test is to create an environment in which a pre-Wi-Fi 6/6E implementation will result in inefficient channel usage, and then to compare that case with the Wi-Fi 6/6E (OFDMA) case. Since OFDMA is a multi-user technology, testing should involve:

- Multi-user, aggregate, and individual throughput in the OFDMA and non-OFDMA cases.
- Individual round-trip and one-way latencies on OFDMA compared to non-OFDMA cases.

Testing Challenges

As a multi-STA test case, testing OFDMA creates several challenges. Among these are:

- Although multiple STAs are required for testing, traditional "virtual STAs" cannot be used since using a single radio interface to emulate multiple client devices which need to have different RF parameters (e.g. RUs) is not feasible.
- The testing space is large, since it is useful to test different combinations of number of users, traffic types, data rates, packet sizes, QoS settings, etc. This means that adequate test coverage can take a long time.
- Though the AP is the controlling entity, it is usually a black box for OFDMA scheduling.
- Traditional sniffers can no longer view all OFDMA packets using monitor mode. Sniffers for OFDMA operate on a per-client basis.

How to test

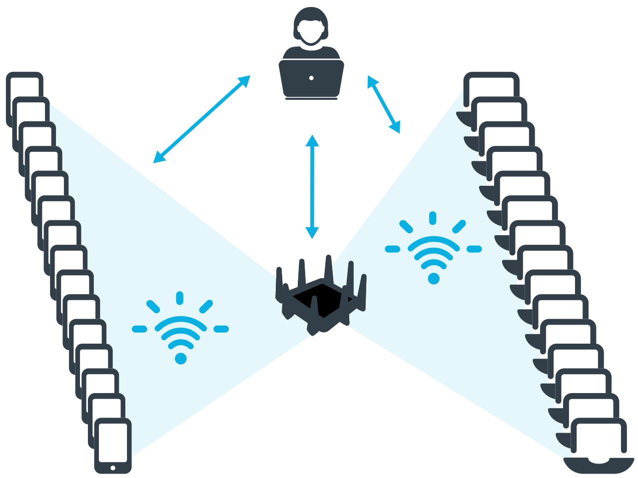

Figure 8

It is certainly possible to put many real STAs, such as Wi-Fi 6/6E enabled mobile phones or laptop PCs, in a lab as an OFDMA testing environment as shown in Figure 8. As with all open-air OTA testing, however, this has the drawback of being completely uncontrolled and not repeatable. In addition, as with all multi-STA testing, the testing complexity in such a model is very high. Somehow control of a large number of off-the-shelf products needs to be provided and, ideally, in a software-controlled way for test automation.

Ideally an OFDMA test will involve testing in a controlled RF environment with multiple STAs which are also controllable and automatable. In addition, the testbed should be able to send traffic of different types (different packet sizes, different QoS settings, different applications) to each STA. And, finally, the testbed should be able to report metrics which will give insight into not only the performance with, and without, OFDMA, but also into the OFDMA operation itself. We will discuss these metrics in more detail below.

What to measure

There are three main classifications of OFDMA testing that can be done.

1. **Sign of life tests**
   Under basic sign-of-life tests, we should simply be able to confirm that OFDMA is happening in the system. For example, we should see evidence of Trigger frames, per-AID Ru assignments, and HE-MU and HE-TRIG PPDUs.

2. **End-user performance tests**
   End user performance tests for OFDMA include comparisons of metrics like throughput and latency in OFDMA compared to non-OFDMA tests. The comparisons require the ability to create scenarios in which OFDMA would be expected to have a positive impact on one, or both, of the above metrics. For example, traffic scenarios in which a large number of devices need to access the channel, but have limited data to send, or tests in which some of the devices have high priority data and other devices do not.

3. **Underlying OFDMA behavior tests**
   In addition to looking at how OFDMA operation affects the end-user experience, some testers may want to look "under the hood" at the OFDMA operation. In this case we can examine information such as the following.

Examples of OFDMA behavior tests

Number of STAs simultaneously allocated OFDMA Resource Units in each Trigger frame, as shown in Figure 9.

Packets to/from each AID separated into uplink/downlink OFDMA (HE_TRIG/HE_MU), single-user (HE_SU), and non-Wi-Fi-6 (Not HE). Shown in Figure 10.

RU allocations per AID for each Trigger frame in the first 1000 packets of a capture. AIDs are differentiated by color. RUs are shown by subcarriers spanned. Shown in Figure 11.

3D representation of RU allocations in all Trigger frames. Each block represents an available RU, and the height indicates the frequency of allocation. For example, the long green box is the 996 tone RU. The tallest boxes are the four 242 tone RUs, showing that those are the most commonly allocated RUs in this test. Shown in Figure 12.

The percentage of subcarriers allocated to users in each Trigger frame. This plot shows that each Trigger frame is allocating about 95% of the available subcarriers (tones) to the STAs. The remaining subcarriers are not used for data traffic. A concerning result would be if this number were much lower, indicating an inefficient allocation of subcarriers. Shown in Figure 13.

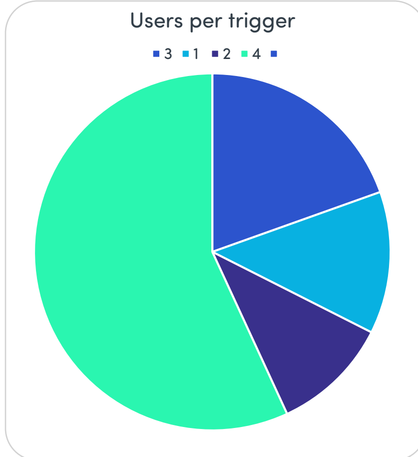

Figure 9

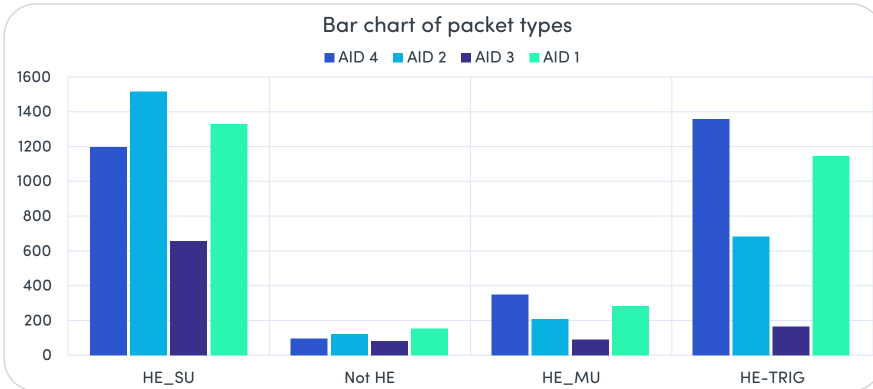

Figure 10

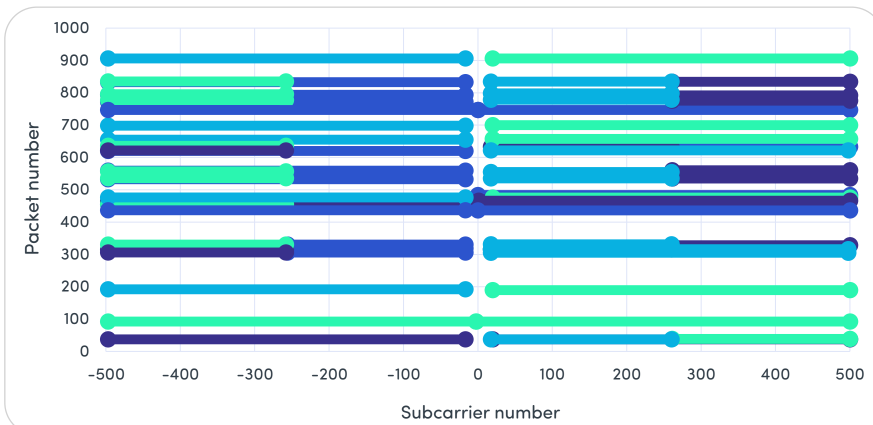

Figure 11

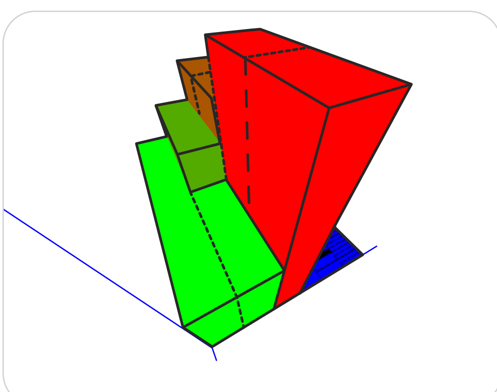

Figure 12

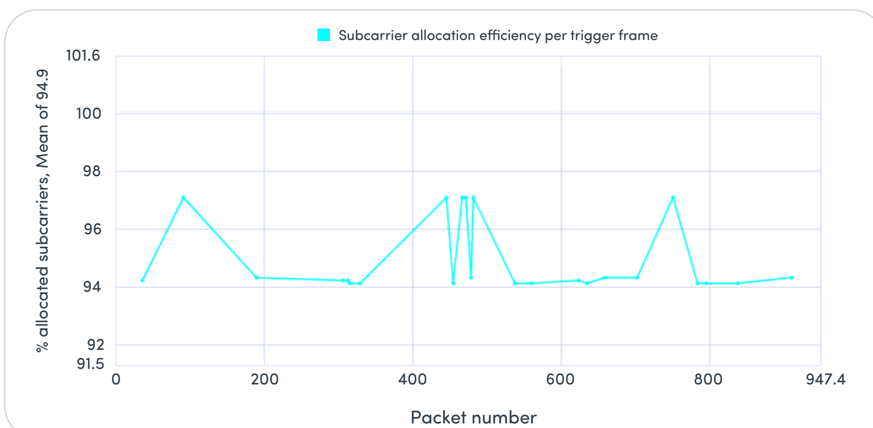

Figure 13

MU-MIMO

MU-MIMO as a spatial domain multiple access technology makes full use of the advantages of multiple antennas, and further improves both transmission rate and spectrum utilization. It allocates different spatial streams for different users, and the receiver must separate the spatial streams of different users by using signal processing technology. The AP collects the channel state information (CSI) from STAs through an HE sounding protocol. STAs are further divided into several groups, and those in the same group can simultaneously transmit data, and the signals of different STAs can be distinguished via spatial streams. Wi-Fi 5 (IEEE 802.11ac) had already adopted DL MU-MIMO, and Wi-Fi 6/6E further improves DL MU-MIMO and introduces UL MU-MIMO, thus enabling symmetrical high throughput. Wi-Fi 6/6E allows up to eight STAs to simultaneously transmit or receive through DL/UL MU-MIMO assuming a Wi-Fi 6/6E AP can support up to 8x8 MIMO.

Both MU-OFDMA and MU-MIMO are for concurrent multi-user service, and an AP can schedule SU-MIMO, MU-MIMO, and OFDMA based on the channel conditions, QoS, device types, service types, and more. There are two key differences between the transmission types:

- OFDMA allows multi-user access by subdividing a channel into different RUs. OFDMA does not change the overall capacity of an AP (that capacity is the capacity of its operating channel) but it makes more efficient use of that capacity.
- MU-MIMO allows for multi-user access by using different spatial streams. MU-MIMO increases the overall capacity of an AP, since the channel capacity is replicated for each spatially separate group of users. See Figure 14.

In other words, OFDMA splits a channel and MU-MIMO creates duplicate channels. Typically, MU-MIMO is used for STAs at short and medium range, while OFDMA can work under all RF conditions, subject to the requirement for time, frequency, and power synchronization. MU-MIMO is more suitable for high-bandwidth data traffic, while OFDMA is effective for low-bandwidth, multi-client service scenarios.

The PPDU frame structure to support MU-MIMO is comparable to that used for OFDMA. For DL MU-MIMO, the AP uses the HE MU PPDU format to support both OFDMA and MU-MIMO. Spatial stream and RU allocation information is embedded in the HE-SIG-B of the HE preamble, and the AP simultaneously transmits frames to multiple STAs through allocated spatial streams. For UL MU MIMO, the AP uses Trigger frames (TF) to trigger multiple STAs and to allocate RUs and spatial streams. After that, the STAs triggered by AP use the HE TB PPDU format to simultaneously transmit frames to the AP through allocated spatial streams.

While so far not a common capability, the specification does allow MU-MIMO and OFDMA to work simultaneously, i.e., multiple STAs can send frames through MU-MIMO in the same RU, to further increase transmission efficiency.

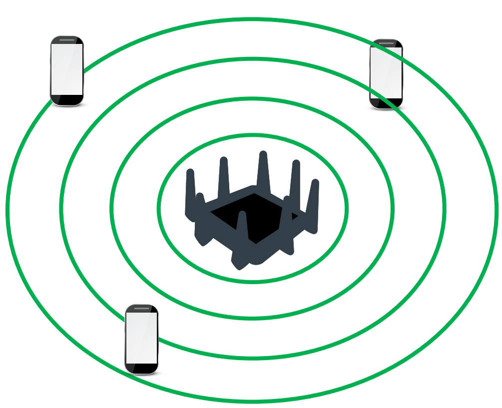

With single-user MIMO, the PHY (RF resource) is shared (time-sliced) among users.

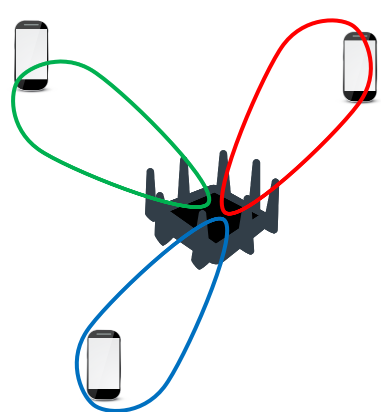

But with multi-user MIMO, a separate PHY can (in theory be created for each (spatially separated) user.

Figure 14

Testing MU-MIMO

What to test

Although both MU-MIMO and OFDMA are multi-user techniques (allowing an AP to communicate with more than one STA simultaneously) they are different in one critical way. OFDMA efficiently uses a single channel of communication; MU-MIMO replicates that channel more than once, so it actually increases the overall capacity of the network.

Therefore, testing MU-MIMO can be somewhat more straightforward than testing OFDMA. With MU-MIMO we should expect to see the increased network capacity. That is, in the simplest case, if the maximum throughput achievable between an AP and a STA is T, then the aggregate throughput achievable between the AP and N STAs could be as high as N*T in the case where MU-MIMO is able to communicate with each of the STAs using its own RF channel.

Testing challenges

As with OFDMA, MU-MIMO is a multi-user technology, so it requires a test environment with multiple STAs. An additional complication is that in the MU-MIMO case, the STAs must be presented to the AP so that they are spatially diverse, enabling MU-MIMO operation. (If the STAs are too close together, the AP will not be able to use beamforming to communicate with them separately, so MU-MIMO will not be possible.)

The main complication for MU-MIMO testing, then, is to be able to create a shielded, controlled, replicable environment in which multiple STAs can be oriented in different spatial orientations relative to the AP.

How to test

The same comments as for OFDMA apply to the MU-MIMO case. Testing done in open-air, with off-the-shelf devices will not be repeatable, or easy to control/automate.

Ideally a MU-MIMO test will involve testing in a controlled RF environment with multiple STAs which are also controllable and automatable. The key capability required is the ability to position those multiple STAs at different orientations relative to the AP so that MU-MIMO operation will be possible.

What to measure

The key metric for MU-MIMO testing is the aggregate system throughput. In the ideal case, this throughput will rise with the number of (independent) MU-MIMO clients, since each client should be able to get its own dedicated RF channel. That being said, in practice products do not actually work that way, since there are a limited number of MIMO streams available to use for both throughput and beamforming. Therefore, it is also useful to be able to look at the MU-MIMO throughput gains as a function of the number of users, and the number of spatial streams that each user requires for throughput. Shown in Figure 15.

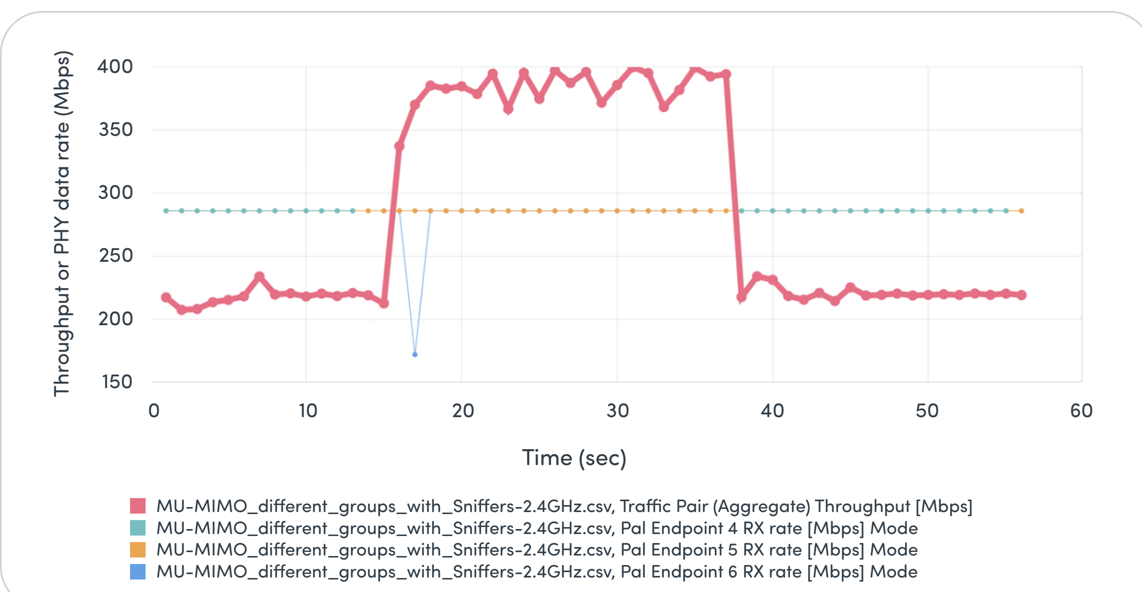

Figure 15

Target Wake Time (TWT)

Wi-Fi was originally developed as a wireless replacement for Ethernet connections (hence the early name of the Wi-Fi Alliance, the Wireless Ethernet Compatibility Alliance) so the original targets for its use were devices like desktop computers that were expected to have access to power, that is, to be "plugged in." Unlike the original designs for cellular standards, power saving was not a primary target because mobile devices were not anticipated to be the primary users of Wi-Fi technology. This changed quickly, and now mobile devices and battery-powered devices of all sorts are major users of Wi-Fi, which makes Wi-Fi's ability to run efficiently in a battery-powered device of critical importance.

Wi-Fi's power save technology has, therefore, improved as the generations of Wi-Fi technology have advanced, and Wi-Fi 6/6E brings another major advance with Target Wake Time, or TWT. TWT has in its sights not only mobile device operation, but IoT device operation as well, in which the expected battery lifetime moves from hours/days to weeks/months/years. As such, TWT is a major advance in power save operation over previous generations of Wi-Fi.

All previous generations of Wi-Fi power save operation were based around the beacon message sent by Wi-Fi APs. These messages are typically sent every 100 msec, and the general idea was that these regular transmissions could be used by STAs as a regularly scheduled time to send or receive data. At all other times, the STA could be "asleep", meaning that it could power down its transceiver, thus saving battery power.

In "legacy" power-save operation, the STA and the AP coordinate so that the AP will "hold" all messages for the STA until the STA is awake and ready to receive them. The AP transmits a DTIM (Delivery Traffic Indication Map) when it sends out beacons, and the STA can tell, from that map, whether the AP is holding any messages for it. If so, it will trigger the AP to deliver those messages, and it can also send any traffic to the AP that it wishes. See Figure 16. Obviously, based on this description, the STA needs to wake up fairly often to check for the DTIM messages. The frequency is configurable, but it is often several times a second. While this is effective at saving battery power, it is not even as good as most cellular technologies.

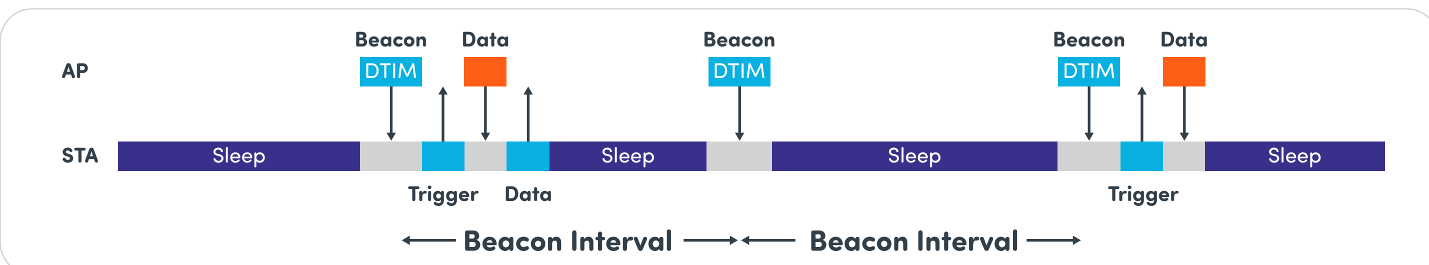

Figure 16

Once Wi-Fi became more embedded into mobile devices, and voice over Wi-Fi became more popular, another power save technology was implemented, this one called U-APSD, or Unscheduled Automatic Power Save Delivery. U-APSD flips the above script and puts the STA in charge of when packets are exchanged between the AP and the STA. In U-APSD, the STA can transmit to the AP whenever it has data to send, and the AP will hold traffic for the STA until that happens. See Figure 17. This has the advantage of not tying the STA so directly to the beacons, which means that a STA can go through several sleep/awake cycles within a single beacon period. This mechanism is useful for the kind of regular traffic generated by voice calls, in which packets are sent at the codec frame rate of, for example, 20 msec.

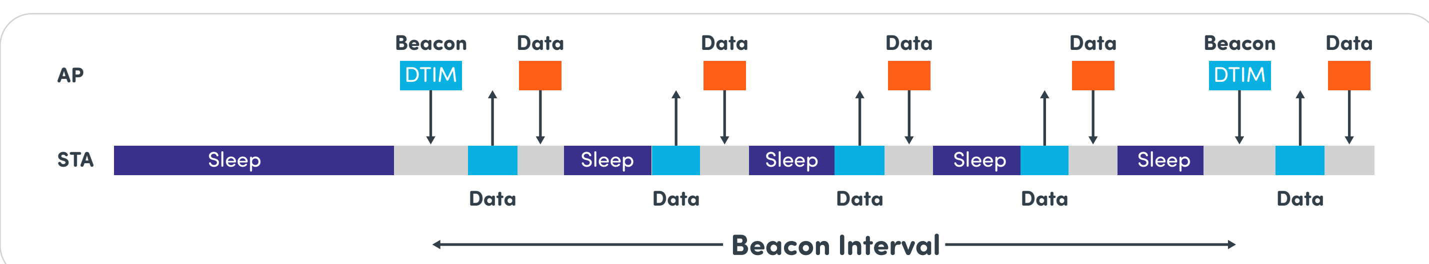

Figure 17

The new Wi-Fi 6/6E technology of Target Wake Time (TWT) takes another step in divorcing the power-save mechanism from the beacon timing. Part of the rationale is that Wi-Fi has moved beyond just being used in mobile devices; now almost every smart device contains a Wi-Fi radio. As a key technology supporting the Internet of Things (IoT) Wi-Fi needs to be able to support power-save operation for devices like wearable medical sensors, rain and moisture sensors, pressure sensors in industrial settings, etc. These devices, and many others, will be battery operated with the expectation that the batteries will last for months, if not years. So, Wi-Fi needed to become very efficient in the use of its radio, which is what TWT accomplishes.

In TWT, STAs and APs negotiate an explicit schedule to determine precisely when, and for precisely how long, a STA will wake up to send/receive messages. At all other times, the device can turn off its Wi-Fi radio. The concept is shown in Figure 18. This timing is completely independent of beacon timing and can also be completely different from one device to another. This has two distinct benefits, one of which has to do with scheduling. This makes TWT fundamentally different from, and more capable than, previous power-save technologies:

1. TWT allows devices to significantly increase their sleep time and conserve battery power. This is an improvement on previous technologies.

2. TWT can reduce contention between devices, which reduces the need to retry packets, saves power, and makes the network more efficient overall. In this way, TWT offers improved packet scheduling, which is something that no previous generation of Wi-Fi power save offered.

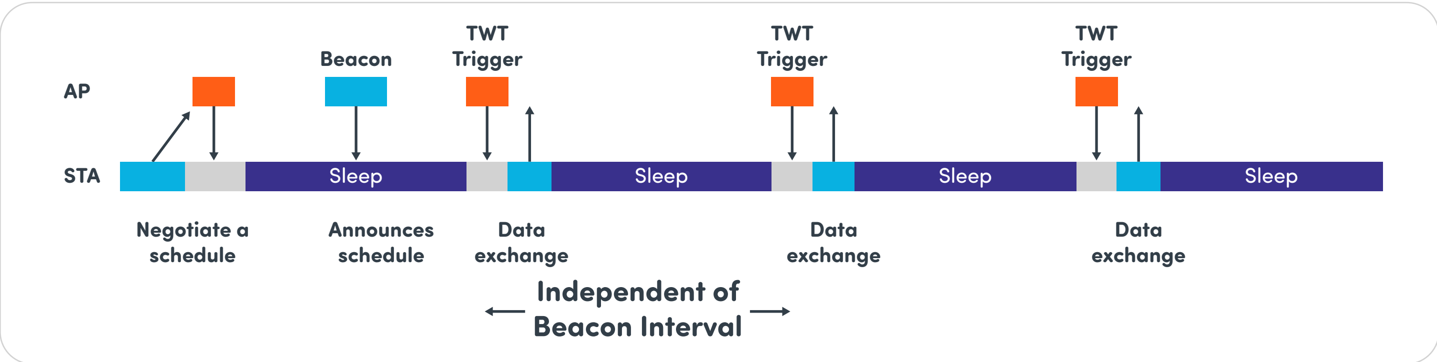

Figure 18

Testing TWT

What to test

TWT is a very flexible protocol, which allows the STA and AP to negotiate "Service Periods" (SPs). These may be service periods that are unique to a particular STA (an individual agreement), or they may be SPs that apply to a group of STAs (a broadcast agreement.) The way that the wake-up times are determined can also differ -- there are both implicit and explicit modes. In the implicit mode, the STA calculates the next wake-up time based on a schedule and its current known timing. In explicit mode, the STA will send messages to the AP and will, in term, receive wake-up time information from it.

And, finally, there are two types of data flows -- unannounced and announced. In the unannounced mode, APs can deliver messages to the STA as soon as the wake-up time has arrived, whereas in the announced mode, APs must wait for messages from the STA before they can deliver any stored messages. Given all these possible combinations, therefore, testing of TWT can be quite complicated. In general, though, a test needs to be able to instigate a TWT frame exchange (and this is done by the STA) and then to confirm that the AP is able to adhere to the negotiated schedule.

Testing challenges

The main complication related to TWT testing comes from the fact that, in this case too, the situation between APs and STAs is completely asymmetric. STAs need to initiate a TWT session; APs need to buffer traffic and adhere to the schedule. So in an AP under test scenario, the main challenge is having a STA device for which the TWT functionality can be controlled. That is you need to have a device that can initiate a TWT session with known parameters, and this is not something that is easily done by an off-the-shelf STA device.

How to test

For an AP under test, there are at least two levels of testing that can be carried out:

1. Sign-of-life testing
   Since TWT support is advertised in the HE MAC Capabilities information element, one simple test is to confirm that the AP is advertising "Responder Support", and that the STA is advertising "Requester Support." Shown in Figure 19.

```
HE MAC Capabilities Information: 0x10401a08010d
   .... .... .... .... .... .... .... .... .... .... .... ...1 = +HTC HE Support: Supported
   .... .... .... .... .... .... .... .... .... .... .... ..0. = TWT Requester Support: Not supported
   .... .... .... .... .... .... .... .... .... .... .... .1.. = TWT Responder Support: Supported
```

Figure 19

2. Performance testing
   Performance testing looks at the actual performance of TWT, and the details of what that testing might look like are described below.

What to measure

TWT testing requires that you be able to negotiate TWT sessions and observe the effect on packet exchange between the AP and STA. So, for example:

- Create a traffic session between the AP and STA
- Negotiate a TWT session
- Look in detail at the packet arrival statistics. Confirm that the sleep period, and the awake time, are as negotiated in the exchange.

These examples can be seen in Figure 20.

- The line plot (labelled STApal-W1) is showing throughput as a function of time. During the test we have initiated three separate TWT sessions. During each session we can see the throughput drop as the STA spends some of its time asleep. The first session is between times 5 and 10 seconds.
- The second session is between 15 and 20, and the third is centered at 25 seconds. For each of these sessions the impact on throughput is different because the TWT session parameters are different.
- A visual indication of these parameters can be seen in the lower half of Figure 20, where we show a bar chart of packets-per-second as a function of time. In this figure, again, the three TWT sessions are clearly distinguishable. (The time scales are different, so the sessions do not appear at the same times.) For example, at around time "6" we can see four large gaps during which no packets arrive (the STA is in sleep mode) and three narrow bands in which packets arrive (the STA is awake and sending/receiving packets.) Similarly we can see a different sleep/awake structure at around time "10", and yet another structure at around time "13".

With more than one STA and more than one negotiated TWT service period, confirm that the two service periods do not overlap, that is, that TWT is addressing channel contention as well as battery-saving sleep time. This is shown in Figure 21.

In this case we have 2 STAs with the same TWT session parameters (same sleep/awake timing), but in the packets-per-second bar chart we can see that the sleep awake times are offset relative to each other. One STA is shown in green and has awake times at 4, 8, 12, 16, etc., while the other is shown in red, and has awake times with a consistent offset to the left of the other STA's times.

BSS-Coloring

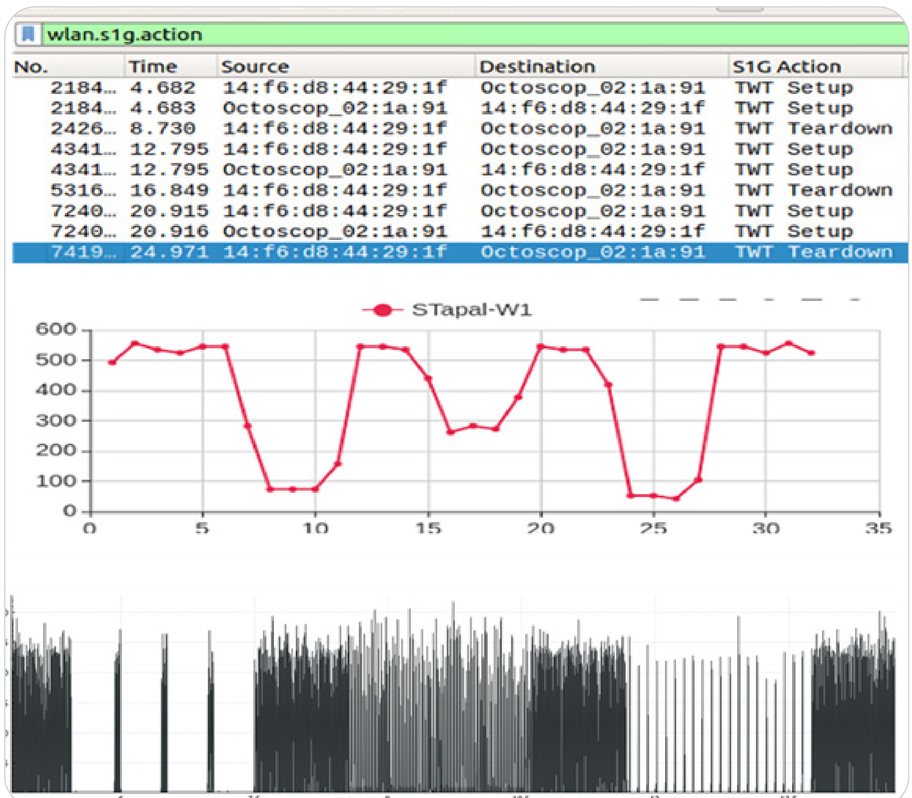

Figure 20

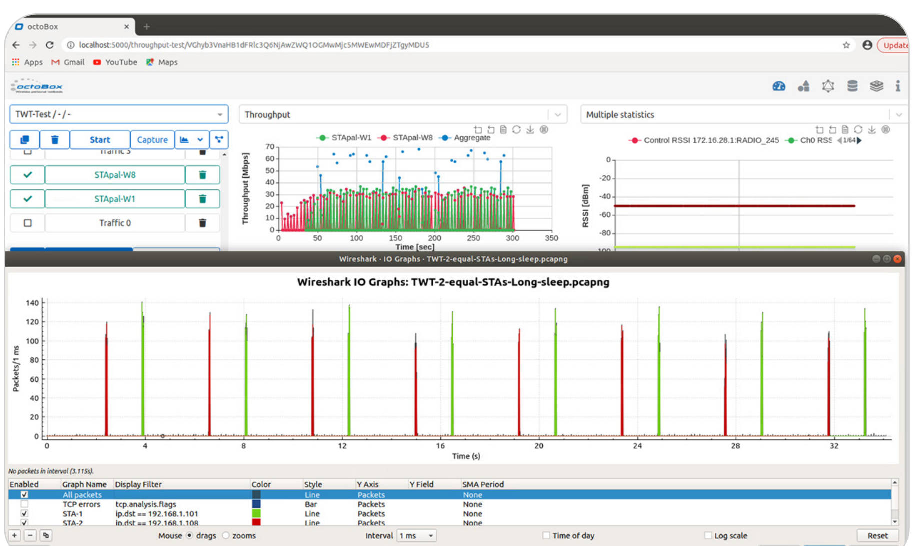

Figure 21

As Wi-Fi has become more and more widely used, there are increasing cases in which a user may find himself in a very densely crowded Wi-Fi environment, that is, one in which there are a large number of APs and STAs which can all "hear" each other. The concept of one radio "hearing" another in Wi-Fi is slightly complicated, and breaks down into two categories:

1. Preamble-detection: If a device is able to receive the Wi-Fi preamble of another device above -82 dBm, it will remain off the channel for the amount of time specified in that preamble (a time specified in what is known as the NAV, or network allocation vector.)

2. Energy-detection: If a Wi-Fi device senses energy above -62 dBm but is not able to find a preamble, the device will assume that it has detected something that is not a Wi-Fi device that is using the same channel, and it will remain off the channel until the energy detected falls below the energy-detection threshold.

The mechanisms outlined above are what make Wi-Fi so adept at operating in unlicensed frequency bands; the technology has built-in sharing mechanisms. While this is one of the reasons for the success of Wi-Fi, these sharing protocols can cause problems in densely deployed Wi-Fi locations. Let's see how.

Imagine the simple scenario shown in Figure 22. In this figure there are two APs (AP1 and AP2) each using the same channel (C). Each AP has attached STAs with which it needs to communicate.

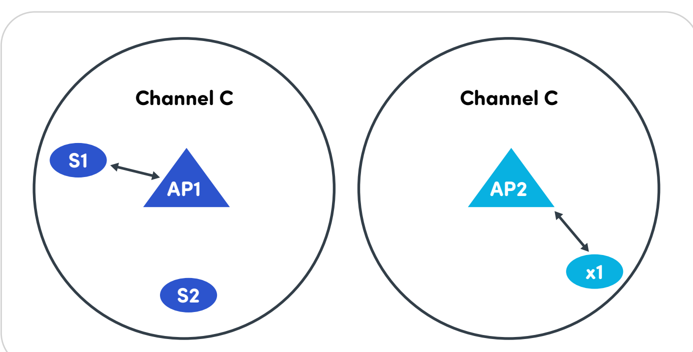

Figure 22

Assume also that AP1 and AP2 are close enough to each other that they can "hear" each other's transmissions. That is, they are within the preamble detection threshold limit. In this case, AP1 and AP2 will each get (all other things being the same) one-half of the airtime to use. They will politely share the channel with each other, each using it about one-half of the time. This means that the capacity of each AP will be about one-half of the expected capacity of the channel. (Notice that for more APs using the same channel and within range of each other, this reduction in capacity continues, linearly.)

The APs described above could be part of the same network (the same Basic Service Set, or BSS), or they could be part of different networks. Traditional Wi-Fi protocols do not distinguish between those two cases.

From the way the figure is drawn, the problem with the sharing mechanism as described may be obvious. Imagine if AP1 were to transmit to S1 at the same time as AP2 transmitted to x1. While it is true that AP1's transmission would cause interference on the AP2-x1 link, this interference would probably be very low. Let's assume, for the sake of discussion, that the interference lowered the SNR on the AP2-x1 link so that the MCS dropped from MCS 11 down to MCS 8. This would take a change of 6 dB for an 80 MHz channel in Wi-Fi 6. That would be a change from a 1200 Mbps PHY rate to an 865 Mbps PHY rate. That is a 28% reduction in throughput, which is significantly better than the 50% reduction achieved by sharing the channel.

BSS Coloring allows for networks to make a tradeoff between time on the channel, and signal-to-noise ratio. In some cases, as described above, it may be better to have access to the channel even with a little interference than it is to stay off the channel and operate in a pristine environment with no interference. BSS Coloring allows each network (BSS) to have a specific color associated with it. Now when a Wi-Fi device detects a preamble, it has another choice beyond simply remaining off the air. If the color of the packet indicates that the device is part of its own BSS, standard channel access mechanisms will continue to apply. However, if the color indicates that the device is part of a different BSS (that is, communicating with a different access point) the device can choose to either remain off the air, or to transmit at the same time, thus trading SNR for channel access.

Testing BSS Coloring

What to test

The primary metric for testing BSS Coloring is to look at the impact on throughput of co-channel usage when there are two networks. The BSS operations should be on the same channel, and then testing should be performed both with, and without, BSS Coloring, and for the BSS Coloring test it should be tested when the BSS Colors to are the same, and when they are different.

Testing challenges

The most challenging part of testing BSS Coloring is that BSS-Coloring requires that the devices be in a specific range of operation -- interference should be above the preamble detection threshold, but below energy detection. This requires the ability to precisely control the relative distance of the devices. So, two separate networks are required so that there can be two separate BSSes. And these networks need to support BSS Coloring and need to have the ability to either enable/disable BSS Coloring, and/or have the ability to directly configure the BSS Color, so that testing can be done with the same, and with different, colors.

How to test

1. Create two independent networks. Ideally the APs used for each network will be the same, with the same firmware, in order to avoid mixing other behaviors with the BSS Coloring testing.
2. Use programmable attenuation to completely separate the two networks.
3. Run traffic on each network and measure the aggregate network traffic.
4. Use programmable attenuation to bring the networks within range of each other.
5. Run traffic on each network and measure the aggregate network traffic.

What to measure

When the two networks are complete separated (no matter what the BSS Coloring status) we should see an aggregate throughput of twice the individual throughputs (and the throughputs on each of the BSSes should be the same.) See Figure 23. In this figure, the lower two lines show the throughput (as a function of time) for the two networks, and the upper line is the aggregate. So each network is getting about 1250 Mbps, and the aggregate is about 2500 Mbps.

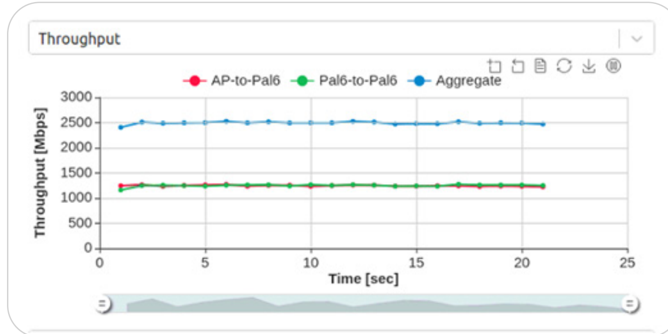

Figure 23

As the two networks are brought within range, there should be measurable differences between how the throughput behaves when BSS Color is operating, versus when it is not. Without BSS Coloring we expect the two networks to begin time-sharing as soon as they are within the preamble detection threshold. That is, there will be a very sharp drop in throughput, as the networks detect each other and begin time-sharing. With BSS Coloring operable, however, we expect that in certain cases the traffic between the two networks will share the airtime, as long as the color is different and the interference level is between the preamble detection and energy detection thresholds. In this case, we expect to see the networks continue to operate simultaneously for some time. While this may, at some point, affect the throughput due to added interference (reduced signal-to-noise ratio and, hence, lower MCS) the effect should be less pronounced than is caused by pure time-sharing. Eventually, as the networks get even closer together, time sharing will take over, even in the BSS-Coloring case, as shown in Figure 24.

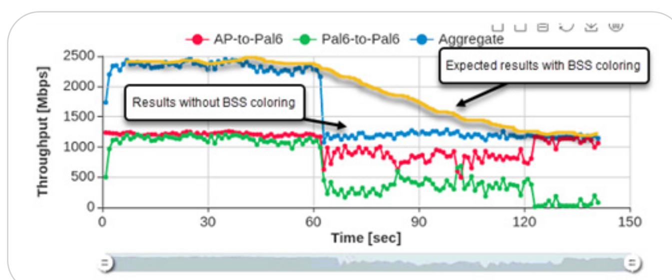

Figure 24

About Spirent

Spirent Communications (LSE: SPT) is a global leader with deep expertise and decades of experience in testing, assurance, analytics and security, serving developers, service providers, and enterprise networks. We help bring clarity to increasingly complex technological and business challenges. Spirent's customers have made a promise to their customers to deliver superior performance. Spirent assures that those promises are fulfilled.

For more information visit: www.spirent.com

Summary

Testing Wi-Fi 6/6E is more complicated than earlier generations but necessary with the emergence of a more robust Wi-Fi platform. With APs having to schedule users simultaneously, a controllable, repeatable environment, like a testbed fitted with many STAs, is needed for OFDMA, MU-MIMO, TWT, and BSS coloring test results. Using a small anechoic chamber environment, testing for Wi-Fi 6/6E can produce multiple, accurate test scenarios without the cumbersome and uncontrolled results of field testing.

Spirent's octoBox automated testbeds are the perfect solution for complex Wi-Fi 6/6E performance testing, incorporating RF chambers and instruments, such as the Pal-6E and STApal-6E, controlled by an integrated server with a browser-based UI and a complete API for test automation. Spirent's state-of-the-art solutions make it possible to conduct repeatable lab tests that have real world relevance, lower costs, and improve test program outcomes while minimizing risk.

A trusted provider for over 25 years, Spirent is the industry leader in Wi-Fi testing technologies. Our team of world-renowned experts are here to help. To learn more about Wi-Fi 6/6E performance testing, Contact Us.

Americas 1-800-SPIRENT
+1-800-774-7368 | sales@spirent.com

Europe and the Middle East
+44 (0) 1293 767979 | emeainfo@spirent.com

Asia and the Pacific
+86-10-8518-2539 | salesasia@spirent.com

Rev A | 10/21 | www.spirent.com