

Navigating Wi-Fi 7

A Deep Dive into Next-Gen Advancements

Table of Contents

Abstract.....3

Wi-Fi Legacy3

Introduction of Wi-Fi 7.....4

New Wi-Fi 7 Testing Challenges23

Testing Systems 24

References.....28

Appendix A.....29

Appendix B..... 30

Abstract

This white paper explores Wi-Fi 7's advancements through a testing system overview and feature introduction. It discusses legacy Wi-Fi, Octobox testbeds, and Octobox software with REST API and ScriptManager integration. Challenges in testing Wi-Fi 7's capabilities are noted.

Wi-Fi 7's features encompass 4096 QAM, new preamble design, 320 MHz bandwidth, multi resource units for concurrent users, and enhanced quality of service (QoS). This includes Wi-Fi 7's latency improvements, restricted service periods for latency assurance, restricted target wake time, and triggered peer-to-peer transmission.

A core focus is multi-link operation (MLO), detailing architecture, reliability vs. latency trade-offs, multi-link device types, channel access, and performance. Multi-link discovery contrasts legacy methods, covering association and security setup.

The white paper serves as a concise guide for Wi-Fi enthusiasts, spotlighting its revolutionary features and testing insights for future wireless technology.

Wi-Fi Legacy

Traditionally, each new generation of Wi-Fi has consistently delivered higher data rates, primarily focusing on individual users. However, Wi-Fi 6 introduced a wide range of new features and enhancements aimed at improving the overall user experience in scenarios with multiple simultaneous users. Notable advancements included orthogonal frequency-division multiple access (OFDMA), which significantly improved latency for users with light data loads. Target wake time (TWT) was specifically designed to schedule traffic in a time-division multiple access (TDMA) manner, but it also found utility in power saving by enabling devices to sleep between scheduled data deliveries.

Additionally, multi-user multiple-input multiple-output (MU-MIMO) was enhanced to support simultaneous uplink and downlink connections, thereby serving multiple users concurrently. The introduction of basic service set (BSS) coloring helped facilitate spatial re-use in dense environments like airports or sports stadiums. While Wi-Fi 6 did extend support for higher data rates through the expansion of bandwidth to 160 MHz and the introduction of 1024 quadrature amplitude modulation (QAM), the primary focus remained delivering improved service to multiple simultaneous users.

In this paper, we will explore the new functionalities that Wi-Fi 7 brings to the forefront and delve into how Keysight is preparing to provide testing solutions tailored for Wi-Fi 7. Wi-Fi 7 is based on IEEE 802.11be, extremely high throughput (EHT), while the term ultra-high reliability (UHR) has been earmarked by IEEE as the designation for the upcoming Wi-Fi 8, which is already in development.

Introduction of Wi-Fi 7

The Wi-Fi generational naming convention is a Wi-Fi Alliance innovation that helps customers distinguish which flavor of Wi-Fi they have. Previously, Wi-Fi flavors were denoted by the last letters of the IEEE task group developing the standard. In the case of Wi-Fi 7 this is the IEEE 802.11be task group. The distinction between these two names is important because the IEEE task group may develop features in the standard that are not adopted by the Wi-Fi Alliance.

The 11be task group is still working on completing the standard, with an indication of the timeline in Figure 1. Nevertheless, silicon vendors have already developed chipsets that are seeing early adoption into the market.

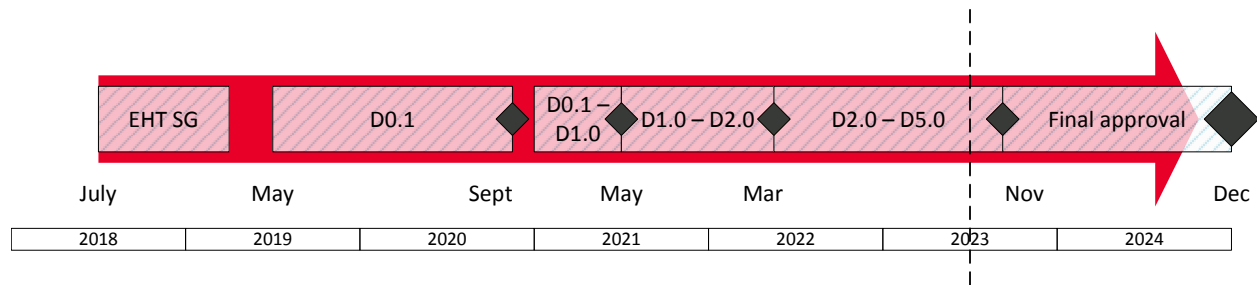


Figure 1. Wi-Fi 7 development timeline

IEEE named 11be ultra-high throughput (UHT) and as UHT implies, throughput is the game. Higher throughput is achieved by two mechanisms. Firstly, the physical layer (PHY) rate is increased by using 4096 QAM and widening the bandwidth to 320 MHz to provide 30 Gbps.

Latency is a key priority for Wi-Fi 7 and there are several mechanisms to improve latency, including a mechanism to provide deterministic latency.

Secondly, the concept of a multi-link device (MLD) is introduced. Multi-link operation (MLO) has several different modes, but essentially it uses a framework to coordinate the multiple radios in a device to act as one.

The combination of these mechanisms leads to throughputs that, in practice, top at 11 to 12 Gbps. And just in case we wanted more antennas, 802.11be supports 16 antennas, although this is unlikely to find its way into many actual implementations of Wi-Fi 7.

4096 QAM

Wi-Fi 7 introduces modulation coding scheme (MCS) 12 and 13, which uses 4096 QAM to reach 12 bits per symbol. 4K QAM, as it is sometimes called, leads to a very dense constellation, as can be seen in Figure 2.

The density of this constellation places several challenges on the Wi-Fi radio, both on the transmission (Tx) side as well as the receive (Rx) side.

4K QAM theoretically requires an error vector magnitude (EVM) ratio of -38 dB or better to ensure that the constellation points are as close to pinprick size as possible. This imposes what some view as almost military grade linearity requirements on the Tx and Rx.

Channel noise effectively blurs the pinpricks so that they overlap with one another, making demodulation impossible so 4K QAM operation is limited to operation where the signal-to-noise ratio (SNR) is better than 42 dB.

Achieving 42 dB SNR means being very close to the AP to limit path loss, but vendors are achieving successful operation at distances of up to 18 feet. This is because they exploit the diversity gains provided by beamforming in a multipath environment to increase the SNR.

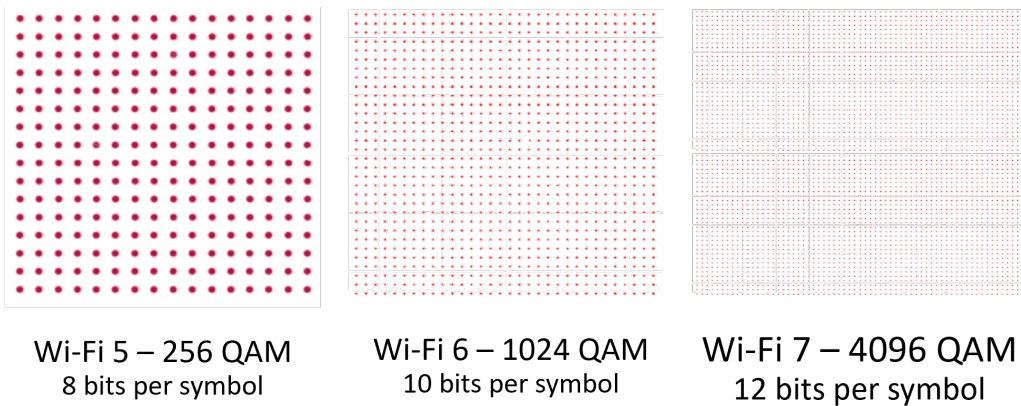


Figure 2. Wi-Fi generational increase in constellation size

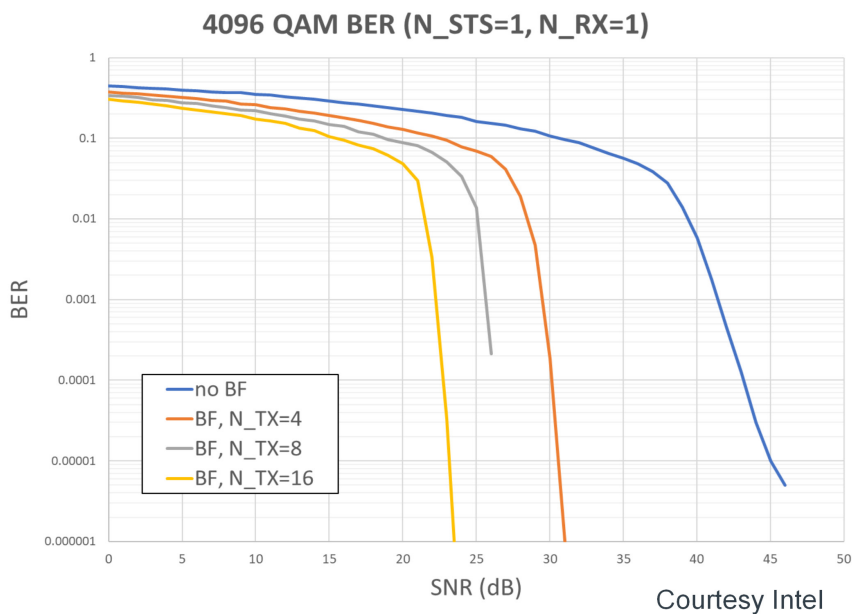


Figure 3. Diversity gains using multiple antennas

The diversity gains achieved by using 4 antennas when compared with 1 antenna is nearly 13 dB, which is highly significant and in line with communications theory. Increasing to 8 or 16 antennas gives only incremental gains.

4K QAM needs very good SNR to work, and so path loss between chambers will be critical for proper operation. Several chipset vendors claim that beamforming is absolutely necessary, but it is not clear if they are talking about Rayleigh fading channels or additive white Gaussian noise (AWGN) channels. Some experimentation will be necessary.

New Preamble Design

Every new amendment to 802.11 kept backward compatibility but introduced new additions to the preamble. This arrangement introduces overhead, probably is not sustainable, and makes the auto detection more and more difficult. 802.11be defines a new universal signal field (U-SIG) in the preamble. The U-SIG is two symbols in length¹ and contains both version-dependent and version-independent information that facilitates legacy as well as future PHYs.

The extremely high throughput (EHT) field contains information for the receiver to demodulate the physical layer protocol data unit (PPDU) and is followed by the EHT-STF and EHT-LTF. The STF and LTF are used to determine timing sync and perform channel estimation for correct demodulation of the transmission.

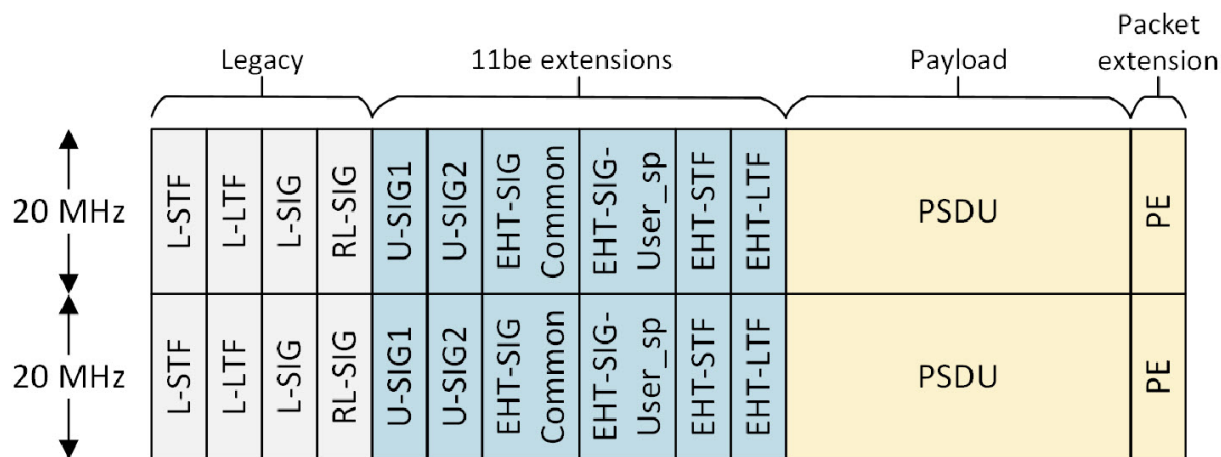


Figure 4. EHT preamble

¹ Symbol length is 4 us to keep backward compatibility.

320 MHz Bandwidth

With the recent opening up of the 6 GHz band, more countries are adopting it, as seen in Figure 5. For the latest view, visit [Regulations Enabling 6 GHz Wi-Fi..](#)

The 6 GHz band provides several GHz of contiguous spectrum and allows 320 MHz channels, which were difficult to achieve in 5 GHz.

The introduction of 6 GHz has led to what some believe to be a more consistent channel numbering system so that, for example, 20 MHz channels have different numbering than 40 MHz channels and there are two sets of three 320 MHz channels, chosen to fit with the regulatory requirements of some countries.

With such wide channels, it is possible that interferers may be present in the band, or indeed there may be sections of the channel that are disallowed because of incumbents.

The probability of encountering interferers is proportional to the width of the channel, making noise-free very wide channels very difficult to find. Wi-Fi 7 brings “preamble puncturing” whereby, in a wide channel, those problematic sections of the channel are not used (i.e., punctured). This greatly improves the chances of using wide channels.

Countries Enabling Wi-Fi in 6 GHz (Wi-Fi 6E)

- Adopted 5925-6425 MHz
- Adopted 5925-7125 MHz
- ▨ Adopted 5925-6425 MHz, Considering 6425-7125 MHz
- Considering 5925-6425 MHz

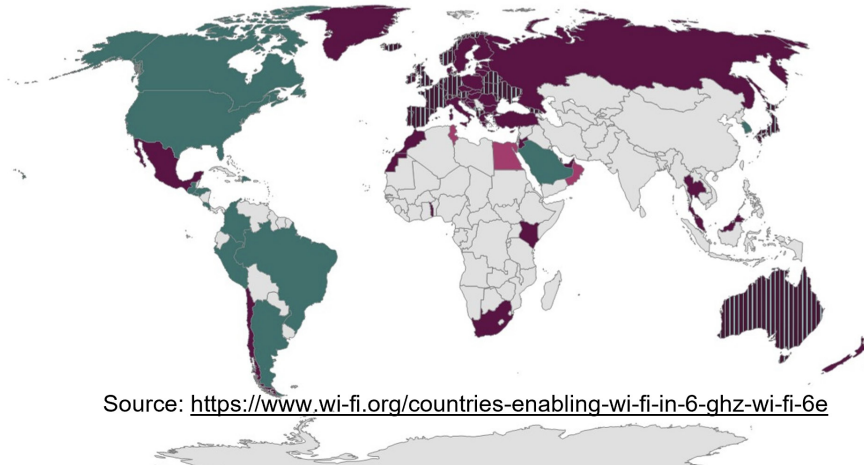


Figure 5. Worldwide 6 GHz adoption

Puncturing resolution is 20 MHz, and puncturing is indicated in the Disabled Subchannel Bitmap, where each bit is set to 1 for those 20 MHz subchannels that are punctured. For example, the Disabled Subchannel Bitmap for channel 95 would be 0000 1000 0011 0000.

The ability to puncture in the 6 GHz band is a mandatory requirement. Puncturing in the 5 GHz band is optional.

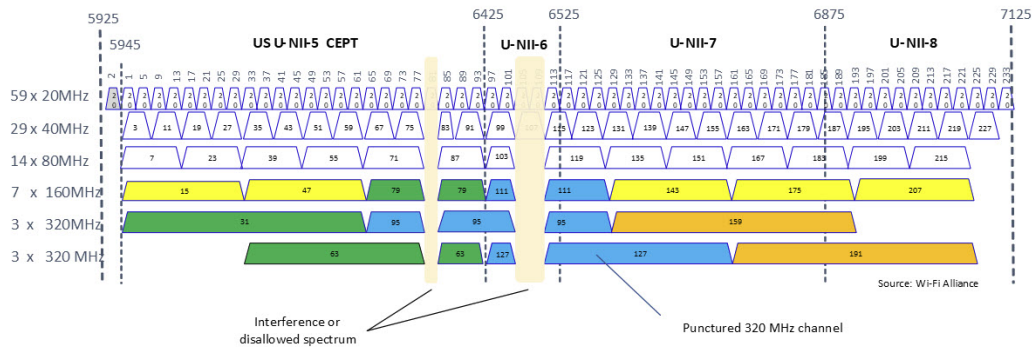


Figure 6. Illustration of puncturing

Multi RU

Wi-Fi 6 OFDMA allows multiple users to access the channel simultaneously by allocating resource units (RUs), or portions of the spectrum to each user. This is done dynamically on a packet-by-packet basis and has the advantage of reducing channel contention, which reduces latency.

However, for wide channels, when there is a large discrepancy in the amount of data each device needs, there may be RUs that are not allocated. This wastes bandwidth.

Wi-Fi 7 allows the unused spectrum to be reclaimed by allowing up to two RUs to be allocated to a device. It is worth noting that the combinations of RU sizes are limited to keep things practical — so not all combinations of all RU sizes are possible.

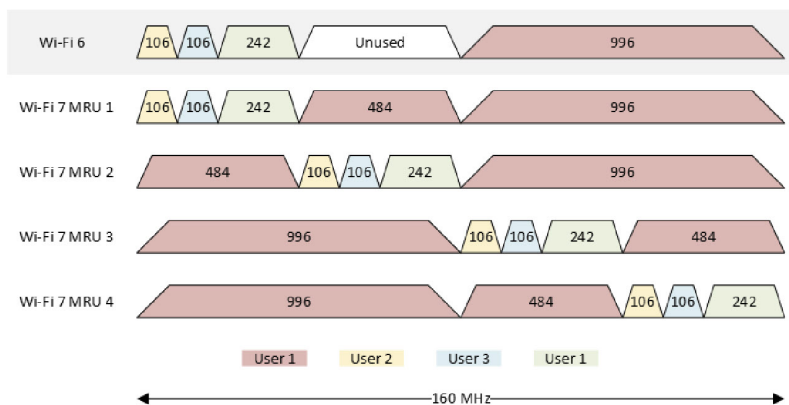


Figure 7. Examples of Multi RU operation

512 MPDU Compressed Block-Ack

Wi-Fi 7 introduces a feature known as “512 compressed block-ack,” enabling the transmitter to consolidate as many as 512 MAC protocol data units (MPDUs) into a single frame. Likewise, it allows the receiver to acknowledge up to 512 MPDUs within a single block acknowledgment (BA) frame. This is a notable enhancement compared to Wi-Fi 6, where aggregation was limited to 256 MPDUs per frame. The introduction of 512 compressed block-ack significantly reduces protocol overhead and enhances the transmitter’s performance, particularly when transmitting data at high throughput PHY rates using a 320 MHz channel width and MCS 13 modulation in Wi-Fi 7.

Enhanced QoS

Wi-Fi 6 Latency Improvements

More people are using internet applications such as teleconferencing and gaming that are latency sensitive. The Wi-Fi industry has sought to control latency by various means for many years. Data rates have increased, which helps a bit, but with growing popularity, multiple simultaneous users tend to worsen this performance.

The main reason for this degraded performance is that many users are seeking access to the channel, albeit with small packets, but at a high rate. The result is confusion on the channel with so many devices having to back off, perhaps even into exponential backoff.

Wi-Fi 6 introduced OFDMA as a way of making channel usage more efficient by assigning RUs to each user, grouping the data, and serving multiple users for a single channel access.

OFDMA may or may not have marginal gains in terms of throughput depending upon the type of traffic. However, the primary purpose of OFDMA is to reduce latency for small packets, and it achieves this objective very well, as can be seen in Figure 8.

In Figure 9, we see the cumulative probability distributive functions (CDFs) of 12 non-AP STAs sending isochronous UDP traffic with a mean data rate of 30 Mbps, a standard deviation of 2 Mbps, and a frame rate of 60 FPS to simulate simultaneous audio / video transmissions.

Note how, in the legacy case, the distribution is spread over a wide range. The tails of the PDF reach as much as 150 ms, and in the truncated view seen here, none of the STAs achieve even 90% certainty of reaching 30 ms one-way delay (OWD). This is accompanied by large OWD standard deviations, which would be extremely uncomfortable for the user.

This is in direct contrast to the case where Wi-Fi 6 employs OFDMA. Here, the PDF is much more orderly; the 97th percentile is generally below about 20 ms. The standard deviation of the spread is minimized, but there is still a significant number of samples in excess of 30ms.




Figure 8. OFDMA minimizes channel contention for multiple users (source: WFA)

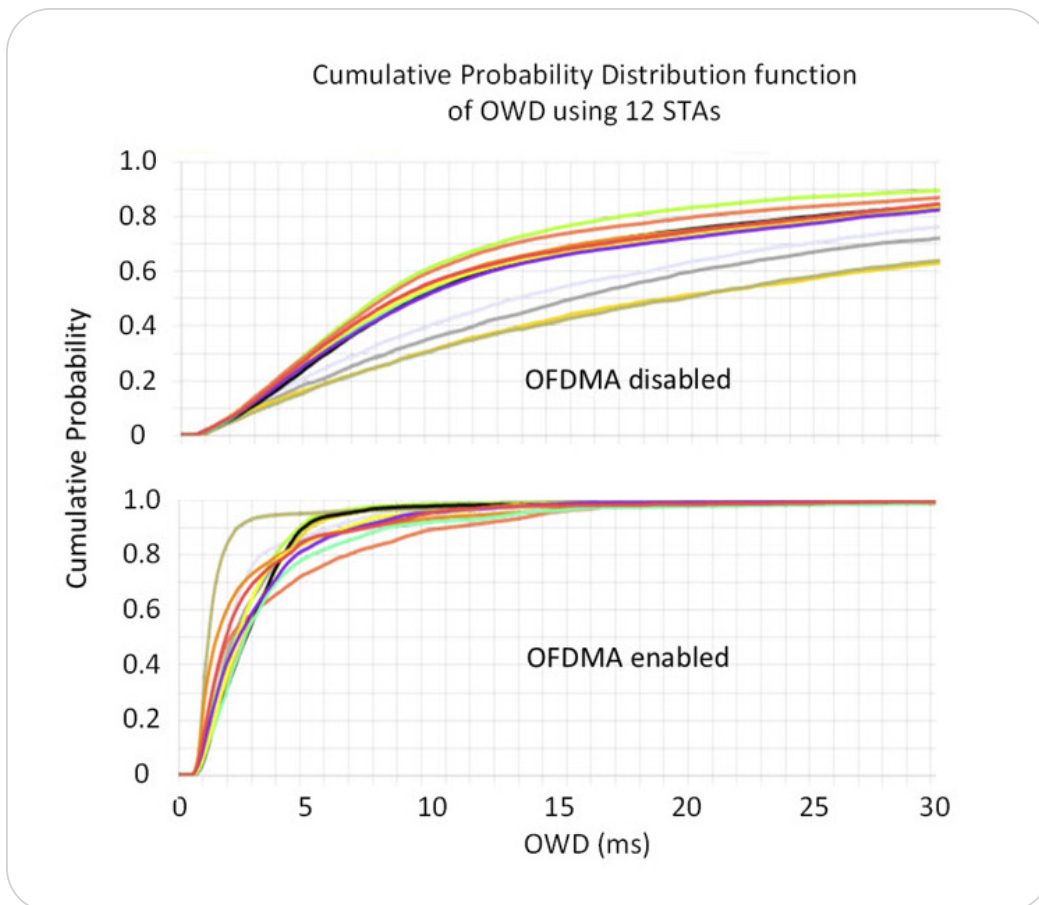


Figure 9. Wi-Fi 6 latency improvement compared to legacy

Wi-Fi 7 Restricted Service Periods for Latency Guarantees

While OFDMA significantly improves the OWD situation for multiple users, it does not guarantee performance. There are many applications such as gaming, industrial control, and IoT where delay must be deterministic and guaranteed to be less than some threshold.

Wi-Fi 7 introduces the concept of restricted service periods (RSP), which are periods of time where only certain devices may access the medium if they belong to an appropriate membership group.

In this scheme, the non-AP STAs with special QoS needs send their requirements to the AP, which in turn advertises the RSP that only members of each designated RSP group are allowed to use. This guarantees channel access for short periods at regular intervals where the latency sensitive data can be exchanged.

Consider the practical example of a gamer who requires very regular repeated access for small packets — mouse clicks. The requirement is for a deterministic latency of 99.9% to be less than 10 ms. This can be set up as the diagram indicates with 1 ms RSPs repeating every 10 ms. Additionally, an EHT AP can advertise quiet periods where non-AP STAs may not access the medium.

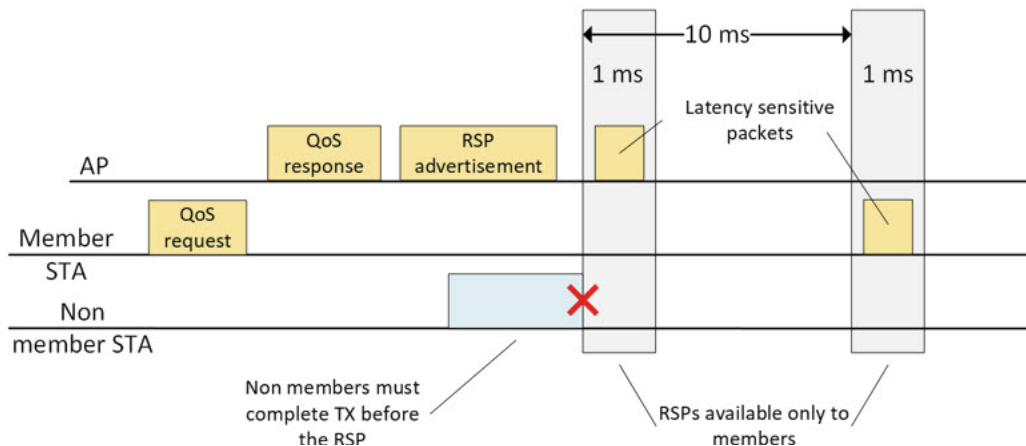


Figure 10. Enhanced QoS using restricted service periods

Figure 11 is an exemplary illustration of the PDFs and CDFs of packet-by-packet OWD of an AP sending to 12 non-AP STAs, one of which is a member of the group discussed above.

Here we see its OWD is confined to a narrow interval of delay about 10 ms wide. Its probability of falling within this region is better than 99%, as can be seen in the CDF, making the OWD of this device highly deterministic. While the other non-AP STAs sometimes achieve better than 10 ms OWD, the probability of this is much lower; indeed, the distribution, which is truncated in this example, extends beyond 100 ms. These delay spreads are common in large populations of devices because the distribution of legacy OWD is highly non-deterministic.

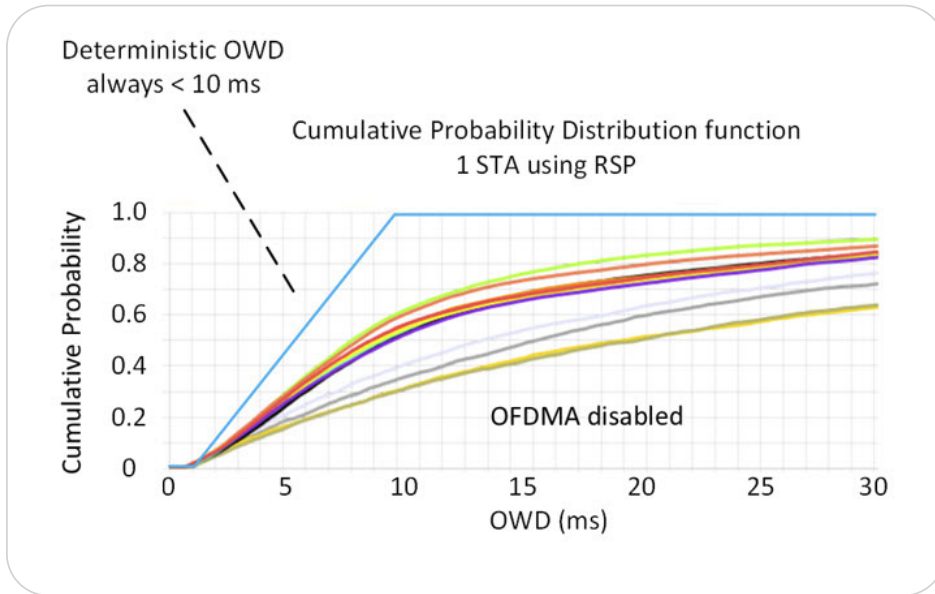


Figure 11. Comparison of OWD spreads for devices with and without deterministic QoS

Restricted Target Wake Time

Wi-Fi 6 introduced target wake time (TWT) as a way of scheduling transmissions in a kind of TDMA-like approach. TWT is most often used as a power-saving mechanism, allowing the STA to sleep until the targeted time. The Wi-Fi 7 restricted TWT extends capabilities to provide reservation mechanisms for more predictable latency and generally higher reliability for latency-sensitive traffic.

The EHT AP announces service periods (SPs) that a non-AP STA can sign up to become a member. This means that it can attempt channel access in these restricted slots as a limited member with fewer other devices with which to contend. The member non-AP STA or the AP must cease transmission before the end of the restricted period.

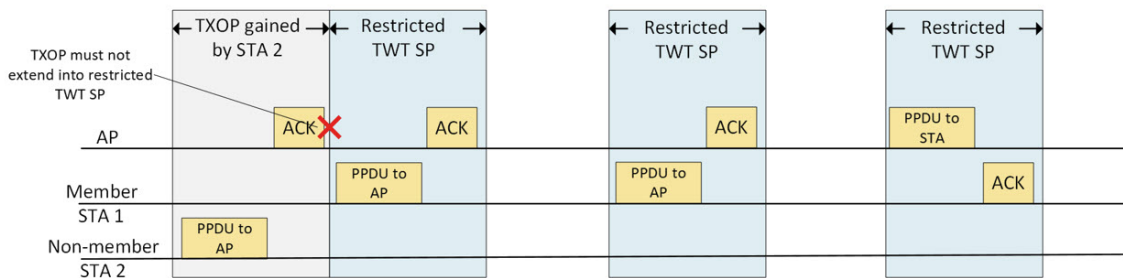


Figure 12. Restricted TWT SP operation

Triggered Peer-to-Peer (P2P) Transmission

There are many applications where communication is between two peers. Printing from a phone, virtual reality applications, and Miracast are examples. Since the content is generated on one STA and consumed by the other STA, there is no real value in relaying it through the AP. Indeed, relaying through the AP introduces delays and doubles the airtime used, which impacts other users needlessly. Although, it could be argued that the AP has a valuable role insofar as it “controls” the network and the relaying function helps coordination between devices, minimizing collisions and the hidden STA problem.

Wi-Fi 7 recognizes the value of managing the network of many users and introduces triggered P2P operation where the AP issues a multi-user request to send (MU-RTS).

Whereas the legacy RTS was a first-person request to all users, “I would like to have the channel for a Tx,” the MU-RTS is a third-person request, “STA 1 and STA 2 would like to have the channel for a Tx.”

So, the AP reserves a TXOP for STA 1 and STA 2 to communicate, as shown in Figure 13.

This mechanism efficiently manages the other devices because they will hold off for the TXOP while STA 1 sends directly to STA 2, thus minimizing collisions and avoiding the data occupying the medium twice.

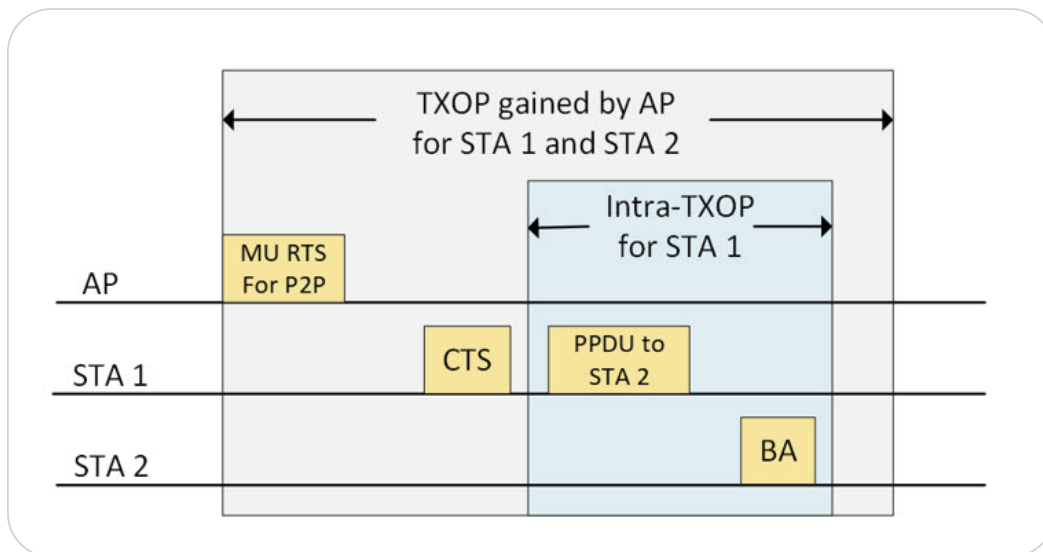


Figure 13. Triggered P2P operation

Multi-Link Operation

Wi-Fi is currently able to sustain multiple radio links on different radios, either in the same band or across different bands; however, there is no coordination between the radios, and there are inefficiencies.

Wi-Fi 7 takes a coordinated approach to using multiple links, resulting in a number of benefits. The multi-link device (MLD) views the individual radios as one radio with one MAC address and schedules communication according to the requirements at the time.

For example, multi-link operation (MLO) could be used to achieve additive throughput by aggregating the links and distributing the load over the different links. MLO can also be employed to enhance reliability by sending redundant data on each link, increasing the probability that a packet will get through correctly.

Latency improvement and load balancing can be achieved by allocating links according to the user priority or degree of time sensitivity of the data.

Multi-Link Architecture

The multi-link device (MLD) consists of several so-called affiliated radios. Each radio has its own PHY and lower MAC, but they are coordinated by a unified upper MAC (UMAC). The UMAC approach was chosen to preserve the long-established 802.11 MAC PHY concepts. From the outside, the MLD device presents itself as a single MAC address; thus, the higher layer protocols view it as a single device.

As in the legacy MAC, the UMAC handles fragmentation, packet reassembly, duplication detection, dynamic link switching, and ACKs. The sequence numbers for the packets are generated uniquely from the same sequence number space, which simplifies the above tasks. It also handles packet re-transmission, which can occur on any link, regardless of the original link the packet was transmitted on.

Association and authentication on each link may occur independently or jointly. In the former case, capabilities for each link are exchanged on each respective link. In the latter case, the capabilities of each can be communicated in a combined fashion on one link.

MLDs with multiple radios allow links to operate simultaneously; however, this is not a requirement of an MLD. An MLD could have only one radio, but the multiple links that it accesses would be one at a time. Some radios capable of multiple streams can configure each RX chain as a 1x1 on each channel / band / link and listen to incoming packets on each channel. Data transmission can only be done on one link at a time. This provides most of the benefits of a multi-radio MLD, especially when receiving data.

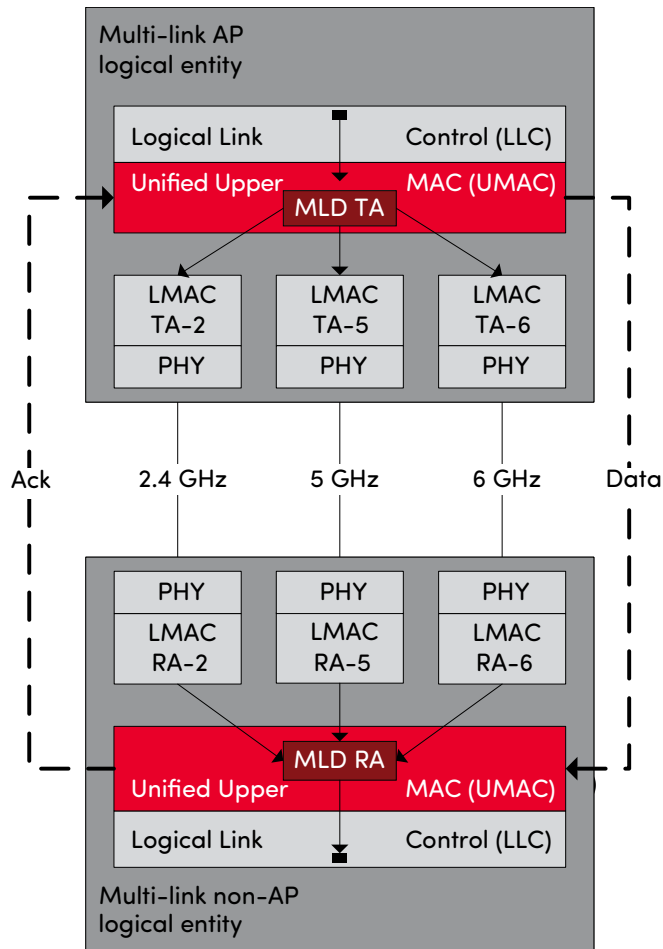


Figure 14. MLD architecture

Reliability Versus Latency

Packets can be allocated to links in various ways depending on the requirements of the application. Where reduced latency or increased throughput is required, unique data packets are assigned to each link.




Figure 15. MLO reduced latency or increased throughput

In situations where high reliability is required, packets can be repeated on each link. The UMAC performs the task of removing duplicate good packets, notes packets that did not get through on any link (lost packets), and composes an appropriate Block Ack at UMAC level.




Figure 16. MLO for increased reliability

These two modes can be used dynamically as channel conditions change, reducing latency in good conditions and extending range in poor conditions due to effective channel diversity.

Note the elegance of the UMAC using sequence numbers for all packets from the same space. This means that the UMAC does not need to be signaled about which mode is being used; indeed, these modes can be intermixed at will. The UMAC will determine duplicates based on their sequence number and discard them.

Types of MLDs

Multi-link devices can be created in various configurations depending upon device requirements. All MLDs operate on multiple links. Links are logical, so an MLD may not have individual radios for each link. If an MLD does have multiple radios, they may or may not be able to Tx and Rx simultaneously. Figure 17 is a hierarchical representation of MLD types as summarized in Table 1.

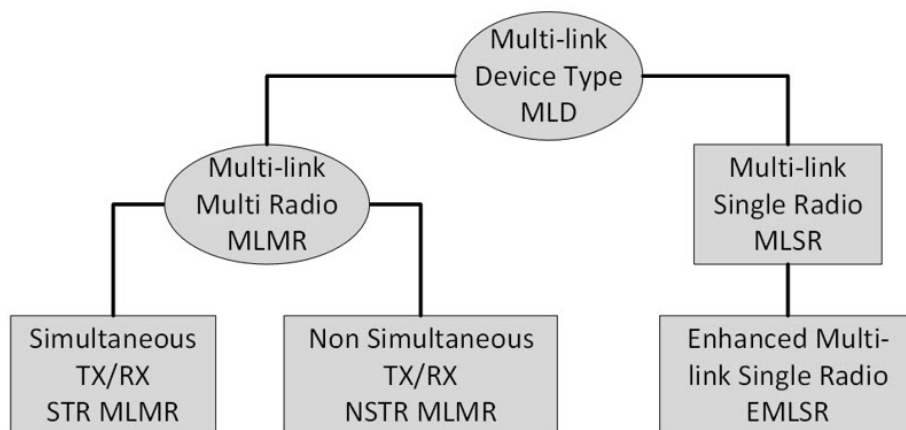


Figure 17. Multi-link device hierarchy

Type of MLD	Number of Radios	Characteristics
Multi-link single radio (MLSR)	1	There could be multiple links, but only one link can be active at a time
Enhanced multi-link single radio (EMLSR)	1	Ability to listen to multiple links simultaneously, but Tx on only one link at a time.
Simultaneous transmit and receive multi-link radio (STR MLMR)	> = 2	Simultaneous Tx/Tx, Rx/Rx and Tx/Rx over multiple links
Non-simultaneous transmit and receive multi-link multi-radio (NSTR MLMR)	> = 2	Simultaneous Tx/Tx or Rx/Rx over multiple links
Enhanced multi-link multi-radio (EMLMR)	> = 2	MLMR with additional to dynamically reconfigure spatial multiplexing on each link

Table 1. Multi-link device type

Multi-Link Channel Access

Multi-link operation allows the individual links to operate asynchronously of one another. This has advantages because the probability of finding a TXOP for an individual link is higher than finding a TXOP for all links simultaneously.

To be able to operate effectively, each local transmitter should not interfere with the local receivers. For many devices, this requirement can be achieved given that the frequencies of the links are generally widely spread — 2.4 GHz, 5 GHz, 6GHz. These MLDs are called simultaneous transmit receive (STR) devices.

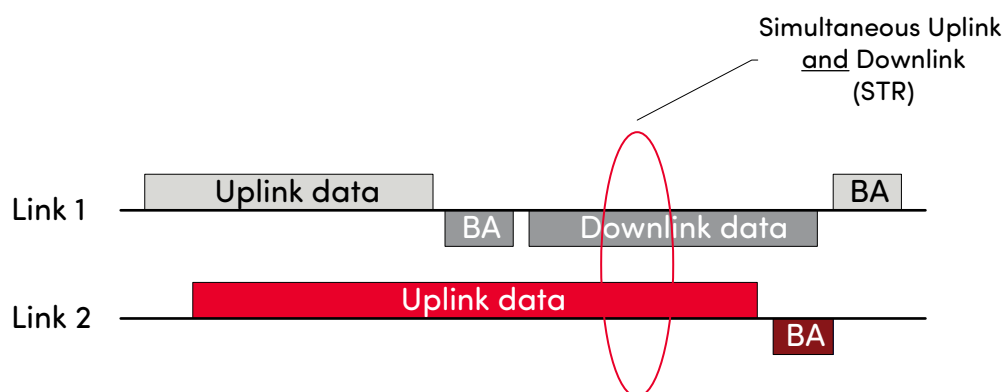


Figure 18. Simultaneous asynchronous operation STR

However, there may be devices where RF or physical constraints result in local Tx/Rx interference — these are termed non-STR (NSTR) devices, and, in this case, data on the links must be synchronous.

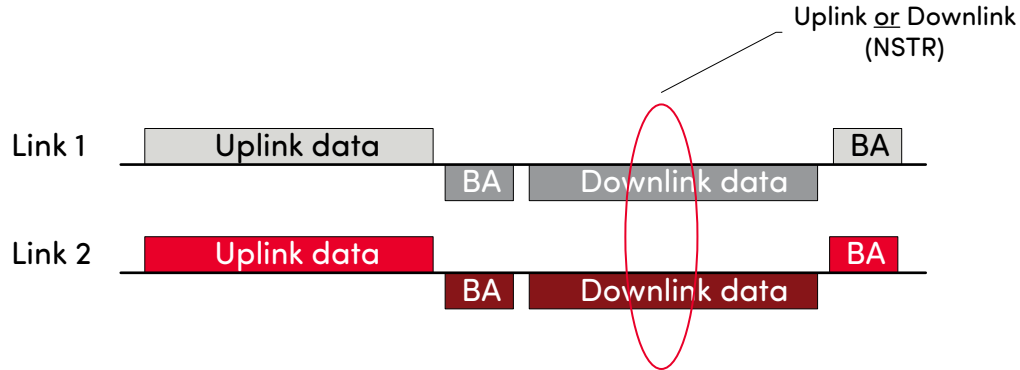


Figure 19. Non-simultaneous synchronous operation NSTR

Multi-Link Performance

As expected, the different types of MLDs offer different quality of performance. OBSS is a common real-life type of interference and shows the throughput of a legacy single radio link in comparison to a two-link multi-link single radio (MLSR) and a two-link simultaneous Tx/Rx multi-link multi radio (STR MLMR). Clearly, both MLDs offer improvement; the STR MLMR device performs at roughly double the single link. Interestingly, the MLSR offers improved performance, too. This is because using multiple links improves frequency diversity.

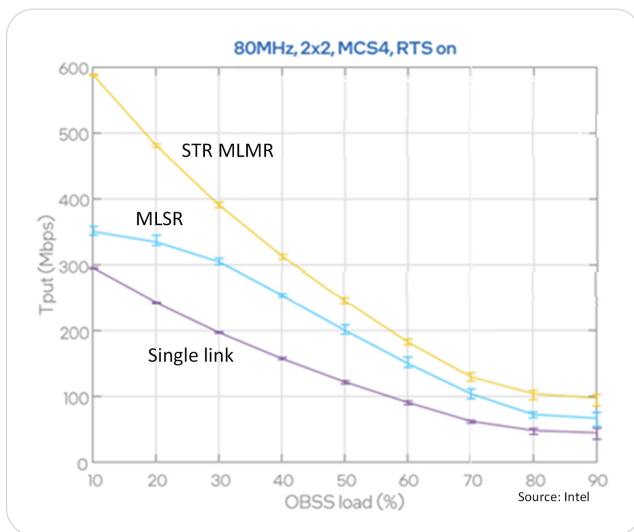


Figure 20. Performance comparison of different MLDs

Multi-Link Discovery

Legacy Link Discovery

In legacy Wi-Fi, a non-AP STA discovers an AP by either passively listening for beacons or by actively sending probe requests on each of the different channels. Passive scanning is a slow process; active scanning is somewhat faster because if the STA does not get a response very quickly, it moves to the next channel.

These mechanisms work reasonably well in 5 GHz, but in 6 GHz there are potentially 59 20 MHz channels to scan, which would be very time-consuming and not acceptable. Active scanning represents a significant waste of bandwidth because although the requests are short, they are management frames sent at a low data rate. Actively scanning 59 channels wastes a lot of airtime, given that potentially only one of the 59 requests would be successful.

6E Has Four Methods for AP Discovery

1. Fast initial link setup (FILS) discovery frames, which are short, contain only the bare details of information, but are broadcast by the AP every 20 ms. This is a passive scan method, but it is faster than looking for beacons.
2. Unsolicited probe response (UPR) frames containing the same information as a beacon are transmitted every 20 ms by the AP. This is faster than listening for beacons but very wasteful of airtime.
3. Preferred scanning channels (PSCs) are the only channels in 6 GHz where probing by the STA is allowed. These channels are spaced at 80 MHz, so allowed channels are 5, 21, 37, etc.
4. Out of band discovery (OBD). There are very few 6 GHz-only APs on the market — they nearly all have 2.4 and 5 GHz — which provides an opportunity for optimization. Here either FILS or UPR can be used in 5 GHz to furnish the AP information for 6 GHz. Indeed, this information could also be included in the beacon, adding a small amount of extra information to the often long beacon and reducing overhead.

Many non-AP STA manufacturers are very conservative about allowing access to 6 GHz because their devices could be highly mobile and potentially be used in a geographic location where 6 GHz is not allowed. They tend to require use of the OBD method because the passive methods are slow, and they cannot use PSC since they do not know if 6 GHz is allowable. Some manufacturers take it one step further — they listen out for any APs in 5 GHz, taking note of their country codes and deduce allowable 6 GHz access from the most restrictive country code they hear.

The discovery phase for a Wi-Fi 7 MLD is more complicated because it needs to obtain information for multiple APs in multiple links.

Wi-Fi 7 Multi-Link Discovery

Recall that an AP MLD may have several affiliated APs. Wi-Fi 7 uses several mechanisms to simplify the discovery process. Each affiliated AP is required to broadcast unsolicited probe response frames and include a reduced neighbor report (RNR) in its beacons. These provide only basic information such as BSSID, channel, operating class, etc., but do not provide the complete information such as operational parameters and capabilities to limit time on the air.

A non-AP STA that wishes to obtain detailed information about the AP MLD uses a multi-link probe request (MLPR) to obtain the full information from any one of the APs affiliated with the AP MLD and will be returned the complete information about all the links the AP MLD supports.

This mechanism can be thought of as an expanded version of OBD mentioned above and represents the most efficient means of discovering the multiple links that the AP MLD supports at one time on one link.

Figure 21 is an illustration of Wi-Fi 7 MLD discovery. In this case, the non-AP MLD passively scans 2.4 GHz because there are fewer channels to scan and minimizes the initial discovery phase.

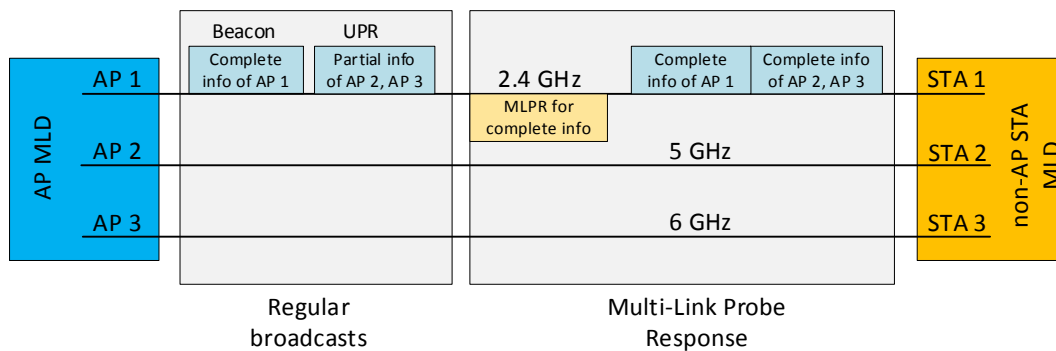


Figure 21. Example of MLD discovery

STA 1 will discover the complete information for AP1 and partial information about AP2 and AP3. This should be sufficient for the non-AP STA MLD to determine if the AP is of interest, whereupon it can request complete information about all the APs and prepare to move to the next step, which is association.

Multi-Link Association

Association in legacy Wi-Fi consists of the STA and AP sharing information, such as capabilities using a request-response protocol, and allows association of one link at a time. Associating MLD devices by performing this procedure would be too time-consuming, so Wi-Fi 7 has introduced a more efficient method where the (re)association request and response procedure is augmented to include the information of all links at the same time. This procedure occurs only once on only one of the links.

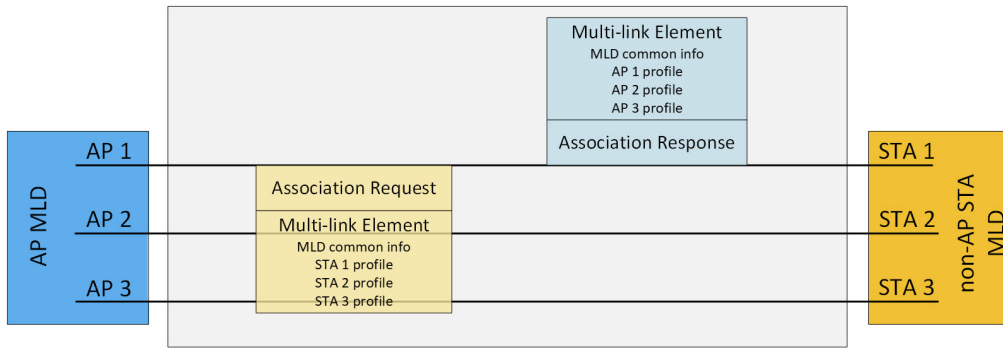


Figure 22. Multi-link association

Multi-Link Security Setup

Wi-Fi 7 requires the implementation of WPA3™ for safeguarding communications among Wi-Fi 7 devices. WPA3 security for Wi-Fi 7 devices requires:

- More robust key management procedures
- Protected management frames
- A cryptographic mechanism to verify the contents of beacon

Wi-Fi 7 also requires support for Wi-Fi Enhanced Open™, which provides unauthenticated data encryption for users on open networks.

Security is set up using the WPA3 four-way handshake. The pairwise master key (PMK) is known to both the authenticator, in this case AP 1, and the supplicant, in this case STA 1.

The AP generates an ANonce, which it sends to the STA in message 1. The STA generates an SNonce and generates a pairwise temporal key (PTK) and sends this together with its SNonce and the message integrity code (MIC) to the AP in message 2.

The AP then derives the PTK and, if needed, the GTK, IGTK, and the BIGTK (defined below). These are encrypted using the key encryption key (KEK), which is derived from the PTK and sent to the STA in message 3. This process is sometimes called wrapping.

Once this is done, the STA can install all the keys and responds with message 4, whereupon the AP installs the set of keys. This completes the security setup for AP 1 and STA 1 in link 1.

The other links, 2 and 3, will all use the same PMK and PTK and packet numbering (PN) space.

The group temporal key (GTK) is used to protect data traffic, the integrity group temporal key (IGTK) is used to protect management frames, and the beacon integrity group temporal key (BIGTK) is used to protect beacon frames.

In Wi-Fi 7, each link has a unique set of these keys. The different APs affiliated with the AP MLD use different GTK / IGTK / BIGTK's in each link. Each AP and the corresponding associated non-AP STA maintains a single PN / IPN / BIPN for each set of keys in each link.

Since these keys have, in our example, been set up in Link 1, AP 1 will generate a unique set of keys for Link 2 and Link 3, and they will be provisioned by the group key handshake protocol.

These keys, as the name implies, are temporal and are updated from time to time using the group key handshake.

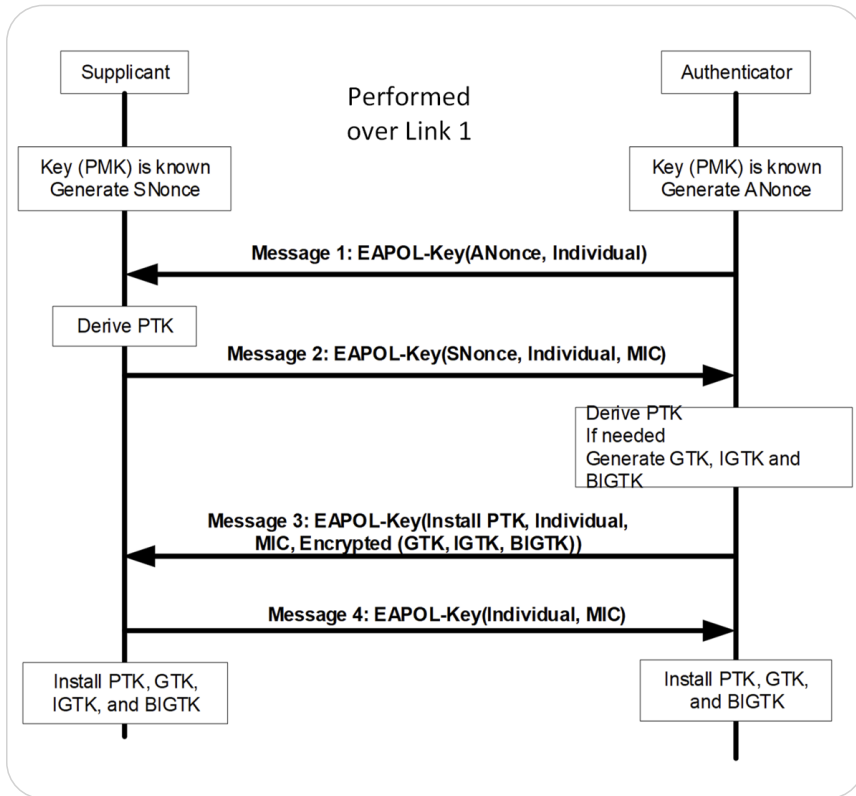


Figure 23. Four-way handshake for security setup

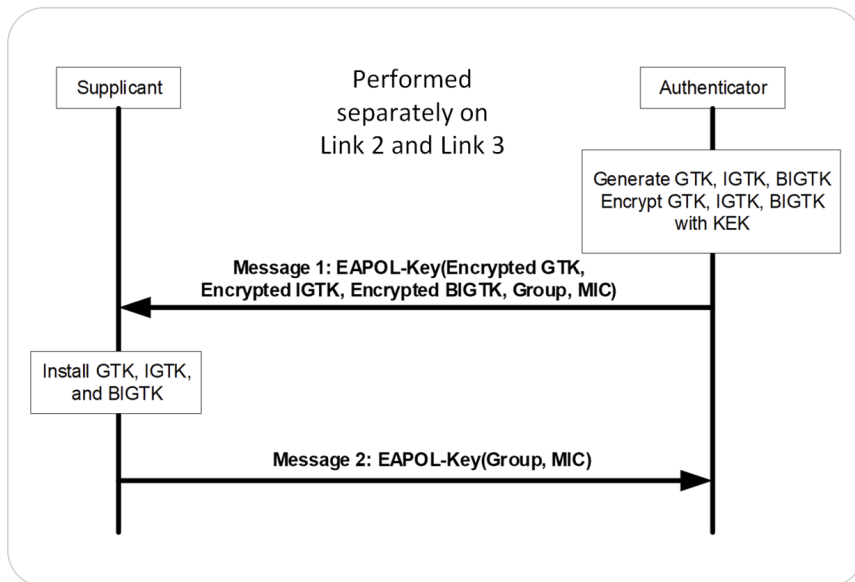


Figure 24. Group key handshake

NSEP

Wi-Fi 7 also has a reservation mode for national security and emergency preparedness (NSEP), where the AP verifies the authority of non-AP STAs to use the NSEP priority access. This access operates in an on-demand fashion where it can be invoked by either the AP or the non-AP STA.

The NSEP features in Wi-Fi 7 prioritize authorized stations for wireless medium access. Wi-Fi 7 APs with NSEP indicate priority access through beacon and probe response messages, and STAs associate with these APs to access priority channels. NSEP allows secure activation and deactivation of priority access based on critical needs. This mechanism builds on Wi-Fi Multimedia™ (WMM®) concepts and assigns relative priorities, ensuring that NSEP users' devices have better access to the wireless medium.

New Wi-Fi 7 Testing Challenges

Many of the new Wi-Fi 7 features, such as 320 MHz bandwidth, multi RU, and QoS enhancements, can be thought of as iterative extensions or enhancements to many of the Wi-Fi 6 features and are easily accommodated with simple additions to existing test methodology.

The new 4096 QAM requires extremely good signal quality to function correctly. Good signal quality comprises three parts. The transmitter must be extremely linear so as not to distort the constellation — this quality is assessed by performing an error vector magnitude (EVM) measurement. Signal to noise (SNR) must be better than 43 dB, which requires a very quiet RF environment with minimal RF path loss. Another factor often overlooked is the channel tracking algorithm in the chipset. A real RF channel can exhibit fading in frequency, fading in time, and is

subject to angular rotation. These impairments must be removed or corrected by tracking the channel state, so clearly, the quality of the channel tracking algorithm is very important.

Multi-link operation represents a significantly large area where enhancement is needed. MLO uses multiple radio links to form a single logical data path. This means that observing activity of a multi-link device requires simultaneous observation (sniffing) and measurement (KPIs) of each radio link. These observations and measurements must all be exactly synchronized to a central time base for correct data fusion.

Wi-Fi 7 has several new protocols and packet types, which must be decoded in a sniffing capture. Fortunately, many of these dissectors are already in place in the open-source Wireshark community.

Testing Systems

Traditional testing methodologies in cabled environments introduce certain complexities, particularly when dealing with MU-MIMO or multiple spatial streams. These complications were somewhat manageable for legacy Wi-Fi implementations. However, with Wi-Fi 6 and 7's emphasis on accommodating numerous devices within a network, as well as multiple overlapping uncoordinated networks (OBSS), the setup process for cabled testing becomes increasingly convoluted and burdensome.

Octobox Testbeds

The Octobox system rose to this challenge by using small shielded anechoic chambers to isolate the devices so that the RF paths could be controlled. The Octobox system also introduced directional antennas, which, with appropriate antenna placement, is key to being able to establish spatial diversity and overcomes many of the subtle complications mentioned earlier.

With the Octobox testbed, it is easy to connect a station (STA) with two antennas to an AP with four antennas in a very natural way, without having to re-wire things. This system has become an industry de facto standard, which is used by hundreds of users and has already been adopted in standards organizations such as the Wi-Fi Alliance, Broadband Forum, ETSI, and the Open Wi-Fi organizations.

The Octobox testbeds for Wi-Fi 7 have several enhancements to leverage this technology. Specifically, RF paths have been redesigned to give lower RF path loss so that 4096 QAM operation can be achieved.

The Octobox SynchroSniffer uses IEEE 1588 precision time protocol to synchronize each of its several sniffer probes so that their individual captures can be merged in a time aligned manner for analysis in Wireshark. This methodology will allow seamless sniffing of MLO operation.

The Octobox software is already equipped to measure KPIs of multiple devices simultaneously and synchronize the data for data fusion and storage, so measuring MLO performance is already technically achievable.



Figure 25. Octobox Stack-Max testbed




Figure 26. Octobox small anechoic chamber with turntable

Octobox Software

To deal with the multiple requirements of Wi-Fi 6, more highly instrumented devices were introduced, sometimes in the basement of a chamber termed a smartBox. Another good example is the palBox, which contains 16 standalone STAs necessary for MU-MIMO and OFDMA testing, but also contains a device that can field a large number of virtual STAs (vSTAs). The Octobox Stack-Max testbed could have well over two thousand devices, each with their own IP address. Scaling to this extent requires serious consideration for the testbed graphical user interface (GUI) if it is to be intuitive to use.




Figure 27. A view of the Keysight Octobox user interface

The GUI is organized in four main panes. The first shows an iconized view of each traffic pair in use and gives a visual indication of the health of the device, its current connection, and other information. When dealing with many devices, it is useful to quickly make sure they are all acting as expected. The pane on the left shows the configuration of the traffic pairs and other instruments such as turntable, attenuators, sniffers, and so on. On the top, the attenuation of all the RF paths is shown in iconized form so that, with a single glance, one can see which RF paths are open and which are closed.

Finally, the central pane displays many real-time performance indicators such as throughput, MCS, NSS, one-way delay, and a host of other performance indicators that are configurable. All these performance indicators are stored in a database for later offline analysis if needed.

The GUI is an important tool for setting up and running tests, and it is especially useful for iterative debug sessions where different parameters and configurations can be managed very easily, and the changes in operation can be seen in sniffer captures.

REST API

In many cases, testing is done under automation. These are custom tests that are often written as Python scripts, and the Octobox software has a rich, easy-to-use REST API, which is well documented using Swagger. The REST API is not program language specific, so it is ideal for a wide variety of programming languages such as Matlab, Octave, Jenkins, and others.

ScriptManager

Managing a large suite of automation scripts can sometimes become burdensome, and Keysight has introduced ScriptManager to organize these tests. ScriptManager runs on a dedicated ScriptMachine, which enables easy configuration of the test, and the collection of results, and even plotting of graphs in real time, using JSON files. ScriptManager provides easy-to-use scheduling of tests with various error recovery mechanisms to deal with failed tests and keep executing the remaining tests so that it can be used unattended in regression or overnight runs.

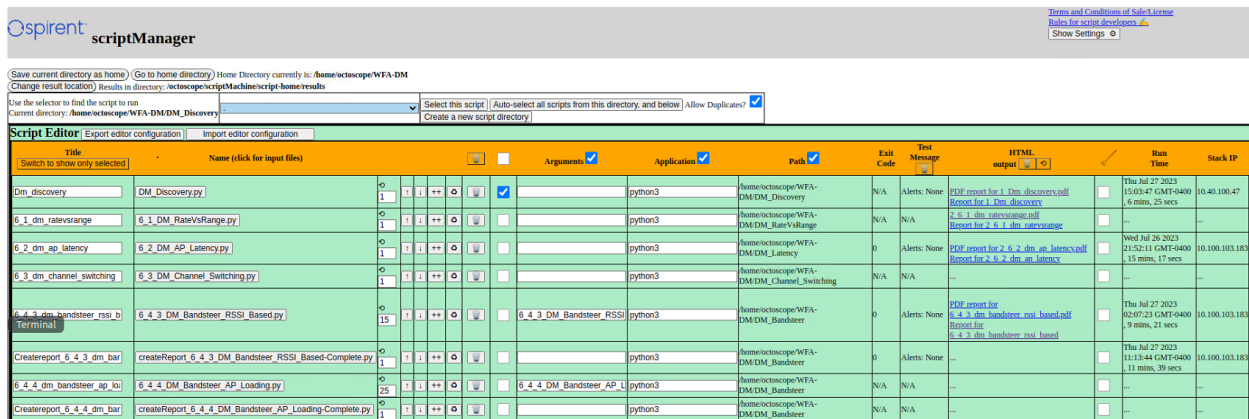


Figure 28. ScriptManager provides organization and automation of test suites

References

1. [Wi-Fi Unleashed: Wi-Fi 7, 6 GHz, and beyond](#). Carlos Cordeiro, PhD. Intel Fellow and Wireless CTO, Intel.
2. [Key advantages of Wi-Fi 7](#). Mediatek.
3. [Current Status and Directions of IEEE 802.11be, the Future Wi-Fi 7](#). Evgeny Khorov, et al.
4. [Next generation Wi-Fi, Wi-Fi 7 and beyond](#). Carlos Cordeiro, PhD. Intel Fellow and Wireless CTO, Intel.
5. [Are you ready for Wi-Fi 7—the next generation of Wi-Fi \(Analyst Angle\)](#). RCR Wireless.
6. [Overview and Performance Evaluation of Wi-Fi 7](#). Cheng Chen, Carlos Cordeiro, PhD, et al.
7. [Our innovative Wi-Fi 7 solutions set the standard for next-generation Wi-Fi](#). Qualcomm.
8. [Wi-Fi 7 Technology White Paper](#). New H3C Technologies Co., Ltd.
9. [RUCKUS is embracing the Wi-Fi 7 revolution—and here's why you should too](#). Ruckus Wireless.
10. [HPE Glossary: Wi-Fi 7](#). HPE.
11. [11-19-0314-01-000m-beacon-protection](#). Emily Qi, et al.
12. [11-19-0773-08-00be-multi-link-operation-framework](#). Po-Kai Huang, et al.
13. [11-20-0751-01-00be-multi-link-setup-clarifications](#). Rojan Chitrakar.

Appendix A

An Update to Wi-Fi Terminology

It is sometimes useful to consolidate the new terminology. Legacy Wi-Fi used the terms STA and AP where the role of station and access point were clear, the STA associates to the AP. As Wi-Fi has progressed, there are now cases where APs connect to other APs, and a tightening up of terminology is necessary.

- An AP always connects to a STA. We distinguish between a non-AP STA (e.g., laptop) and a STA (which could be part of an AP).
- A multi-link device (MLD) is a device with more than one associated STA. It has a single MAC data service access point (SAP) and logical link control (LLC).
- An AP MLD is an MLD where each STA inside the MLD is an AP.
- A non-AP MLD is an MLD containing a number of non-AP STAs.
- A multi-link multi radio (MLMR) MLD can operate on multiple links simultaneously.
- MLD devices may have less radios than links. For example, the definition of an MLD allows a single radio device to operate on a single link at a time. A multi-link single radio (MLSR) device cannot operate on more than one link simultaneously. However, many MLSRs have at least 2x2 MIMO capability and 802.11be defines an enhanced multi-link single radio (EMLSR) that can use the 2x2 functionality to listen to two links simultaneously, allowing the device to perform clear channel assessment (CCA) and receive data frames on the link that is otherwise idle to achieve enhanced throughput and latency reduction that is close to a concurrent dual radio device.
- Devices are further categorized into simultaneous transmit and receive (STR) and non-STR (NSTR).
- An enhanced MLMR (EMLMR) device can configure different spatial multiplexing on each link of the MLMR.

Appendix B

Multi-Link Architecture — Diving a Bit Deeper

It is sometimes instructive to investigate the details of an MLD device to understand the functional blocks and look at the relationship between an MLD and the individual affiliated devices. Figure 30 is an illustration an AP MLD with two affiliated APs.

The incorporation of an MLD brings in the unified upper MAC as discussed earlier. Functionally, it is equivalent to the legacy upper MAC except for the last two blocks.

This means that the device can be used in two different ways. Firstly, as an MLD with its own unique U-MAC address, the incoming data is mapped to the different links in a way that is invisible to the upper layers. Or it could be used in legacy mode by using the individual radios with their own MAC addresses, which are different to the MLD MAC address.

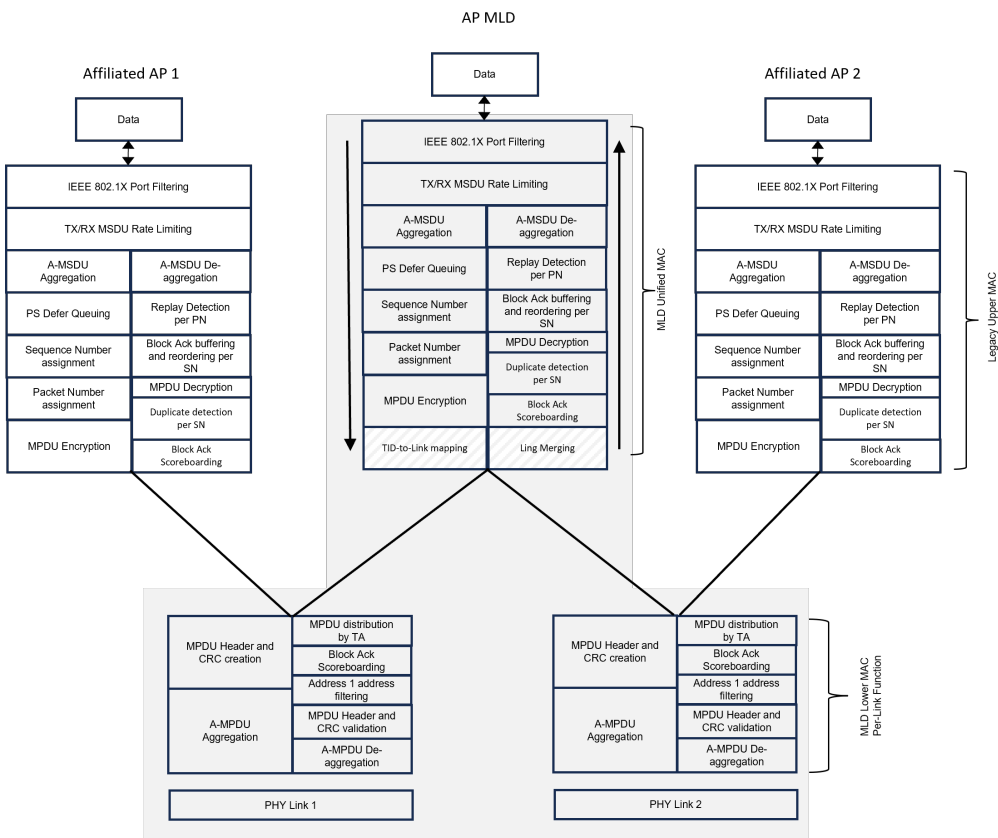


Figure 29. Legacy MAC and U-MAC illustration

Possibly the most important of these MLD U-MAC functions is the traffic identifier (TID)-to-link mapping block.

This is the block that takes the fully assembled MPDU and sends it to the appropriate lower MAC and PHY link. If increased reliability is required, this block could send the MPDU to both link lower MACs so that it is sent on both links similar to Figure 16.

Conversely, if minimum latency is required, it could split up the MPDU into parts and send each part to a different link as indicated in Figure 15. In the case where the MPDU is split, each portion has its own header and CRC applied so that on the receive path the link merging function can join them up again into the original MPDU.

Keysight enables innovators to push the boundaries of engineering by quickly solving design, emulation, and test challenges to create the best product experiences. Start your innovation journey at www.keysight.com.



This information is subject to change without notice. © Keysight Technologies, 2026, Published in USA, June 1, 2026, 3126-1207.EN