
WHITE PAPER

802.11AX



TABLE OF CONTENTS

THE WI-FI MARKET AND THE GENESIS OF 802.11AX	3
TECHNICAL FEATURES OF 802.11AX	7
USAGE MODELS: WHAT CAN WE DO WITH 802.11AX?	27
BACKWARDS COMPATIBILITY, DEPLOYMENT CONSIDERATIONS AND UPGRADE STRATEGIES	28
MANDATORY AND OPTIONAL FEATURES	28
PERFORMANCE ESTIMATES	29
CHANNELIZATION	32
CONCLUSION	32
APPENDICES (EXISTING FEATURES AND NEW 802.11AX FEATURES DEFERRED TO WAVE 2)	33

THE WI-FI MARKET AND THE GENESIS OF 802.11AX

Wi-Fi is by now the established way to access the Internet, whether at home or at work, from PCs or cellphones. In 2018, 19 years after the first meeting of the Wi-Fi Alliance, Wi-Fi will carry more than 50% of all Internet traffic.




Figure 1a: Wi-Fi Device Annual Shipments





Figure 1b: Wi-Fi Device Cumulative Shipments

With around 18 billion Wi-Fi devices shipped, 8 billion still in use and 3 billion new ones added every year, it is difficult to find anywhere without a Wi-Fi signal. Even cellphone networks, which have been improving speeds and capacities with the LTE build-out, small cells and flat-rate data plans, rely on Wi-Fi to meet the traffic requirements of their subscribers. A cellphone today without integrated Wi-Fi would be unthinkable.



STRATEGY ANALYTICS


Source: Strategy Analysis, February 2018

Figure 2: Average Monthly Data Usage by Operator
January 2018 (Android users, in MB)

Wi-Fi standards originate in the Institute of Electrical and Electronics Engineers (IEEE), where the 802.11 working group meets 6 times a year, with many interim conference calls for specialist task groups, to update and extend the technical standards that underpin Wi-Fi. Once the IEEE has completed the standard, the focus shifts to the Wi-Fi Alliance, the industry trade forum that owns the 'Wi-Fi' trademark, where a series of plugfests support the drafting of a test plan and interoperability certification program. This is how the industry ensures that Wi-Fi clients work with Wi-Fi access points, across all the different vendors in the ecosystem.

The last major 'PHY' or physical-layer certification was 802.11ac, with 'wave 1' commercial shipments commencing in 2014 and 'wave 2' shipments in 2016. But the work on 802.11ac goes all the way back to 2008: the gestation period for this work can be long.

So it was that, even before 802.11ac wave 2 equipment started shipping, the IEEE started work on the next 'PHY' standard, designated '802.11ax'. The project formally kicked-off in March 2014, and as of early 2018 is progressing through a series of 'letter ballots': the scope of the standard is now set, and with each revision the details become increasingly solid. Final approval in the IEEE is expected late in 2019, but the standard will be effectively frozen many months before that.



Source: 650 Group, 2017

Figure 3: Enterprise Access Point Shipments

Earlier physical-layer amendments to 802.11 set a precedent where the Wi-Fi Alliance started its work in parallel with the IEEE, accelerating time-to-market, and 802.11ax follows this timeline: work on a 'Wi-Fi CERTIFIED AX™' certification program is already underway; the first plugfest was in early 2018 and the certification is expected to launch some time in 2019.

This overlap of the certification path with the standardization effort is important to shrink time-to-market, and as the standards organizations and equipment vendors have experience with prior physical-layer amendments, the risks are understood and can be minimized.

Design Goals of 802.11ax

When deciding how to improve Wi-Fi beyond the current release, 802.11ac, the IEEE and Wi-Fi Alliance surveyed Wi-Fi deployments and usage, to identify impediments to wider use and causes of dissatisfaction among user communities.

The conclusion was to depart from previous upgrade paths, which advanced peak data-rates under 'good' field conditions, and to focus more on 'actual' field conditions, and how to improve not just peak performance, but average and worst-case performance in real-world conditions.

These real-world conditions have changed over the years, due in no small part to the success of Wi-Fi. Access points are everywhere, even covering many outdoor spaces. In many areas, congestion has become a serious problem.

Examples include busy airports and train stations, multi-dwelling apartment buildings and even school and university settings. All are characterized by overlapping

coverage from many access points, whether managed in the same network or uncoordinated, all serving many data-hungry client devices. So the IEEE and Wi-Fi Alliance set out to improve performance for everyone, especially in areas of overlapping coverage: in some places, interfering signals can be reduced by coordinating between access points, while in others, protocol enhancements make the Wi-Fi signal more resistant to interference.

But Internet service for cellphones and PCs is not the only use for Wi-Fi. The growing market for Internet-of-Things (IoT) sensors is using Wi-Fi for connectivity in many places, but a few limitations have restricted its adoption. So new features in 802.11ax allow efficient allocation of low data-rate connections, improve the battery life of IoT sensors, and extend the range of Wi-Fi signals.

Wi-Fi is also used by wireless Internet service providers (WISPs) and for outdoor point-to-point links, and here 802.11ax includes features to extend range, increase data-rates and reduce the effects of interference.

Timelines

The procedures for developing 802.11ax broadly followed prior practice for 'PHY' protocols like 802.11n and 802.11ac. This entailed a certain amount of parallel development, with the Wi-Fi Alliance starting work on certification tests before the IEEE has completely finished the underlying specification. Commercial pressures, as for prior physical-layer protocols, will push access point and device vendors to release 'pre-standard' equipment ahead of the Wi-Fi Alliance certification: expect commercially-available equipment in mid- to late-2018, ahead of the certification launch in mid-2019.

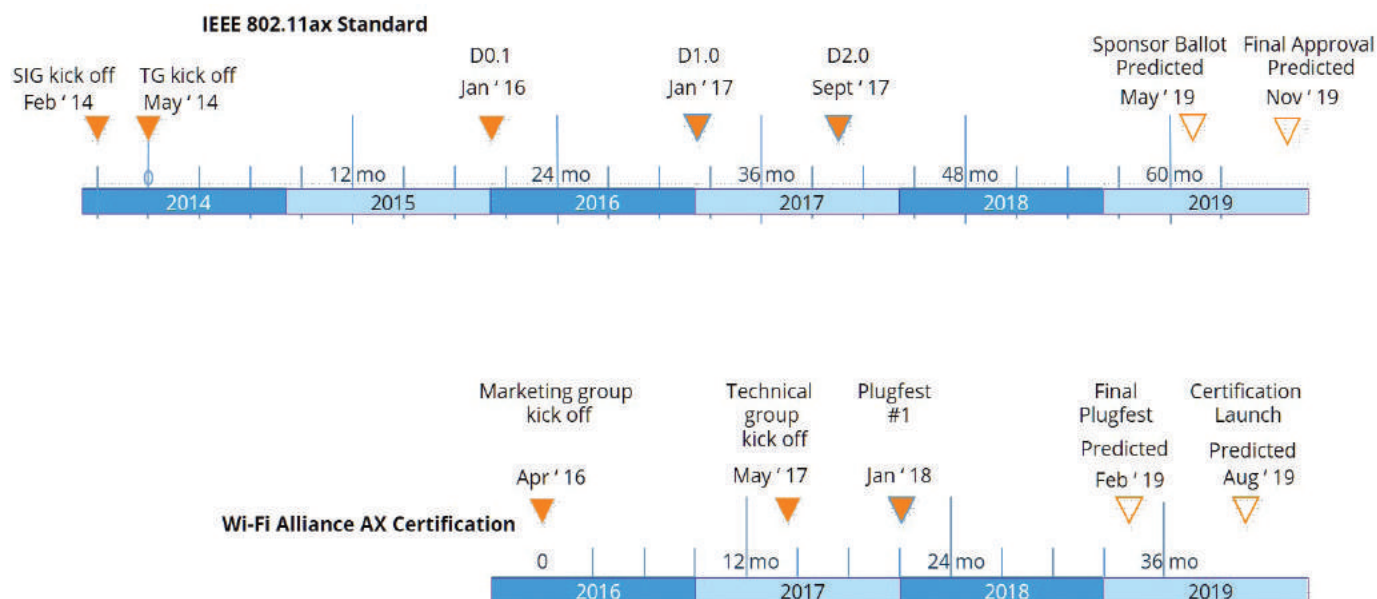


Figure 4: Enterprise Access Point Shipments

While this schedule-compression is not an ideal way to roll out a complex new protocol, the precedents of 802.11n and 802.11ac indicate that the risk of early 802.11ax equipment becoming orphaned is very low. Vendors have successfully met these challenges before.

802.11ax is already split into two 'waves' as was the case for 802.11ac. The precise distribution of features is not yet final, but this paper will focus on those we expect to see in wave 1, dealing with wave 2 features in an appendix. There will be perhaps two years between wave 1 (timeline shown above) and wave 2 availability.

Timing upgrades

New physical-layer Wi-Fi standards require new hardware, hence Aruba is often asked by our customers 'when is the right time to upgrade?' The answer, which has been unchanged from 802.11n to 802.11ac wave 1 and 802.11ac wave 2, is to upgrade when you are ready. There will always be a new, better Wi-Fi access point on the horizon, about every two years, as standards advance, silicon becomes more powerful and equipment vendors add features. After reading a paper like this, some customers are able to identify a particular feature they will need, and will decide to wait for early 802.11ax shipments so they can take advantage – but most will be driven by budgetary and construction deadlines, and other events, and our advice is to use the best technology available when it's time to make the purchase decision.

Spectrum and Regulations

While there are many initiatives and lobbying efforts aiming to open up new spectrum and ease constraints on the broad adoption of unlicensed wireless technologies like Wi-Fi, few changes have been completed in time to affect 802.11ax. Regulations for the 2.4 GHz and 5 GHz bands have not significantly changed since 802.11ac.

One new development is that, while 802.11ac was specified for operation only in the 5 GHz band (802.11n protocols apply in 2.4 GHz), 802.11ax applies to both bands: it can replace all Wi-Fi in use today. Although many consider the 2.4 GHz band so oversubscribed in heavily populated areas as to be unusable, the Wi-Fi community feels there are still many opportunities for this band particularly for IoT where its superior propagation characteristics can be exploited.

If, as is anticipated, regulatory changes allow new spectrum to be apportioned for unlicensed or lightly-licensed use, which is suitable for Wi-Fi, the IEEE and Wi-Fi Alliance will be able to extend the 802.11ax specifications for operation in these new bands.

And, while discussing spectrum, it is important not to neglect the third Wi-Fi band. The 'WiGig' protocol, using the 60 GHz band, is now adopted by both the IEEE and Wi-Fi Alliance, and has significant commonality with Wi-Fi in 2.4 and 5 GHz, allowing connections to be seamlessly switched between different bands. However, due to different characteristics at millimeter-wave frequencies, WiGig has a different physical-layer specification and is not part of 802.11ax.

Convergence with LTE and 5G

This is a time of considerable upheaval in the broad communications industry. On one side, cellular operators (mobile operators) are in the midst of their 4G buildout, but are already preparing for 5G. They see opportunities not just in the traditional Internet-to-cellphone market, but many others such as IoT, smart cities, fixed wireless broadband access to residences, and managed services for enterprise customers. This, and the availability of new licensed spectrum and new technologies, has made 5G considerably broader in scope than the preceding 2G, 3G, 4G generations. In order to accommodate these new use-cases and markets, the 5G standards organizations have expanded their scope.

Meanwhile the radio technologies proposed for 5G and 802.11ax share many characteristics: multi-user MIMO, spatial diversity, beamforming, OFDMA, channel aggregation and others. This is partly because the drivers for spectral efficiency, high data-rates, long-range and good battery life are common to cellular, private and consumer networks, and also because these are state-of-the-art radio techniques that one would expect any new radio standard to incorporate.

But another aspect of 5G is perhaps more interesting. In their quest to cover new markets, especially enterprise networking, the 5G standards now include detailed specifications for integrating Wi-Fi and associated authentication protocols into 5G networks. This, along with lightly-licensed spectrum suitable for either LTE/5G or Wi-Fi use, offers a much broader swath of possibilities for the networks of the future.

So both at the radio and the system level, 5G and Wi-Fi are moving closer together, driven by increasing overlap between the markets targeted by the 3GPP and Wi-Fi Alliance.

Related Wi-Fi standards and certifications

As a new PHY with some MAC modifications, 802.11ax has a few pre-requisites. The Wi-Fi Alliance will require all Wi-Fi CERTIFIED AX™ equipment to be Wi-Fi CERTIFIED AC™ and also Wi-Fi CERTIFIED N™. It will also require the Wi-Fi CERTIFIED Agile Multiband™ certification, a group of features allowing clients broader visibility into network loading, and the ability to move (or be moved) to the optimum band and access point. Most access points and client devices already support the functions required for Agile Multiband, although it is a relatively recent program and not all contemporary equipment has the certification.

And all Wi-Fi equipment will have to meet new security standards for authentication, authorization and encryption. The long-standing WPA2 certification will be replaced during 2018 with WPA3, and it is anticipated that all 802.11ax equipment should also be WPA3-compliant, to support best security practices.

New features in 802.11ax

802.11ax major features

1. Downlink and uplink OFDMA
2. Downlink* and uplink multi-user MIMO
3. Higher order modulation
4. Advanced OFDM and coding
5. Outdoor operation
6. Reduced power consumption
7. Spatial re-use
8. Transmit beamforming*
9. Single-user operation*

(* not new in 802.11ax)

Figure 5: 802.11ax major features (both waves)

There are in excess of 50 features in the IEEE 802.11ax standard: not all will be adopted by the Wi-Fi Alliance. The following is a high-level summary of features (including both wave 1 and wave 2 features).

- **Downlink and Uplink OFDMA:** OFDMA is one of the more complex features in 802.11ax. It allows a single transmission (for downlink OFDMA, the access point transmits) to be split by frequency within a channel, such that different frames addressed to different client devices use groups of subcarriers. Uplink OFDMA is equivalent to downlink OFDMA, but in this case multiple client devices transmit simultaneously, on different groups of subcarriers within the same channel. Uplink OFDMA is more difficult to manage than the downlink variety, as many different clients must be coordinated: the access point transmits trigger frames to indicate which sub-channels each client can use.
- **Downlink and uplink multi-user MIMO:** The downlink version extends an existing 802.11ac feature where an access point determines that multipath conditions allow it to send, in a single time-interval, frames to different client devices. 802.11ax increases the size of downlink MU-MIMO groups, allowing more efficient operation.

Uplink multi-user MIMO is a new addition to 802.11ax, but is deferred to wave 2: like uplink OFDMA, the access point must coordinate the simultaneous transmissions of multiple clients.

- **Transmit Beamforming:** This is another existing feature where an access point uses a number of transmit antennas to land a local maximum signal on a receiver's antennas. It improves data-rates and extends range.
- **Higher-Order Modulation:** 802.11a/g introduced 64-QAM, and 802.11ac 256 QAM: in 802.11ax, the highest-order modulation is extended to 1024-QAM. This increases peak data-rates under good conditions (high SNR).
- OFDM symbols, subcarrier spacing and FFT size are all changed to allow efficient operation of small OFDMA sub-channels: these changes allow an increase in the length of guard interval without loss of symbol efficiency.
- **Outdoor Operation:** A number of features improve outdoor performance. The most important is a new packet format where the most sensitive field is now repeated for robustness. Other features that contribute to better outdoor operation include longer guard intervals and modes that introduce redundancy to allow for error recovery.
- **Reduced Power Consumption:** Existing power-save modes are supplemented with new mechanisms allowing longer sleep intervals and scheduled wake times. Also, for IoT devices, a 20MHz-channel-only mode is introduced, allowing for simpler, less powerful chips that support only that mode.

- **Spatial re-use:** When contending for a transmit opportunity, a device is allowed to transmit over the top of a distant transmission, which would previously have forced it to wait. This increases network capacity by allowing more simultaneous transmissions in a given geographic area.

In historical context, it can be seen that the new features in 802.11ax are mostly extensions or improvements on previous work – with the standout exceptions of OFDMA and spatial re-use, which are new territory.

TECHNICAL FEATURES OF 802.11AX

The following section deals with the major technical improvements in detail.

New subcarrier spacing and symbol duration

The OFDM symbol is the basic building-block of a Wi-Fi transmission. It is a small segment in time of the modulated waveform of a subcarrier, carrying information: the more variants of a symbol are available, the more information (binary bits) it can carry. The fundamental characteristics: fast Fourier transform (FFT) size, subcarrier spacing and OFDM symbol duration are linked, given a fixed channel width. In 802.11ax, the subcarrier spacing is reduced by a factor of 4x while the OFDM symbol duration increases by 4x.

802.11n (2008):	802.11ac (2012):	802.11ax (2018):	Goals of the 802.11ax project:
<ul style="list-style-type: none"> • 2.4 and 5 GHz supported • Wider channels (40 MHz) • Better modulation (64-QAM) • Additional streams (up to 4) • Beam forming (explicit and implicit) • Backwards compatibility with 11a/b/g 	<ul style="list-style-type: none"> • 5 GHz only • Even wider channels (80, 160 MHz) • Better modulation (256-QAM) • Additional streams (up to 8) • Beam forming (explicit) • MU-MIMO • Backwards compatibility with 11a/b/g/n 	<ul style="list-style-type: none"> • 2.4 GHz and 5 GHz supported • OFDMA uplink and downlink • Extends and generalizes OFDM • Introduces the concept of Resource Units (RU's) • Massive parallelism • Better modulation (1024-QAM) • Uplink MU MIMO • Spatial re-use (BSS color) • Backwards compatibility with 11a/b/g/n/ac 	<ul style="list-style-type: none"> • Enhance operation in 2.4 & 5 GHz bands (802.11ac was only 5 GHz) • Increase average throughput per station by at least 4x in a dense deployment scenario (802.11ac specified aggregate throughput without a specific scenario) • For outdoor and indoor networks • Scenarios include wireless corporate office, outdoor hotspot, dense residential apartments, stadiums • Maintain or improve power efficiency of client devices

Figure 6: Wi-Fi Standards Progression

Table 1: OFDM characteristics from 802.11ac to 802.11ax

	802.11ac	802.11ax
Bands	5 GHz only	2.4 GHz and 5 GHz
Channels	20, 40, 80, 80+80, 160 MHz	20, 40, 80, 80+80, 160 MHz
FFT Sizes	64, 128, 256, 512	256, 512, 1024, 2048
Subcarrier spacing	312.5 kHz	78.125 kHz
OFDM symbols	3.2 usec	12.8 usec
OFDM symbol cyclic prefix	0.8 or 0.4 usec	0.8 or 1.6 or 3.2 usec
Highest modulation	256 QAM	1024 QAM
Spatial streams	1-8 (not implemented beyond 4)	1-8 (may be implemented)


802.11ax

Figure 7: OFDM symbol duration & subcarriers

The primary impetus for a change in subcarrier spacing was to allow OFDMA (see later) to extend to small sub-channels. Each sub-channel requires at least one (usually two) pilot subcarriers, and with a 2 MHz minimum sub-channel size, a smaller subcarrier spacing loses a much smaller percentage of the overall bandwidth to pilots.

There are other advantages. The number of guard and null subcarriers across a channel can be reduced as a percentage of the number of usable subcarriers, again increasing the effective data rate in a given channel. The figures above show a ~10% increase in usable subcarriers compared to 802.11ac, after allowing for the 4x factor.

The longer OFDM symbol allows for an increase in the cyclic prefix length without sacrificing spectral efficiency, which in turn enables increased immunity to long delay spreads, especially in outdoor conditions. The cyclic prefix can be reduced to a smaller percentage of the symbol time, increasing spectral efficiency even while more robust to multipath conditions. And it reduces the jitter-sensitivity of uplink multi-user modes.

There are, of course, some side effects. The frequency accuracy required to successfully demodulate more closely-spaced subcarriers is more stringent. Also, the Fast Fourier Transform (FFT) requires a slightly more complex chip. But since it has been around 20 years since the 312.5 kHz / 64 point FFT was first used in 802.11, these effects are considered manageable.

OFDMA advantages

Orthogonal frequency division multiple-access (OFDMA) is one of two multi-user modes in 802.11ax, the other being MU-MIMO (downlink-only in wave 1). OFDMA is a technique that has been used in other systems, like cellular-LTE, for many years. It works by dividing a transmission across the frequency dimension, with pairs of devices assigned to transmit and receive in sub-channels or Resource Units (RU's) of the main RF channel.

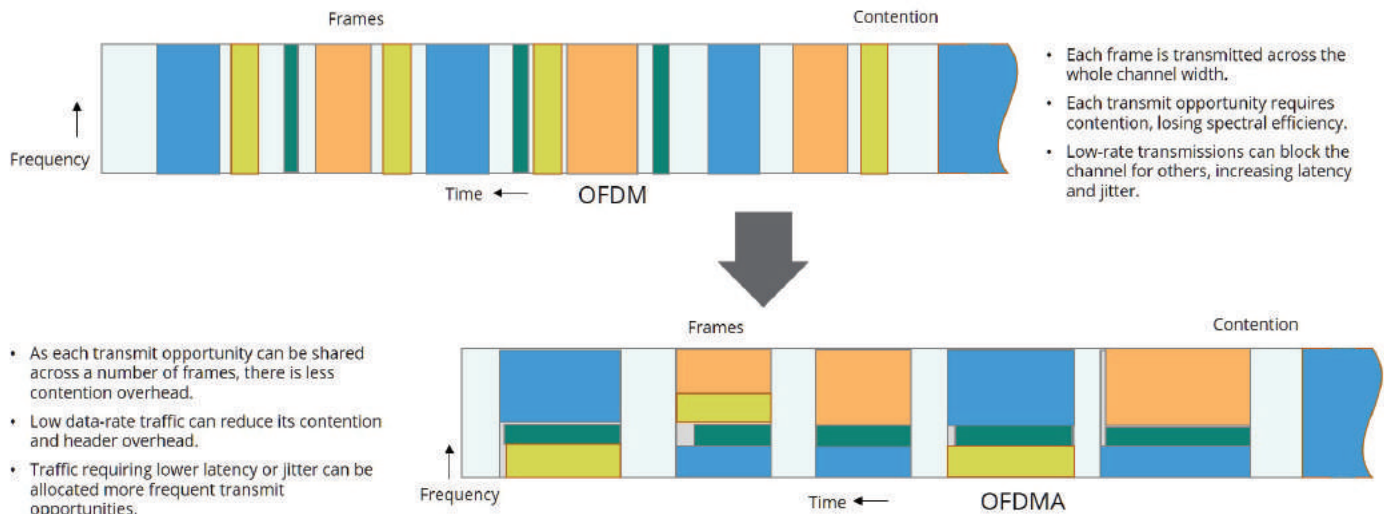


Figure 8: OFDMA compared with single-user OFDM

This allows an access point (for downlink OFDMA) to bundle several frames together in different sub-channels in a single transmit opportunity, while its clients tune their radios to different sub-channels to receive their respective transmissions.

At first glance, OFDMA offers no advantage over full-channel single-user OFDM (which is still available in 802.11ax). Given that the transmit link-speeds do not differ, when considering a long time-period covering many transmissions, each station transmits the same amount of data: when OFDMA allocates 1/2 of the channel, the transmission takes 2x the time and nothing is saved. However, a closer examination shows a number of efficiency improvements.

In the 802.11 CSMA/CA channel access protocol, each transmit opportunity negotiation loses time to contention. That is time lost on the medium, reducing overall capacity and spectral efficiency. When OFDMA is used, transmissions are bundled together, reducing the number of transmit opportunities necessary to move a given amount of data, and increasing efficiency. Also, CSMA/CA becomes less efficient as the number of clients increases—if 5 clients can achieve 100 Mbps each, 50 would not be able to achieve 10 Mbps—and one goal of 802.11ax is to improve performance in large-scale deployments with dense client populations. OFDMA is especially useful in managing large numbers of clients fairly, and the reduction in contention overhead means there is little deterioration in capacity as client numbers increase.

There are also advantages for less-capable stations. As link-speeds have increased, some devices struggle to transmit at the maximum rates. Whereas with full-channel OFDM, they have to do the best they can, perhaps not filling the medium, OFDMA allows them to cap their maximum rates. This allows for simpler hardware implementations and potentially longer battery life, similar to the 20 MHz-only (see later) concept for IoT sensors.

OFDMA also offers opportunities for applying QoS, particularly for traffic that demands low latency or jitter. While a device might have to wait a long while in a single-user OFDM system for a transmit opportunity, OFDMA allows it to transmit 'little and often', reducing latency and jitter.

But OFDMA brings some subtlety. As can be seen above, the frame-by-frame transmission of 802.11 requires that, when the access point contends for a transmit opportunity, it must bundle up a number of frames of different lengths. Where frames are shorter than the length of the transmit opportunity, padding is added and this, of course, is usable bandwidth lost. Also, as hinted earlier, each OFDMA sub-channel must reserve one or two subcarriers for pilot tones, unusable for data transmission. So the access point must calculate the optimum use of OFDMA, taking into account its offered load and the frames in its buffers, as well as client distribution and link-speeds.

OFDMA thus opens up many new dimensions to traffic management, but it also requires more sophisticated control mechanisms, as the access point must choose how to allocate sub-channels, and coordinate with its clients on use; more on these later.

Downlink OFDMA

802.11ax introduces OFDMA in both the downlink and uplink directions. There are some differences: since the access point is in control of all transmissions for downlink OFDMA, it may be simpler to implement.

Dealing first with downlink OFDMA transmission, and leaving control till later, we see that an access point first contends for a transmission opportunity in the usual way. It then assembles a number of frames for different clients, but modulated over the allocated sub-channels.

When a frame is shorter than the longest frame of the bundle, padding is added to bring the length up. This lost bandwidth can be reduced by allocating a smaller

sub-channel to the frame, so it takes longer to transmit, but then more frames must be added to the bundle to use the full channel bandwidth.

The smallest allocated sub-channel in 802.11ax is 26 subcarriers (2 MHz). There are 9 available 26-subcarrier sub-channels in a 20 MHz channel, allowing up to 9 different frames and recipients to share a transmission.

The IEEE uses the term “Resource Unit” (RU) to refer to sub-channels. The 26-subcarrier unit above is known as RU-26, for example: the full set is RU-26, RU-52, RU-106, RU-242, RU-484, and RU-996.

Allocated sub-channels

In OFDMA, uplink or downlink, sub-channels are defined in the standard. Channels are sub-divided in binary fashion, with the 26-subcarrier smallest block being used to fill holes where the channel does not divide exactly.

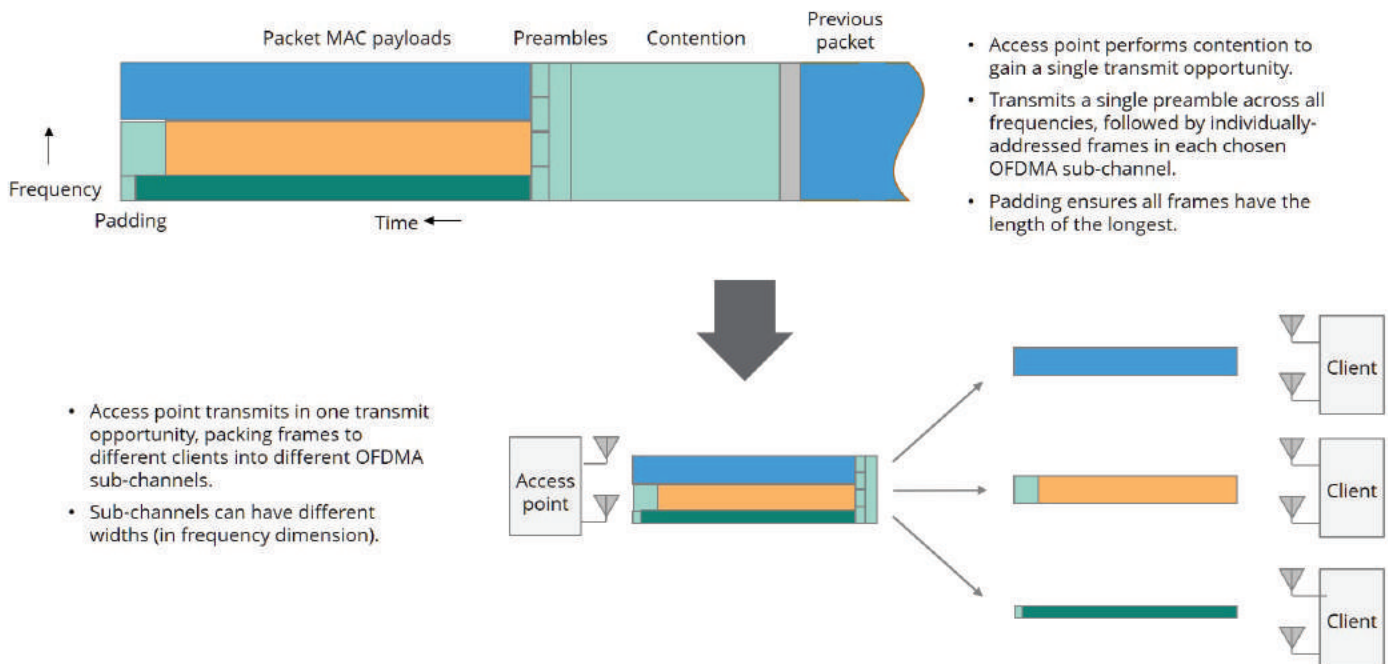


Figure 9: Downlink OFDMA transmission

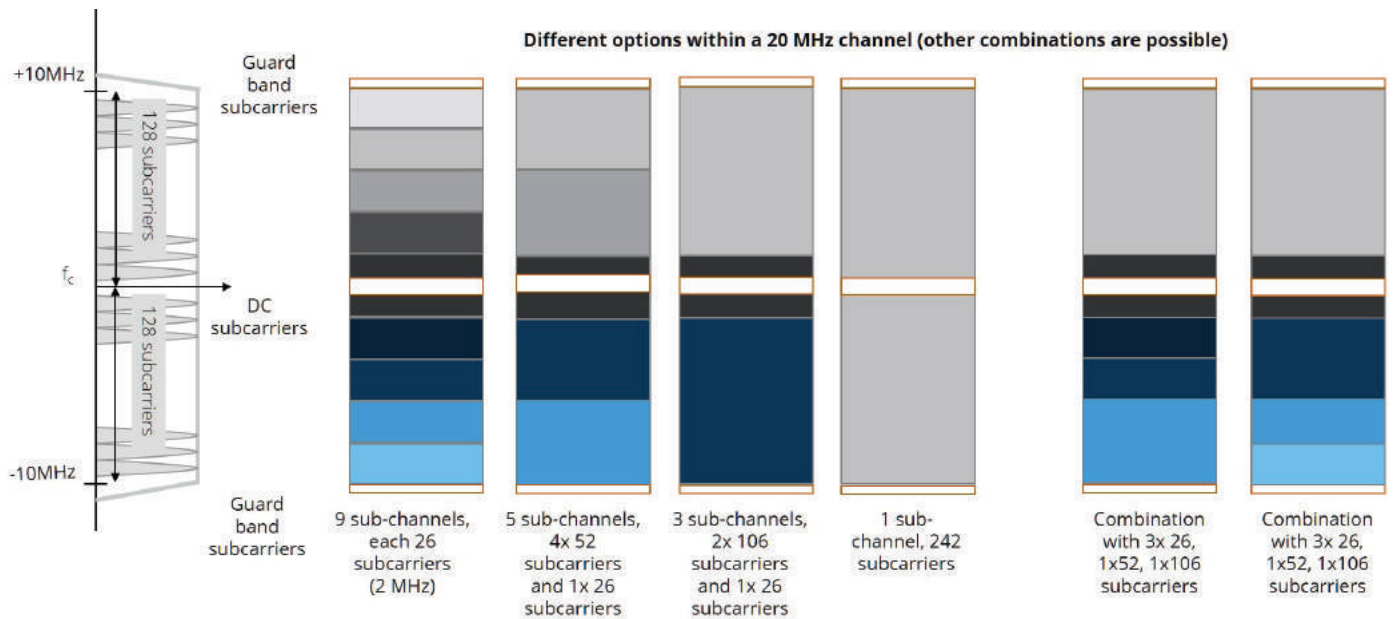


Figure 10: Sub-channel allocation for OFDMA

As in prior generations of OFDM, not all subcarriers in a channel can be used for data. Some subcarriers are unused for guard-band purposes, so as not to interfere with transmissions in adjacent channels, or between

sub-channels. Others are used for DC or pilot tones, to provide a frequency reference and allow accurate demodulation of the signals.




Figure 11: Sub-channel allocation for OFDMA

Usable sub-channels, subcarriers and data-rates for OFDMA

The table below lists the menu of options for RU-N (e.g. RU-26) sub-channels in OFDMA. These RU's can move

around but only in certain configurations as specified in 802.11ax.

Table 2: Usable sub-channels, subcarriers and data-rates for OFDMA

Channel (MHz)	Divided into sub-channels	Bandwidth of largest sub-channel	Data subcarriers including pilots	Data subcarriers excluding pilots	Lower-end guard subcarriers	Higher-end guard subcarriers	DC subcarriers	Pilot subcarriers	Null subcarriers	Total subcarriers
20	9x 26	1.88 MHz	234	216	6	5	7	18	4	256
	4x 52, 1x 26	3.91 MHz	234	216	6	5	7	18	4	256
	2x 106, 1x 26	7.97 MHz	238	228	6	5	7	10	0	256
	1x 242	18.3 MHz	242	234	6	5	3	8	0	256
40	18x 26	1.88 MHz	468	432	12	11	5	36	16	512
	8x 52, 2x 26	3.91 MHz	468	432	12	11	5	36	16	512
	4x 106, 2x 26	7.97 MHz	476	456	12	11	5	20	8	512
	2x 242	18.3 MHz	484	468	12	11	5	16	0	512
	1x 484	36.6 MHz	484	468	12	11	5	16	0	512
80	37x 26	1.88 MHz	962	888	12	11	7	74	32	1024
	16x 52, 5x 26	3.91 MHz	962	888	12	11	7	74	32	1024
	8x 106, 5x 26	7.97 MHz	978	936	12	11	7	42	16	1024
	4x 242, 1x 26	18.3 MHz	994	960	12	11	7	34	0	1024
	2x 484, 1x 26	36.6 MHz	994	960	12	11	7	34	0	1024
	1x 996	76.6 MHz	996	980	12	11	5	16	0	1024

Uplink OFDMA

OFDMA works in the uplink direction much as for the downlink, except that the client devices transmit and the access point receives.

The difficult functions are for the access point to calculate the best grouping of clients, and then to signal when each should transmit, and on which sub-channel. More on that later.

Also, synchronization of preamble symbols in the uplink direction is complex because each preamble is transmitted across a full 20 MHz channel. This was an implementation decision in the IEEE, and it requires that all preamble

waveforms are synchronized in time, frequency and amplitude when received at the AP's antennas. This has driven a number of new requirements for Wi-Fi devices including calibrating signal strength measurements, local oscillator requirements and others, which may be useful in other areas.

Even within an OFDMA packet body, it is very important that the transmitter maintains frequency accuracy, transmitter linearity and other parameters to avoid causing interference to transmissions in adjacent RU's: implementation of OFDMA is more complex than the simple diagrams above would indicate.




Figure 12: Uplink OFDMA transmission

Downlink multi-user MIMO transmission

Downlink MU-MIMO was introduced in 802.11ac (wave 2) and is becoming widespread in current access points and client devices. It extends concepts of spatial diversity and beamforming to support simultaneous transmission from an AP to a number of clients.





Figure 13: Downlink multi-user MIMO transmission

MU-MIMO is only possible where propagation characteristics allow the AP to identify that a transmission optimized for one client or group of clients will not be heard at a significant signal strength by another client, and vice versa. These are the conditions that allow it to build separate data frames for each client group, and transmit them simultaneously.




Figure 14: Downlink Sounding

In order to identify candidates for MU-MIMO, the AP performs sounding operations, sending null frames from all its antennas to clients, which then return responses with matrices of the measured receive levels for each AP-antenna to client-antenna pair. Sounding is used for beamforming as well as MIMO. Multi-user sounding in 802.11ac could be time consuming because the beamforming report matrix can be large, and the client devices had to stagger their responses to avoid interference: the new 802.11ax multi-user control protocol makes it much more efficient with simultaneous responses.




Figure 15: Sounding—beamforming explicit feedback

Experience with 802.11ac MU-MIMO in real-world deployments revealed some limitations. For instance, it was not always possible to form usable groups, and even with a 4-antenna AP, gains over single-user mode were sometimes modest: in 802.11ax, the larger MU-MIMO groups (increased from 4 to 8 clients) will allow considerable improvement. As can be seen from the diagram above, 802.11ax can accommodate large numbers of client devices by grouping clients and dealing with groups sequentially. The example shows grouping of clients for beamforming reports, but the concept is also extended to other packet types.

And any link- or transport-level protocol like TCP/IP that includes acks will gain from the improved downlink performance but may still be bottlenecked by the uplink: this will be solved when uplink MU-MIMO is added in 802.11ax wave 2.

Uplink MU-MIMO and combined MU-MIMO and OFDMA

Both uplink MU-MIMO, and packets that combine OFDMA and MU-MIMO are deferred to wave 2. Even downlink OFDMA and MU-MIMO together are not supported in wave 1, although the combination opens even more possibilities for the AP scheduler.

Packet preambles

In the 802.11 protocol, packet preambles contain information for the receiver to synchronize to the incoming signal, and identify the sub-channels and format of the packet to follow. The information below is encoded by the PHY layer of the transmitter.

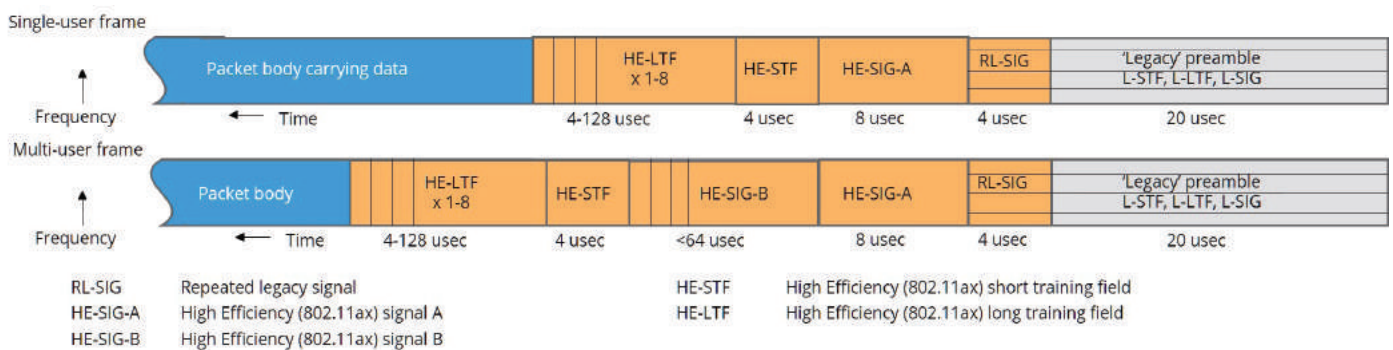


Figure 16: Preambles and training sequences

The 802.11ax frame starts with the 'legacy' preamble for backwards-compatibility: these fields have been used since before 802.11n and allow older devices to recognize there is an 802.11 frame on the air. This allows the CSMA/CA protocol to continue functioning in the presence of 802.11ax transmissions.

The next field, RL-SIG, would be the beginning of the frame body in older protocols like 802.11g. It identifies the frame to follow as 802.11ax rather than pre-802.11n. The 'legacy' preamble and RL-SIG field are transmitted in parallel in all 20-MHz sub-channels used for subsequent transmissions, for backwards-compatibility.

The subsequent fields are used for 802.11ax purposes ('HE' is 'High Efficiency', the IEEE 802.11 name for 802.11ax) and use a mix of symbol formats, with 'legacy' modulation used for low-rate fields and for backwards compatibility, while other fields use the new, close subcarrier spacing and longer OFDMA symbol of 802.11ax.

First is the HE-SIG-A field, which contains information about the packet to follow, including whether it is downlink or uplink, BSS color, modulation MCS rate, bandwidth and spatial stream information, and remaining time in the transmit opportunity. This field has different content for single-user, multi-user and trigger-based frames, and is repeated in the 'extended range mode' of 802.11ax.

The HE-SIG-B field is only included for multi-user packets. It has information common to all recipients, and other fields that are user-specific, so its length depends on the number of users receiving the transmission. When OFDMA is used, the HE-SIG-B client-specific fields are sent concurrently in each sub-channel used for the subsequent packet transmission. More on this later.

The HE-STF training field allows receivers to synchronize to the timing and frequency of the incoming frame before decoding the packet body, while the HE-LTF is important for channel estimation, enabling beamforming and MIMO spatial diversity.

Packet tail—padding, tail bits and packet extension

The new structures and applications of 802.11ax mean some new fields are added to the end of the packet.




Figure 17: Padding, tail bits and packet extension

Padding may be added after the packet payload. It is required when OFDMA is used and the frame, as built by the transmitter, is not quite long enough to fill the negotiated transmit opportunity. The calculations to determine optimal bandwidth utilization are performed by the AP, and it varies the sub-channel, MCS rate and transmit power for the frames grouped in a transmission to ensure that all transmissions start and end simultaneously. This is important because the other devices on the channel, including pre-802.11ax devices, must see signals at a certain power level filling the channel in order for their CSMA/CA contention mechanisms to work correctly. Padding can be included in the forward error correction (FEC) calculation, or added after the calculation.

If the AP is doing a good job (assuming the system is operating at capacity), very little padding will be used. If it has a 'short' frame, it can always reduce the MCS rate to improve the transmission's error rate, driving a longer duration.

Tail bits may be added after the data field. They are only necessary when BCC error correction is used, not for LDPC. This field existed prior to 802.11ax. (Binary convolutional codes (BCC) were used in early 802.11 standards for error correction. As data-rates increased, the BCC decoder became complex, and now higher data-rates use low-density parity check (LDPC) coding, a lower-complexity alternative.)

The packet extension field may be added at the end of the frame. It is used to allow extra time for the receiver to process the frame's contents before responding with a frame of its own, recognizing for the first time in 802.11 that some chips may move certain functions to slower software layers rather than fast-calculating hardware. A client requiring extra time to process received frames must signal its requirements to the AP: the allowed values for packet extension are 0, 4, 8, 12 or 16 usec.

Packet aggregation

Packet aggregation was introduced in 802.11n and has become widespread, particularly where streaming video is carried over Wi-Fi.

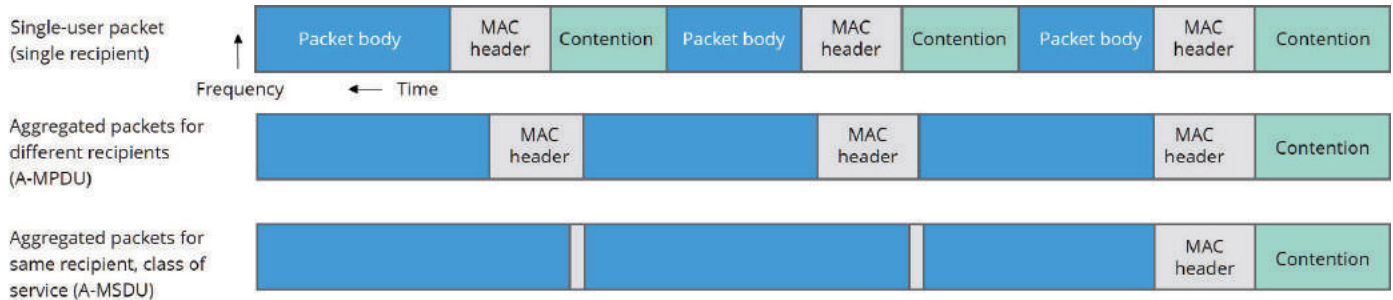


Figure 18: MAC Aggregation

The value of MAC aggregation lies in more efficient use of the air, for higher throughput and more capacity. This stems from two effects.

A-MSDU aggregation requires a full MAC header only on the first packet of the sequence, reducing header overhead. This is a significant effect, but eliminating per-packet contention is bigger: with both A-MSDU and A-MPDU aggregation, the transmitter is able to negotiate a transmit opportunity covering many packets, greatly reducing contention overhead.

Packet aggregation is not changed for 802.11ax, but it still plays a significant part in optimizing network capacity. It works in conjunction with MU-MIMO, and with OFDMA: for all the OFDMA illustrations in this paper, packets within sub-channels will often be aggregated packets.

Control for multi-user modes

802.11ax includes two multi-user modes: MU-MIMO, which exploits diversity in space, and OFDMA in the frequency dimension. Both modes allow simultaneous bi-directional communication between an AP and multiple client devices, and 802.11ax provides common control mechanisms.

Downlink and uplink are different: the former has no prior signaling, the AP just starts to transmit in appropriate modes, and receivers synchronize as the packet arrives. But multi-user uplink traffic requires a special 'trigger' frame where the AP allocates MU-MIMO groups and OFDMA Resource Units to its clients, and informs them of the allocation, and this in turn requires that the AP polls clients for their uplink traffic requirements.

Downlink multi-user control

There is no prior signaling for downlink multi-user control: all relevant information is contained in the packet header, specifically in the HE-SIG-B field, which is only included in downlink multi-user frames.

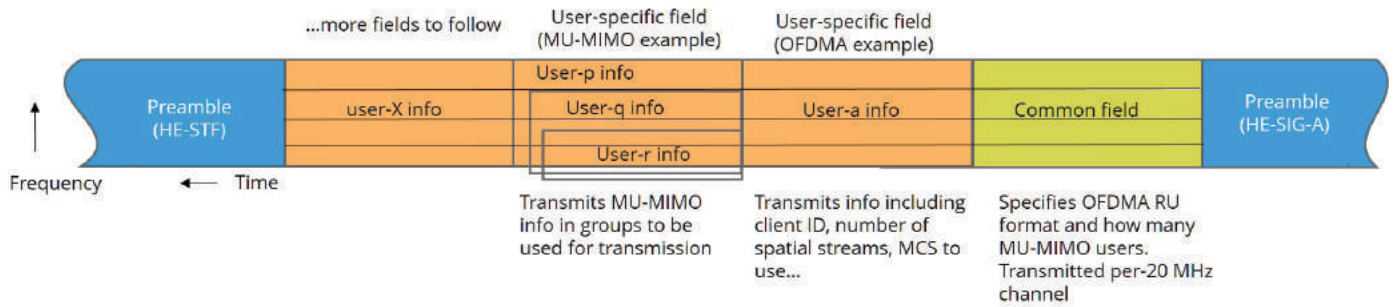


Figure 19: Control for multi-user downlink modes

The HE-SIG-B is a complex field. It has variable length, depending on the number of clients the AP is addressing, and two different types of information, common and user-specific.

The common field identifies the structure of OFDMA sub-channels or RU's that will be used, e.g. 18x 26 RU or 2x 242 RU. It includes other information that is common to all transmissions.

A number of user-specific fields follow the common field. The AP uses these fields to identify exactly how it will be transmitting to each client, including the number of spatial streams, the MCS it will use and whether it will use beamforming.

The 802.11ax specification requires the transmitter to form the HE-SIG-B field simultaneously in multiple 20 MHz channels, taking up the total bandwidth of the allocated channel. Thus, if the AP is using an 80 MHz channel, it will transmit 4 HE-SIG-B fields, one in each 20 MHz subchannel.

The HE-SIG-B fields provide all the information a client device needs to discover it is the intended recipient of the frame to follow, and the information it needs to receive and decode that frame.

In Figure 20, downlink shows how multi-user frames follow a simple format, no trigger or signaling frames are necessary. However, it becomes more difficult to manage acknowledgements, as these are uplink transmissions and, in multi-user mode, require coordination and a trigger frame from the AP.

Options for the trigger frame are BlockAck request (MU-BAR), buffer status report (BSRP), bandwidth query report (BQRP) and uplink multi-user response scheduling control fields in the basic packet preamble.

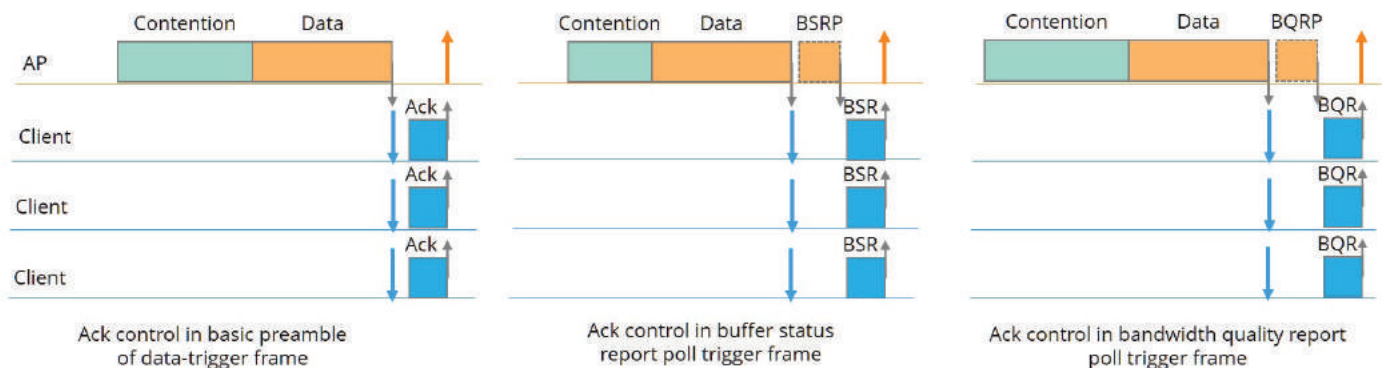


Figure 20: Control for multi-user downlink modes

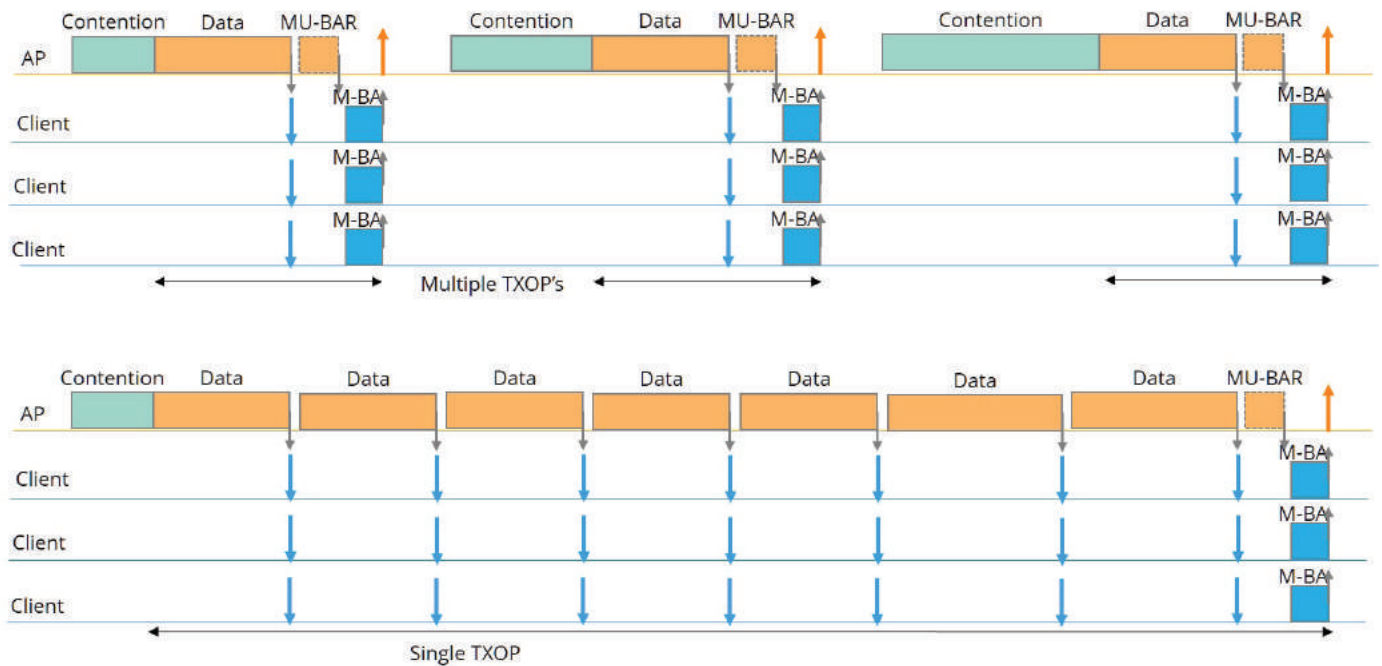


Figure 21: Control for multi-user downlink modes

BlockAck's can be used in conjunction with downlink MU modes, allowing a set of MU BlockAck's to be deferred to the end of a group of downlink data frames, transmitted in the same transmit opportunity up to the TXOP limit of 4.096 msec. This minimizes the overhead from contention and multiple Ack's.

(802.11ac introduced downlink multi-user MIMO, but had no uplink multi-user mode, so recipients of a downlink MU transmission had to ack one after the other, wasting time on the air. The 802.11ax approach is an improvement on 802.11ac.)

Uplink multi-user control

The uplink is more complicated than the downlink, as the AP first has to discover what traffic clients are ready to transmit. Following this, it must calculate the optimum allocation of MU-MIMO groups and OFDMA RU's, then signal the allocation information to its clients and synchronize them to transmit simultaneously.

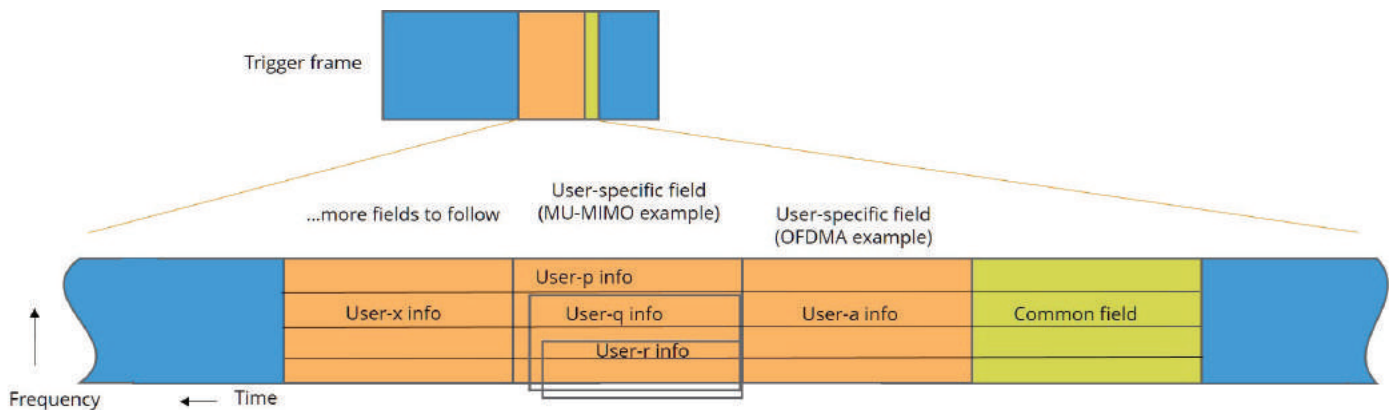


Figure 22: Control for multi-user uplink modes

The trigger frame format is shown above. It contains the following information.

- The length of the uplink transmission window
- Which client devices are to transmit
- Which OFDMA RU's are to be used by each client device
- How many spatial streams are to be used by each client device
- Which MCS modulation level is to be used by each client device
- Whether functions like STBC are to be used on the uplink
- Required signal strength at the AP for the client's transmission (this is calculated by the client using the AP's transmit power level, the client's receive RSSI level and an assumption of channel reciprocity)
- (Uplink MU-MIMO is deferred to wave 2, but all the necessary fields are already defined in the 802.11ax standard.)

This is a very versatile frame because it can be concatenated with several other functions, as listed below.

- **Basic Trigger frame:** This has no extra functions. It specifies how and when client devices should respond.
- **Beamforming Report Poll (BRP):** This solicits beamforming reports from client devices. User Info fields specify how the beamforming report is formatted. There is no Common field in this frame.

- **Multi-user BlockAck Request (MU-BAR):** This trigger frame requests a BlockAck from multiple client devices simultaneously. User Info fields specify the frames that are to be Ack'd.
- **Multi-user Request To Send (MU-RTS):** This trigger frame is used to clear the air before a transmission, in the same way as single-user RTS-CTS.
- **Buffer Status Report Poll (BSRP):** This trigger frame allows the AP to find what traffic client devices have queued to transmit, allowing the AP to schedule uplink traffic efficiently.
- **Bandwidth Query Report Poll (BQRP):** This trigger frame requests client devices to report on the occupancy of 20 MHz RF channels, allowing the AP to control uplink channel use efficiently.
- **Group Cast with Retries multi-user BlockAck Request (GCR MU-BAR):** This intimidating frame is used when the AP is building a multicast group and solicits a BlockAck from each member of the group.

A trigger frame of some kind is required any time an AP wishes to initiate a set of uplink multi-user frames.

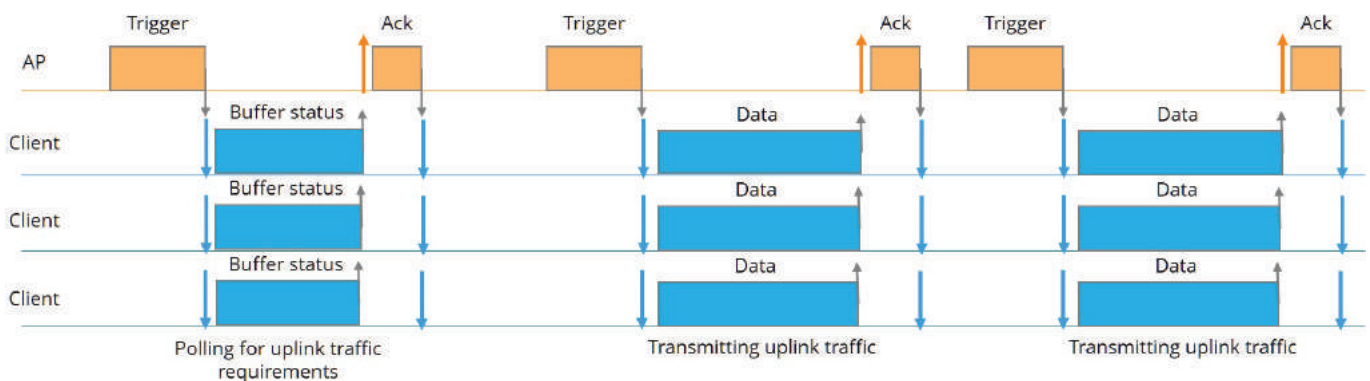


Figure 23: Control for multi-user uplink modes

The trigger frame is used to map clients to their respective OFDMA RU's and MU-MIMO groups, and includes timing and MCS modulation-rate information, as well as transmit-power guidance.

At the designated time, client devices start transmitting in their assigned RU's or MIMO groups.

The AP usually transmits an acknowledgement frame following uplink data frames. This can be either a multi-user transmission individually addressing clients with block-ack (BA) frames, or a new 'multi-STA BlockAck' frame contained in a pre-802.11ax, or an 802.11ax frame.

AP scheduling for 802.11ax multi-user modes

With the addition of OFDMA and uplink MU-MIMO, the AP must perform additional functions not required in previous generations of 802.11. Scheduling of the downlink and uplink traffic becomes critical for optimal performance in a heavily-loaded system.

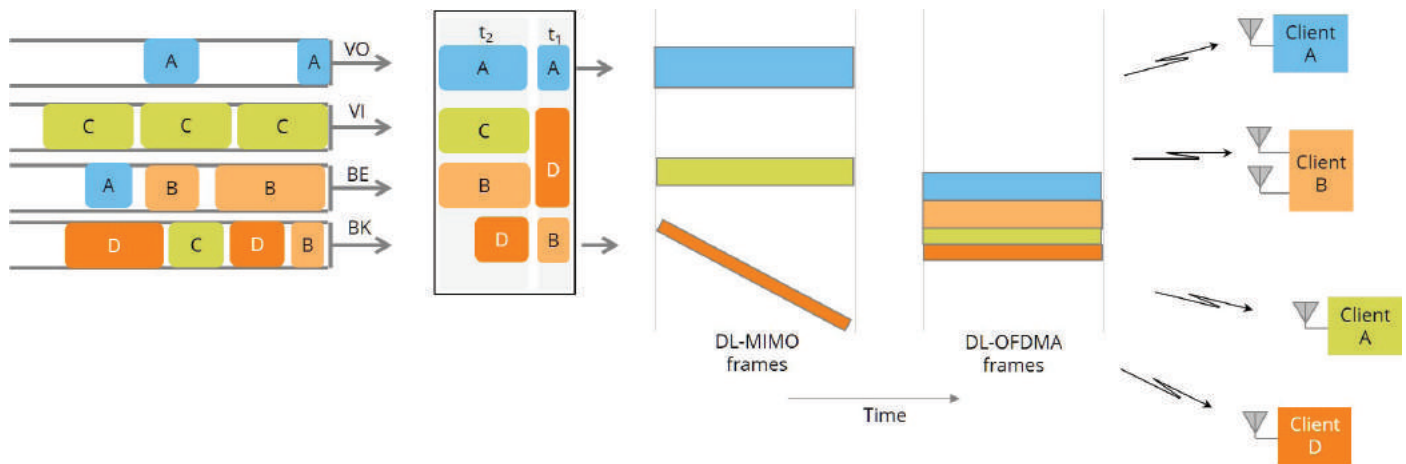


Figure 24: AP scheduling for downlink multi-user modes

When 802.11ac introduced downlink MU-MIMO, the AP took on the task of monitoring its buffers of downlink traffic and deciding how to group the various packets to make best use of the distribution of MU-MIMO groups across its client population. Sometimes it would be optimal to reach into the buffers to fill out a transmission group, for instance. Identifying groups becomes more complex in 802.11ax: taking account of signal strength (near-far distance from the AP) may improve efficiency gains, and in the future, as data from real-world deployments accumulates, there will be opportunities for big-data and machine-learning to analyze performance and improve scheduling algorithms.

With OFDMA, the downlink traffic-grooming problem gains a new dimension: now the AP must look ahead in its buffers, take account of both MU-MIMO groups and OFDMA channels (and which of its clients are 802.11ax-capable) and re-order and group packets.

But the uplink is even more complicated. Even though uplink MU-MIMO will be deferred to 802.11ax wave 2, uplink OFDMA will be a significant feature of the initial wave of 802.11ax equipment. Multi-user operation requires the AP to learn of its clients' buffer states and traffic streams, then make equivalent calculations to the uplink, then signal and coordinate uplink multi-user transmissions for optimum system performance.

This function is already used in cellular systems, where base stations incorporate considerable expertise and intellectual property in their scheduler algorithms. We should expect to see similar developments in 802.11ax APs.

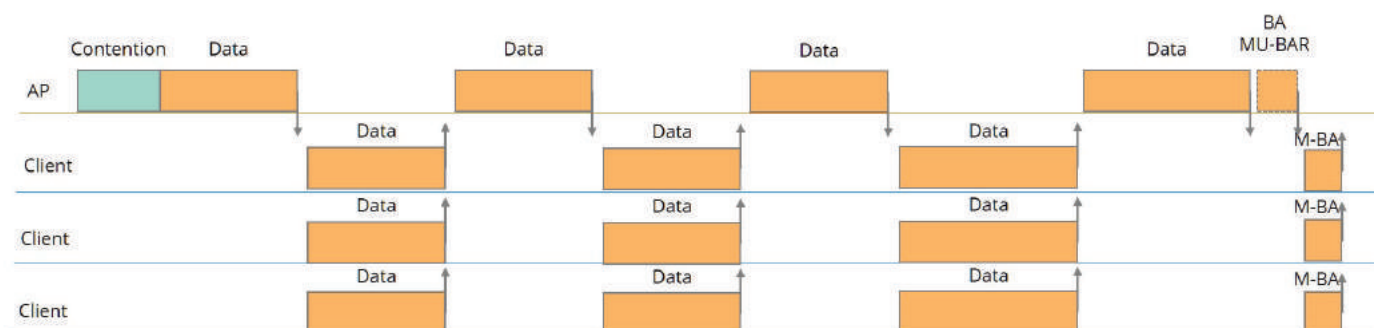


Figure 25: Downlink-Uplink multi-user cascading frame exchanges

The most efficient multi-user scheme available to the AP cascades uplink and downlink sets of multi-user frames, multiplexed in space with MU-MIMO (downlink only in wave 1) and OFDMA as shown above. The downlink packets include Ack's and triggers, and the uplink are trigger-based frames which also carry Ack's, and all are controlled and orchestrated by the AP.

It is interesting that the new multi-user modes in 802.11ax, along with efficient scheduling by the AP, allow a Wi-Fi system to act in a near-cellular (TDD + TDM/TDMA + OFDMA) fashion. The AP can schedule consecutive multi-user transmit opportunities for uplink and downlink and, with appropriate traffic, the per-packet overhead associated with prior 802.11 protocols becomes so small as to nearly disappear. Meanwhile the scaling of client numbers and granularity of OFDMA bandwidth assignment allows a very broad range of client-density and traffic scenarios to be accommodated.

8-antenna access points and client devices

The 802.11ac standard extended to 8 the maximum specified number of antennas an access point or client could use. In the event, while there are many 4-antenna 802.11ac access points, no equipment vendors have ventured beyond 4.

The 802.11ax standard keeps the upper limit of 8 antennas, and it is quite likely that, in the same way that 5G is embracing 'massive MIMO', there will be opportunities to build innovative 802.11ax products with up to 8 antennas. Benefits beyond MU-MIMO include beamforming and MRC with more antennas, and more efficient spatial grouping of 1- and 2-antenna client devices by the AP.

High-density AP and client situations, performance-sensitive applications and point-to-point links are some of the possible scenarios. There will still be high-volume production access points for consumer and enterprise use with 4 antennas or fewer, but these will now be mid-range products.

It seems less likely that clients will increase their antenna count. Many smartphones and tablets support 2 spatial streams, and this seems adequate for their performance needs: the gains from extra antennas will be in the overall capacity an AP can support: more clients at higher data-rates than before.

High-order modulation

It is now traditional for a new 802.11 physical-layer amendment to bump up the highest modulation level, and 802.11ax adds two 1024-QAM rates on top of 802.11ac.

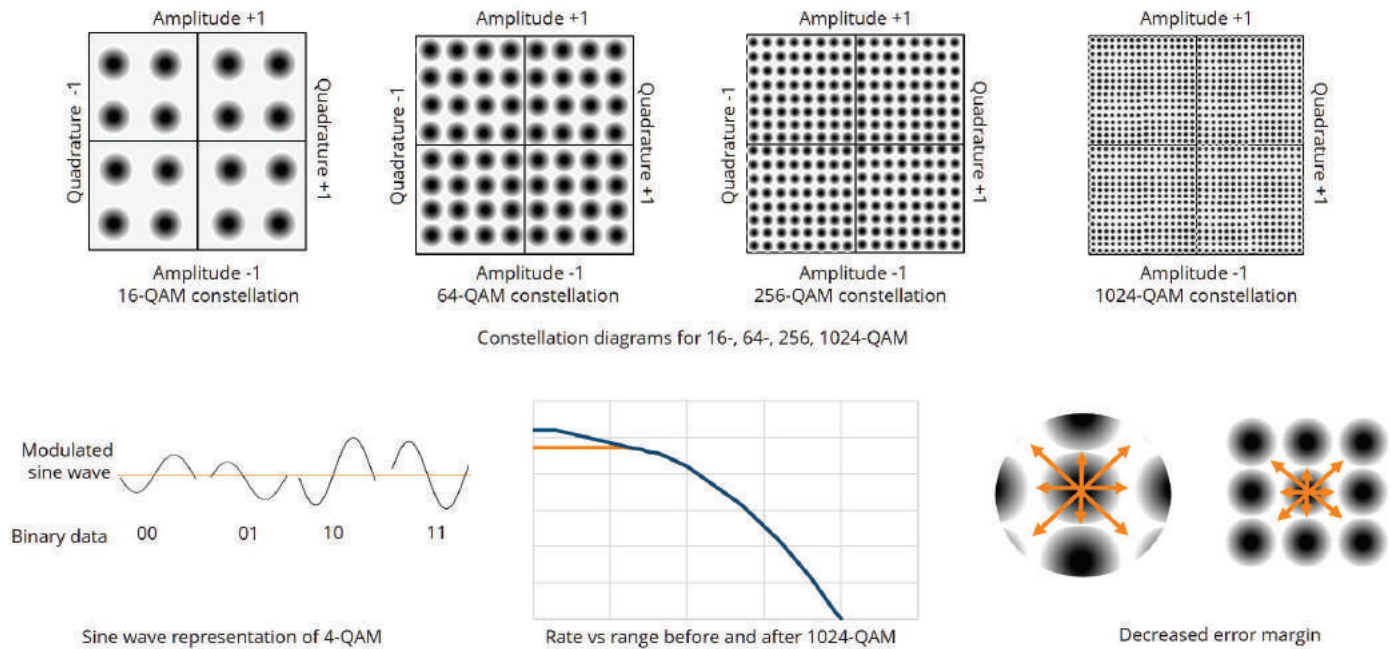


Figure 26: 1024-QAM modulation

The move from 256-QAM to 1024-QAM increases the number of bits carried per OFDM symbol from 8 to 10, for a data-rate and spectral-efficiency boost of 25%. But, as before, the improvement only works for the cleanest conditions, where the signal level is high and the noise low. This is because the receiver has to make a decision about the modulation level, choosing one of 32 states along each axis (amplitude and phase or quadrature) rather than one of 16 for 256-QAM or one of 8 for 64-QAM.

The chart below shows that the receive power level required to decode an 80 MHz, 1024-QAM 5/6, MCS-11) frame is close to -45 dBm, a very high level. This emphasizes how some 802.11ax rate tables are shown below (160 MHz channels are just 2x 80 MHz).

The first table shows data-rates for 20, 40, 80 MHz channels (in Mbps, with short guard interval).




Figure 27: 802.11ax receive sensitivity requirements

Table 3: 802.11ax selected rates (Mbps, short GI)

MCS	Modulation & Rate	20 MHz 1x SS	20 MHz 2x SS	20 MHz 4x SS	20 MHz 8x SS	40 MHz 1x SS	40 MHz 2x SS	40 MHz 4x SS	40 MHz 8x SS	80 MHz 1x SS	80 MHz 2x SS	80 MHz 4x SS	80 MHz 8x SS
0	BPSK 1/2	8.6	17.2	34.4	68.8	17.2	34.4	68.8	137.6	36.0	72.1	144.1	288.2
1	QPSK 1/2	17.2	34.4	68.8	137.6	34.4	68.8	137.6	275.3	72.1	144.1	288.2	576.5
2	QPSK 3/4	25.8	51.6	103.2	206.5	51.6	103.2	206.5	412.9	108.1	216.2	432.4	864.7
3	16-QAM 1/2	34.4	68.8	137.6	275.3	68.8	137.6	275.3	550.6	144.1	288.2	576.5	1,152.9
4	16-QAM 3/4	51.6	103.2	206.5	412.9	103.2	206.5	412.9	825.9	216.2	432.4	864.7	1,729.4
5	64-QAM 1/2	68.8	137.6	275.3	550.6	137.6	275.3	550.6	1,101.2	288.2	576.5	1,152.9	2,305.9
6	64-QAM 3/4	77.4	154.9	309.7	619.4	154.9	309.7	619.4	1,238.8	324.3	648.5	1,297.1	2,594.1
7	64-QAM 5/6	86.0	172.1	344.1	688.2	172.1	344.1	688.2	1,376.5	360.3	720.6	1,441.2	2,882.4
8	256-QAM 3/4	103.2	206.5	412.9	825.9	206.5	412.9	825.9	1,651.8	432.4	864.7	1,729.4	3,458.8
9	256-QAM 5/6	114.7	229.4	458.8	917.6	229.4	458.8	917.6	1,835.3	480.4	960.8	1,921.6	3,843.1
10	1024-QAM 3/4	129.0	258.1	516.2	1,032.4	258.1	516.2	1,032.4	2,064.7	540.4	1,080.9	2,161.8	4,323.5
11	1024-QAM 5/6	143.4	286.8	573.5	1,147.1	286.8	573.5	1,147.1	2,294.1	600.5	1,201.0	2,402.0	4,803.9

(The maximum rate for a 160 MHz channel with 8x SS and MCS 11 is now 9607.8 Mbps.)

And the table below shows data-rates for sub-channels (in Mbps, with short guard interval).

Table 4: 802.11ax selected rates (Mbps, short GI)

MCS	Modulation	RU-26	RU-52	RU-106	RU-242	RU-484	RU-996
		2.0 MHz	4.1 MHz	8.3 MHz	18.9 MHz	37.8 MHz	77.8 MHz
0.0	BPSK 1/2	0.9	1.8	3.8	8.6	17.2	36.0
1.0	QPSK 1/2	1.8	3.5	7.5	17.2	34.4	72.1
2.0	QPSK 3/4	2.6	5.3	11.3	25.8	51.6	108.1
3.0	16-QAM 1/2	3.5	7.1	15.0	34.4	68.8	144.1
4.0	16-QAM 3/4	5.3	13.2	28.1	51.6	103.2	216.2
5.0	64-QAM 2/3	7.1	14.1	30.0	68.8	137.6	288.2
6.0	64-QAM 3/4	7.9	15.9	33.8	77.4	154.9	324.3
7.0	64-QAM 5/6	8.8	17.6	37.5	86.0	172.1	360.3
8.0	256-QAM 3/4	10.6	21.2	45.0	103.2	206.5	432.4
9.0	256-QAM 5/6	11.8	23.5	50.0	114.7	229.4	480.4
10.0	1024-QAM 3/4				129.0	258.1	540.4
11.0	1024-QAM 5/6				143.4	286.8	600.5

Transmit power control in multi-user mode

Multi-user modes in 802.11ax allow much more control over transmit power levels, and most of the control lies at the AP. This should be useful for clients' battery life, and for limiting co-channel interference as, for example, clients currently tend to transmit at maximum power even though APs may be on reduced power in a dense multi-AP deployment, increasing the interference radius.

As a result of the sounding procedure, an AP learns how its clients are receiving its signals, which allows it to estimate the path loss and RF channel conditions. Thus it can adjust its transmit power to target a particular signal level at the client, or more often, a signal-to-noise-and-interference (SINR) level. Since MCS and error rates are related to SINR, it can choose to optimize by reducing the error level, or increasing the MCS and/or transmit power to increase data rates and reduce time on the air.

One interesting possibility is to increase the power transmitted in certain OFDMA RU's, while reducing that used in others. This is interesting because it opens an opportunity for 'water-filling', a technique to allocate resources to the most-effective recipient, but also allows the AP to transmit above the allowed power levels (EIRP) for certain sub-carriers, while reducing power on others. So long as the overall EIRP on a 20 MHz channel is within limits, this configuration would be allowed by regulation.

With the new multi-user signaling mechanisms, the AP can now control the client's transmit characteristics. In wave 1 this is only applicable to OFDMA, but in wave 2, uplink MU-MIMO will be controlled as well. The fields controlling uplink multi-user operation allow the AP to specify the transmit power indirectly, as a desired SINR at the AP, which the client can derive from the sounding estimate of path loss. But more than that, the AP specifies the number of spatial streams to use, the MCS, the OFDMA RU and other features that should be used. The AP may do this to optimize packing of an OFDMA packet, but alternatively it could seek to reduce the transmit power used by particular clients, for battery life or interference optimization purposes.

Power-save mechanisms in 802.11ax

One of the goals of the 802.11ax project is to improve performance by a factor of 4x while keeping power requirements unchanged or improved. With the emerging IoT market, power-save mechanisms at the other end of the performance scale were also a particular focus. Several power-save mechanisms already exist in prior 802.11

standards: these remain, and are supplemented with a new mechanism, 'target wait time' (TWT). (TWT was introduced in 802.11ah, the amendment for low-power, long-range IoT transmission; but since 802.11ah chips and devices have not been widely adopted by the market, it is new to users of Wi-Fi equipment.) TWT is particularly useful for battery-powered devices that communicate infrequently.

The existing 'legacy PS' mechanism has been in use since 802.11b, the first widely-used Wi-Fi standard. Clients can sleep between AP beacons, or multiples of beacons, waking when they have data to transmit (they can transmit at any time, the AP does not sleep) and for beacons containing the delivery traffic information map (DTIM), a bit-map indicating that the AP has downlink data buffered for transmission to particular clients. If the DTIM bit is set for a client, it can retrieve its data by sending a trigger frame to the AP immediately after the beacon. PS is an effective mechanism but only allows clients to sleep for a small number of beacon intervals, usually clients must wake several times per second to read the DTIM.

As explicit voice-over-Wi-Fi support was added with 802.11e, the IEEE recognized that voice-capable devices required a new power save mechanism, as voice packets are transmitted at short time intervals, typically 20 msec. Unscheduled automatic power-save delivery (U-APSD) allows a client to sleep at intervals within a beacon period. As in PS, the AP buffers downlink traffic until the client wakes and requests it. With symmetrical traffic like voice, the client can often send and receive frames in the same waking interval.




Figure 28: Power-save options before 802.11ax

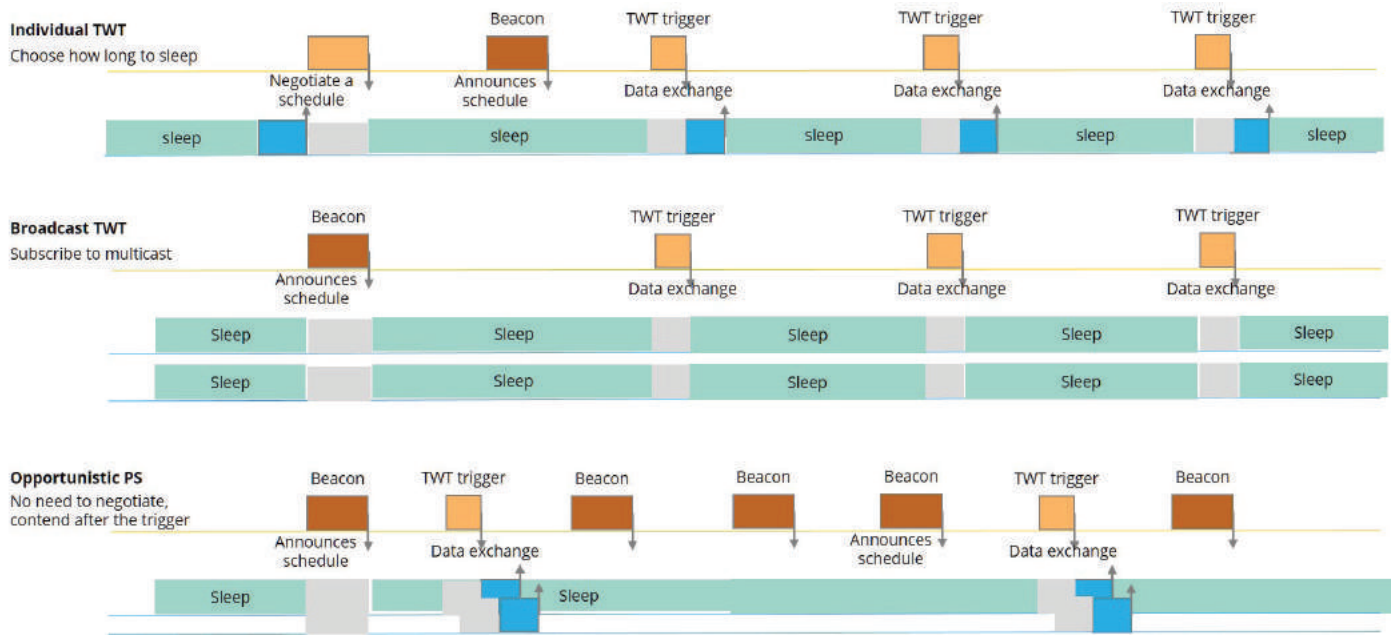


Figure 29: TWT power-save options in 802.11ax

The new TWT mechanism in 802.11ax allows more flexible, long-term and even multi-client sleeping arrangements. First, a negotiation between the client and AP sets up an agreed schedule for the client to wake and communicate. The schedule is often periodic, with a long, multi-beacon interval (minutes, perhaps hours or days) between activities. When its designated time arrives, the client wakes, awaits a polling trigger frame from the AP (required in multi-user mode) and exchanges data, subsequently returning to the sleep state. Since the AP negotiates separately with each client, it can group or separate scheduled transmissions in order to achieve best traffic efficiency or to accommodate traffic requirements from other clients.

The standard allows several variations of the individual TWT described above. Multicast traffic, that many clients wish to receive, can be set up by the AP on a schedule published in beacons. Opportunistic power save allows the AP to publish a schedule of intervals when any client can wake and request a packet exchange, even within OFDMA. And there are mechanisms for unassociated clients to learn when information they might be interested in will be broadcast.

Various multi-user modes can also be used with TWT, making some options rather complicated. But the straightforward goal is to enable flexible and long-term sleep intervals.

802.11ax is very power-aware: in addition to TWT, it has many other features that can extend the battery life of IoT sensors and other clients.

- The uplink/downlink bit identifies the frame as transmitted by and AP or client device. It is included in every preamble. This is useful because client devices do not need to receive frames from other clients, and can switch off their radio circuitry as soon as they see an 'uplink' bit in a preamble.
- The 20 MHz-only option allows new designs for stripped-down chips optimized for long battery life. It also offers a lower-power mode for mainstream Wi-Fi chips in specialized equipment.
- Multi-user signaling allows the AP to indicate to clients what signal strength it needs to see on their transmissions, and to specify an MCS to be used. If the AP has knowledge of its clients, it can optimize these settings for power-constrained devices.
- The BSS-coloring feature allows clients to stop receiving a frame and return to sleep mode as soon as they recognize the frame is not of interest to them. This is a wave 2 feature (see appendix).
- 'Receive operating mode' and 'Transmit operating mode' allow clients to reduce the number of active transmit and receive chains they use for data transmission, as well as the channel width. A client can use this to reduce the peak power requirement for sending and receiving data. Although lower data-rates will mean the transmission time is extended, this can be a worthwhile saving for an IoT sensor.

20 MHz-only operation

Since 802.11ax was developed in the IoT era, it offers more than a nod towards low-cost, battery-powered client devices. Wi-Fi chip shipments number in the billions per year, allowing manufacturers to bring unit costs down to very low points for such sophisticated chips, but Wi-Fi remains a more expensive option for an IoT sensor compared to some of the alternatives such as Bluetooth or Zigbee. And Wi-Fi chips are several times more power-hungry.

Over the years, several specialist chip vendors have brought out modified 802.11 chips for IoT applications (for example, Wi-Fi asset tags) which minimized power requirements in as many ways as possible, but the market for these special chips was low-volume and the designs were constrained by the need to interoperate with other Wi-Fi equipment including access points.

So the developers of 802.11ax, many of whom are employed at chip vendors, sought to close these gaps in complexity and battery-life in the new standard in several ways. The new TWT protocol for power-save should allow Wi-Fi-based IoT sensors to operate at a considerably lower power draw: this will narrow the gap on battery life, but the complexity and footprint of Wi-Fi chips leave them at a disadvantage for very low-cost sensors.

802.11ax seeks to reduce complexity by opening the door for a new class of chip. This will lead to very beneficial consequences if the vendors take their lead from the IEEE and build this new class of Wi-Fi chip.

A 20 MHz-only device is capable of operating in either the 2.4 or the 5 GHz band, but only in 20 MHz at a time, on the designated primary channel. But nearly all other mandatory features of 802.11ax apply, including OFDMA options, allowing such a device to transmit and receive on a much smaller sub-channel. Of course, the 20 MHz channel cap limits the data-rates that can be supported, but this should not be an issue for an IoT application.

(Access points must support the full 20 MHz channels in the 2.4 GHz band and 20, 40 and 80 MHz channels at 5 GHz to be certified by the Wi-Fi Alliance: 20 MHz-only applies just to client devices.)




Figure 30: 20 MHz-only operation

USAGE MODELS: WHAT CAN WE DO WITH 802.11AX?

802.11ax business benefits	802.11ax major use cases
<ul style="list-style-type: none"> • Increased aggregate network throughput • Increased rate at range • Peak link throughput increase • Reduced overhead • Increased efficiency in dense networks • Increased robustness outdoors • Reduced power consumption • Enhanced Wi-Fi coexistence 	<ul style="list-style-type: none"> • Airports and stations • Education • Shopping mall • Wireless office • Smart cars • Stadium • Smart cities • Dense apartment buildings • Dense suburban homes

Figure 31: Usage models

Dense deployments and overlapping access points

Improved performance in dense networks is perhaps the primary goal of 802.11ax. Dense networks can take different forms: large numbers of clients in a small area, or closely-spaced access points, and sometimes overlapping access points that may have common or entirely separate management. The new standard and certification offer solutions for all these scenarios.

For an isolated access point with small numbers of high-rate clients, 802.11ax supports increased data-rates and improved multi-user MIMO operation for simultaneous transmissions wherever possible. Where data streams include many short frames, multi-user OFDMA enables much lower contention and preamble overhead.

When access points overlap, spatial re-use through BSS coloring reduces the common-channel interference radius of an access point, improving simultaneous transmission across a wide area, and hence network capacity.

And OFDMA and multi-user uplink control from the access point should drive performance improvements in situations where client populations are heterogeneous with varying data-rate and frame-length.

Low-power, large-scale: Internet of Things

Whether or not it is called 'IoT', the industry is looking forward to a much wider range of Wi-Fi connected clients in the 802.11ax era. It is hoped that the power-saving features in the new certification, particularly TWT, 20 MHz-only, even some of the multi-user control functions and OFDMA will all contribute to extending battery life far enough to make inroads into the emerging IoT market.

Along with its lower-power requirements, IoT will certainly increase the number of Wi-Fi devices in a home, office building or even retail store. Several features in 802.11ax extend the number of client devices that can be associated with an access point, and more important, the amount of simultaneous active clients. In particular, OFDMA allows the frequency domain to be sliced much thinner, favoring large numbers of devices with low data-rates and long sleep-times, one definition of IoT.

Long-range, outdoor operation

The outdoor point-to-point, point-to-multipoint and mesh markets are often overshadowed by the home and business WLAN segments, but they represent sizeable and consistent markets for Wi-Fi equipment, and will benefit from several improvements in 802.11ax.

Moreover, one of the more significant targets of the 5G project from the cellular world is 'fixed wireless access' (FWA) where wireless mesh networks deliver broadband Internet service to the home in urban and rural settings. This market requires long-reach links, high data-rates and low-cost equipment: 802.11ax improves on the first two, and the third has always been one of Wi-Fi's strengths. Indeed, early FWA entrants are touting their equipment as 'pre-standard 5G' when in fact they are using frequency-shifted 802.11ac: with 802.11ax the case for using Wi-Fi chips will be even stronger. It is hoped that the improvements in long-range outdoor operation in 802.11ax will spur greater penetration of this and other new markets.

BACKWARDS COMPATIBILITY, DEPLOYMENT CONSIDERATIONS AND UPGRADE STRATEGIES

Wi-Fi can boast a near-flawless backwards-compatibility record. Thanks to the legacy training fields in every packet preamble, even 15 year-old 802.11g equipment is able to decode 802.11ax frames. While the 802.11n standard included an optional 'greenfield' mode, it was never implemented in shipping equipment, and since then there have been no 'greenfield' options.

Even the extended-range outdoor features, which because of dual-beacons and other special frames will be incompatible with older equipment, will be protected by co-locating a 'legacy' AP beacon.

This allows different strategies when upgrading an 802.11n or 802.11ac WLAN to 802.11x. Some may like to intersperse new APs in a salt-and-pepper topology, while others may upgrade a whole floor or corner of a building at one time. Either will work.

The increased sustained throughput of an 802.11ax AP may also prompt a backhaul upgrade. While many APs support dual 1 Gbps Ethernet connections, the move to 2.5 and 5 Gbps Ethernet seems to be attractive over the long term.

MANDATORY AND OPTIONAL FEATURES

Table 5: 802.11ax major features: mandatory and optional

Access Point		Client	
Mandatory	Optional	Mandatory	Optional
Downlink OFDMA transmit		Downlink OFDMA receive	
Uplink OFDMA receive		Uplink OFDMA transmit	
Downlink MU-MIMO transmit (if 4+ SS)	Downlink MU-MIMO transmit (if < 4 SS)	Downlink MU-MIMO receive (up to 4x SS)	
Transmit beamforming (if 4+ SS)		Receive beamforming	
SU MIMO transmit & receive with up to 2x SS	SU MIMO with 3+ SS	SU MIMO transmit & receive	
20, 40, 80 MHz operation if supporting 5 GHz		20, 40, 80 MHz operation if supporting 5 GHz	
20 MHz operation if supporting 2.4 GHz		20 MHz operation if supporting 2.4 GHz	
20 MHz-only operation in wideband OFDMA			Individual TWT
Individual TWT		BSS coloring	Spatial re-use
BSS coloring	Spatial re-use	Transmit & receive operating mode	
Transmit & Receive operating mode			
	MCS 8, 9, 10, 11 (256 & 1024-QAM)		
	160 MHz operation (if supporting 5 GHz)		

802.11ax wave 1 and wave 2

It is already established that 802.11ax will roll out with wave 1 and wave 2, but the exact split of features has not been frozen. This is the current view at the time of writing (early 2018).

Table 6: 802.11ax wave 1 and wave 2 features

Wave 1	Wave 2
Downlink and uplink OFDMA	Uplink multi-user MIMO
Downlink MU-MIMO	
Target Wait Time (TWT)	
BSS Coloring	Spatial re-use
20 MHz-only	
	Long-range 802.11ax

In this paper we have taken the wave 2 features and moved them to an appendix, as they probably will not be seen in the field until 2020.

Wi-Fi Alliance certifications for 802.11ax

As of early 2018, the Wi-Fi Alliance plans to certify four different types of equipment under the 'Wi-Fi CERTIFIED AX' certification (these plans are subject to change, certification is expected to launch in mid-2019).

- **Wave 1 AP:** This would be the access points we are familiar with in residential or enterprise environments.
- **Wave 1 client device:** The usual client in smartphones, PCs and other consumer and enterprise devices.

- **Mobile AP:** The certification for 802.11ac included a special option for lower-power, smaller devices that include AP functionality, and 802.11ax continues to include the category. It applies primarily to battery-powered, mobile access points where functionality and performance can be limited. Mobile APs do not need the 'Agile Multiband' certification, and many of the requirements around OFDMA and MIMO are relaxed, allowing Mobile APs to be single-stream devices with data-rates up to MCS5. Mobile APs can be 2.4 GHz-only, or dual-band, but if dual-band they must support 20, 40 and 80 MHz channels.
- **20 MHz only client device:** This category opens the door to the IoT sensor market. Again, many features that are mandatory for standard client devices become optional, allowing simpler devices that have longer battery life but lower functionality.

It is well-known that the total data-throughput of an access point falls as the number of client devices increases. This is explained by contention overhead: as the number of clients competing for a transmit opportunity increases, the average wait also increases, and the amount of air-time used for data transmission decreases.

The graph above shows network capacity increasing at first, as when the client count is low, they cannot fill the available bandwidth. Once past this threshold, capacity falls by as much as 40% when 100 clients are present, and the trend continues beyond that. The orange series shows that Aruba's models predict a significant improvement after 802.11ax. (This, and subsequent graphs are measured for 2x SS operation).

PERFORMANCE ESTIMATES

While 802.11ax is not explicitly targeting peak data-rates, overall aggregate performance is the most important goal. To this end, the following section offers some explanations of current challenges in Wi-Fi performance, along with projections of how performance is expected to improve.

Large numbers of clients




Figure 32: Network (BSS) capacity vs number of clients

Short packets





Figure 33: Network (BSS) capacity vs packet length

As packet length decreases, the preamble and contention overhead stays constant, per-packet. This increases the overhead and decreases the overall access point capacity.

802.11ax includes a number of measures to improve this reduction of capacity with small packet sizes: the orange series above shows the results of Aruba's models.

Narrow channels (OFDMA)



* OFDMA is not quite the same technical effect - model and chart for illustration

Figure 34: Capacity vs channel width

It is generally accepted that, for highest aggregate throughput, a dense WLAN with many APs and clients in a small area should be configured for large numbers of 20 MHz channels, rather than fewer 40 or 80 MHz channels. This is partly because of the co-channel interference issue addressed by BSS coloring, but also because it reduces the number of client devices per AP and minimizes overhead. OFDMA can be seen as a continuation of this trend, and we expect it to contribute to system-level improvements in data capacity through a number of mechanisms.

The graph above is not technically consistent – the increased capacity afforded by OFDMA is not quite the same as the small-channel trend, but it serves to illustrate the effect of OFDMA on network capacity.

MIMO and multi-user effects

The capability to transmit data simultaneously from or to different antennas and devices is very powerful. First introduced with single-user MIMO in 802.11n, it was extended with downlink multi-user MIMO in 802.11ac wave 2, and

again with OFDMA in both directions in 802.11ax. Simulations indicate that MU-MIMO is most effective with long packets and high SNR, while OFDMA is effective over the whole range of client densities and packet length distributions.

Power-saving and battery-life

Simulations predict that use of TWT, especially in conjunction with 20 MHz-only operation and OFDMA will enable prolonged battery-life. Some estimates suggest Wi-Fi sensor battery life could approach Bluetooth Low-Energy (BLE), although this is conjecture until chips are shipped and equipment is built and tested.

Spectral efficiency of 802.11 over the years, bps/Hz

With the recent interest in LTE and 5G waveforms in unlicensed spectrum, the topic of spectral efficiency—the data-rate that can be achieved per Hz of spectrum—has enjoyed renewed interest.

Wi-Fi has a strong record in spectral efficiency, although as can be seen below, the figure is heavily dependent on MIMO effects.




Figure 35: Historical spectral efficiency of 802.11, bps/Hz

CHANNELIZATION

The available RF channels in most countries have not changed very much since 802.11ac was introduced: in the USA, the FCC announced changes in 2014 that now allow outdoor operation in the 5150-5250 MHz band, and lift the temporary ban on operation at 5600-5650 MHz, subject to a modified DFS test.

But, with an eye to the future, the Wi-Fi industry is taking several initiatives to lobby for more unlicensed spectrum that could be used by Wi-Fi.

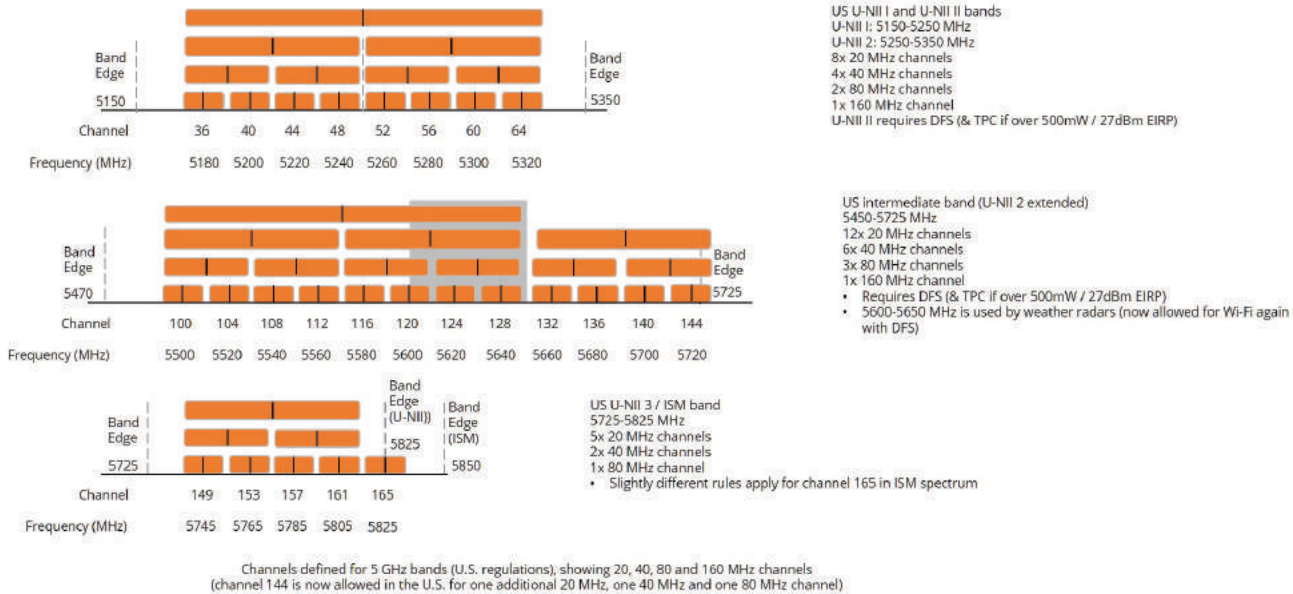


Figure 36: Available 5 GHz Channels for 802.11ax (USA)

CONCLUSION

The coming years hold much promise for Wi-Fi, but also great uncertainty. 802.11ax is the Wi-Fi industry's response to these opportunities and challenges, to take us through the next five years from 2019 to 2024 and the next 802.11 amendment.

Wi-Fi chips are shipping at a rate of 3 billion a year, with an installed base of 8 billion. Every smartphone and PC comes with a Wi-Fi chip, every broadband home Internet connection terminates on Wi-Fi, the technology is established in outdoor point-to-point links and making inroads into the automotive industry and connected factories.

But, as with every successful industry, Wi-Fi is looking for even faster growth. Several avenues are opening, but changes will be necessary in order to meet their needs.

It is universally acknowledged that the Internet of Things (IoT) will grow into a huge market in the next few years. But IoT often requires battery-powered devices, wireless connections reaching out for hundreds of meters, and very

low-cost, small-footprint chips. Wi-Fi's first assault on this market was 'extended range ah' (based on IEEE 802.11ah). It is a comprehensive standard, but has been stalled for some years: for a number of reasons the chipmakers are not making 'ah' chips. As a result, Wi-Fi risks missing a large segment of the IoT market. Several features in 802.11ax make it more attractive for IoT, but it remains to be seen whether Wi-Fi will win a significant footprint in this market.

Meanwhile, regulatory changes are progressing. As a specific example, the 'citizens broadband radio service' (CBRS) initiative from the USA's FCC offers a lightly-licensed band where it was—at one time—hoped that private organizations could easily purchase a semi-exclusive license to use spectrum in the 3.5 GHz band. There has been a lot of activity around CBRS, and some uncertainty as the rules are still not final. Nevertheless, it seems that public and private LTE and 5G services will use this band, while Wi-Fi has no product plans and has ceded the initiative.

Other, more sweeping changes are in motion in spectrum and regulation. Amidst a land-grab by mobile operators, seeking more 5G spectrum, the incumbents—government and military, satellite users, radars and others—are pushing back on spectrum sharing proposals, and Wi-Fi is caught with limited lobbying budgets and two frequency bands, at 2.4 and 5 GHz, that have a relatively large span of spectrum but are increasingly congested. Indeed, one of the threats to Wi-Fi is the growing perception that the 2.4 GHz band is ‘junk’ in high-rise buildings and city centers due to overuse. The industry is not supine: a robust lobbying effort in 2017-19 seeks to open up bands across the 6 GHz range, and these efforts are receiving a sympathetic hearing from regulators. But the outcome of all these threads is uncertain: the race to secure the spectrum Wi-Fi needs for continued success is by no means over.

And the 5G vision—powerful, radical, persuasive and comprehensive—looms over all of Wi-Fi’s short- and medium-term plans. Building on their strength in LTE, mobile operators seek to use 5G to break out of the consumer-phone market, into many of the markets Wi-Fi dominates today including home broadband, connected cars, factories and cities, and enterprise networking. In many ways 802.11ax is Wi-Fi’s answer to the 5G vision, but — because of the differences between the Wi-Fi and mobile operator ecosystems—it is narrower and less detailed. Wi-Fi companies face a short-term perception that 5G is a comprehensive answer to many of the problems they are addressing piecemeal.

Despite all these challenges, Wi-Fi can face the future with great optimism. In a 20-year span, it has emerged from nowhere to become a household name, and a technology used everywhere. The features included in 802.11ax—following considerable debate among the companies that support Wi-Fi—are not only for faster headline data-rates under ‘best case’ conditions, but responses to the practical real-world issues that are often due to Wi-Fi’s great success.

Features such as BSS coloring, multi-user scheduling and backwards-compatibility will greatly improve performance in congested areas, where many uncoordinated Wi-Fi access points and client devices operate in close proximity.

Where access points are under common management, as in sports stadiums, airports, lecture theaters and convention centers, the new features offer even greater control, for higher network capacity and more even performance across the user population.

OFDMA is ideal for clients transmitting short packets, and for low-bandwidth devices such as IoT sensors, which will also benefit from new power-save features including TWT. These and the 20 MHz-only feature should spur a generation of very-low-power chips tailored to this market.

And an often neglected market for Wi-Fi equipment, the outdoor point-to-point and point-to-multipoint wireless, will be boosted by features for extended range and higher interference immunity. Indeed, early contenders in the new wave of fixed-wireless-access networks—targeted by 5G - for home broadband services to rural communities are using Wi-Fi chips in preference to 4G or 5G because of their superior price and performance.

How will it turn out? The Wi-Fi ecosystem, much less-coordinated and regimented than the cellular world, has shown it is able to assemble features in unexpected combinations and solve new problems. The features in 802.11ax give the companies that by now have deep expertise in Wi-Fi technology the tools to move into already-emerging markets and react to new opportunities. The most likely outcome is that history will repeat itself and Wi-Fi will continue its growth amid the coming changes.

APPENDICES (EXISTING FEATURES AND NEW 802.11AX FEATURES DEFERRED TO WAVE 2)

Appendix 802.11ac existing MIMO features carried over to 802.11ax

All of the features used for MIMO and beamforming in 802.11ac are carried forward to 802.11ax: they are explained below for completeness. 802.11ax implements downlink MU-MIMO (with a new control structure) in wave 1, and uplink MU-MIMO in wave 2.

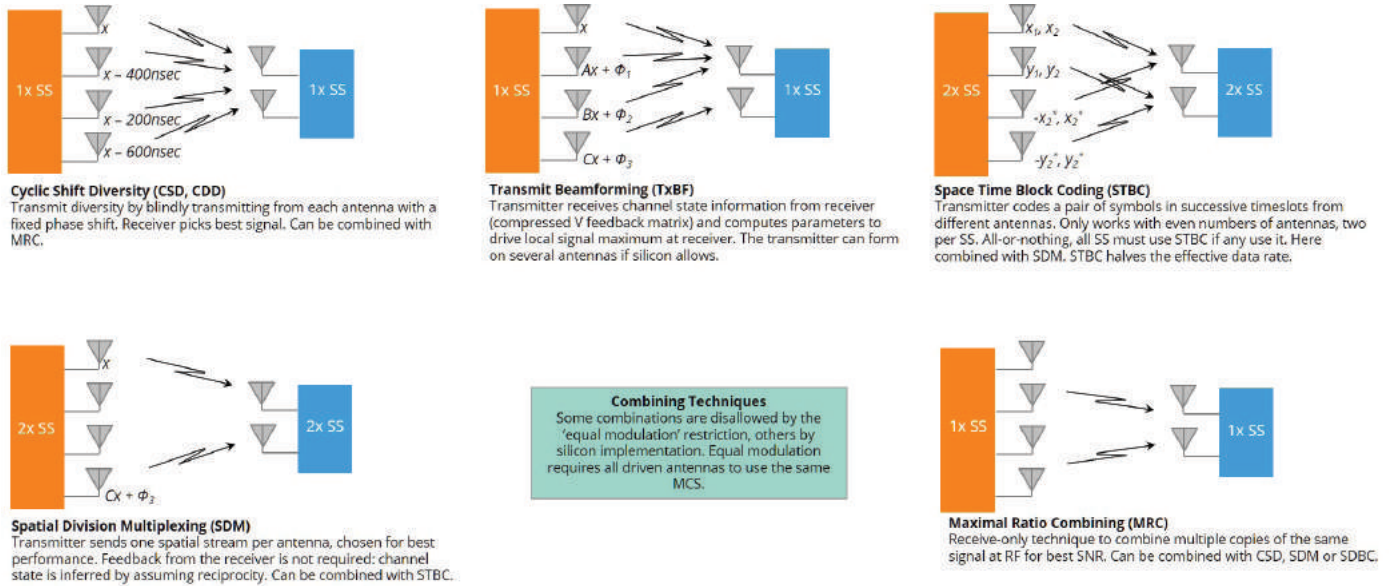


Figure 37: SU MIMO techniques, multi-antenna client

Beamforming in 802.11ax always uses explicit feedback from the client device to the AP to calculate the optimal transmit signal weightings. This is unchanged from 802.11ac.

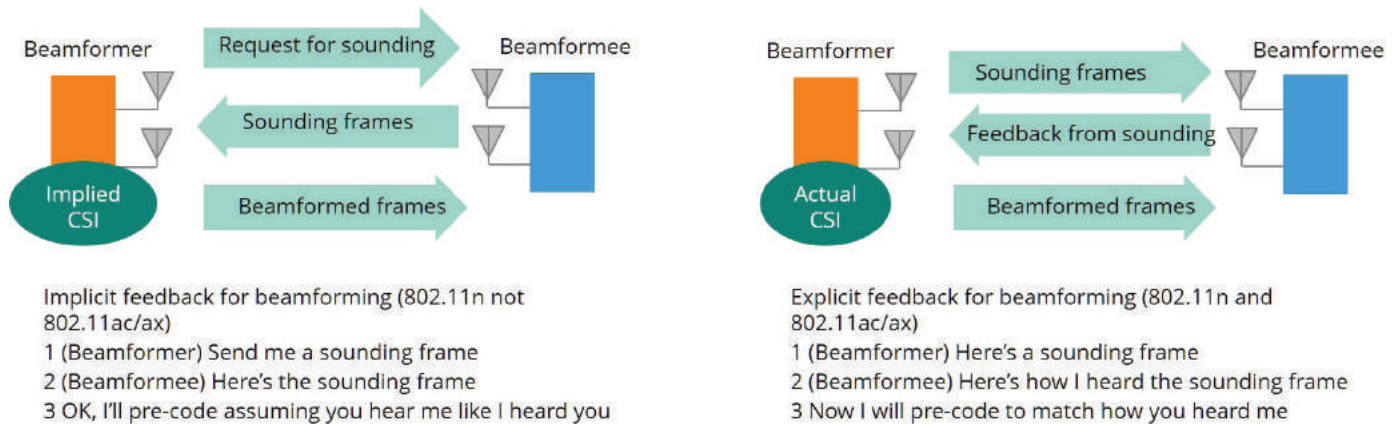


Figure 38: Beamforming feedback

Appendix 802.11ax wave 2 features—uplink MU-MIMO

Uplink multi-user MIMO is simpler than the downlink variety: sounding and antenna weighting is not required.



Figure 39: Uplink multi-user MIMO transmission

The most difficult task in uplink multi-user MIMO is synchronizing the waveforms from different clients as they reach the AP's antennas.

When it identifies such groups, the AP considers several aspects of their signals and traffic, then instructs them to transmit simultaneously as appropriate.

This feature uses the same control mechanisms as OFDMA in the uplink direction, but it is deferred to wave 2 due to its complexity.

Appendix 802.11ax wave 2 features—Combined MU-MIMO and OFDMA

As noted above, the combination of MU-MIMO and OFDMA in a single packet, while quite possible in theory, becomes complicated in practice and support is deferred to wave 2.

Appendix 802.11ax wave 2 features—Spatial re-use

From the beginning, Wi-Fi has used a medium-access protocol called CSMA/CA, where all devices including APs

sense the air when they have a packet to transmit. If they sense energy above a certain power threshold, or a Wi-Fi frame on the air, they defer transmission and use a backoff algorithm to return to the air at a later time, again sensing to ensure it is clear before transmitting. The CSMA/CA protocol has been very successful for Wi-Fi. It is distributed, meaning each device makes transmit decisions independently, which allows for overlap of different, uncoordinated basic service sets (BSS's, or cells, each with an access point).

It is due to CSMA/CA that Wi-Fi access points can be set up adjacent to each other, with no coordination of channel or other configuration, and each can support successful communication with its clients. Wi-Fi is very good at 'sharing' the air and requires none of the cell-by-cell management of, for instance, the cellular network. But when many uncoordinated BSS's overlap in space, as in city centers or apartment buildings, CSMA/CA can be quite inefficient in terms of network capacity.




Figure 40: BSS coloring: co-channel interference

The reason is that wireless signals are not neatly confined, as we usually draw with circles or hexagons, but spread over distance. If an AP and client device use sufficient transmit power to communicate reliably across a BSS's radius, their signals are still quite high at the edge of the 'cell', and spread across neighboring cells as they decay over distance. In an enterprise network, where APs are managed, the RF channel plan tries to separate frequency reuse, so that APs on the same channel are kept far apart and do not interfere with each other. But with 80MHz channels (introduced with 802.11ac) in widespread use, even the 5 GHz band supports only a few non-overlapping channels. Thus, even in enterprise networks neighboring APs on the same channel are sometimes unavoidable, and when APs are uncoordinated, same-channel or co-channel interference can be common. In such scenarios, even though two or more APs and many client devices are present, only one AP or device can transmit at a time, as CSMA/CA causes all the others to defer.

BSS coloring works by distinguishing between 'same BSS' and 'distant BSS' transmissions and applying different CSMA/CA power thresholds. This allows simultaneous transmissions

in the different cells, as, in addition to two power thresholds, each client device keeps two network allocation vectors (NAV's) which tell it how long the medium will be occupied.

The change is not unambiguously positive: there will be cases where simultaneous transmission results in one or both frames failing at the receivers, due to the reduced signal-to-noise-and-interference-ratio (SINR), but retransmissions will allow error recovery and simulations predict significant capacity enhancement in real-world Wi-Fi deployments because of this feature.

Note that while the option to use 'BSS color' labeling in 802.11ax transmissions is part of the standard, the algorithm and mechanisms for configuring APs with appropriate 'colors' is left to equipment vendors. We may see a range of automatic and semi-automatic configuration features as 802.11ax rolls out.

Appendix 802.11ax wave 2 features—Outdoor and long-range operation

802.11ax includes a number of features that are useful for outdoor and long-range operation, chiefly aimed at point-to-point links. The main innovation is the single-user extended-range frame.

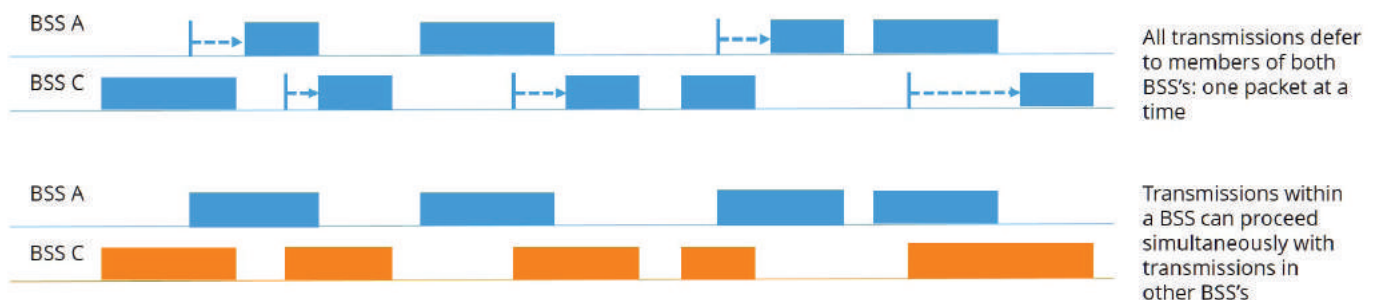


Figure 41: BSS coloring: before and after




Figure 42: Single-user extended-range frame

The frame includes a repeated HE-SIG-A field in the preamble, for better resistance to signal fading and interference. It is specified to use only one spatial stream, and low data-rates for similar reasons.

In long-range operation, the longer cyclic prefix options (1.6 and 3.2 usec) could be used with this frame to combat long-delay multipath.

Other measures to extend range include:

- An option to use a 10 MHz bandwidth for the packet (the 106-RU sub-channel), which would be used with

- The dual-beacon, a repeated beacon (because it cannot be read by mainstream devices, this must be accompanied by a conventional beacon)
- Dual sub-carrier modulation (DCM) where the same signal is replicated across two subcarriers. This is new in 802.11ax.

All these techniques go together, and they will probably all be deferred to wave 2.

Appendix—Abbreviations

3GPP	3 rd Generation Partnership Project
AC	Access Category
ADI	Association Identifier
A-MSDU	Aggregated MAC Service Data Unit
A-MPDU	Aggregated Protocol Service Data Unit
AP	Access Point
BAR	Block-Ack Request
BCC	Binary Convolutional Coding
BQRP	Bandwidth Query Report Poll
BRP	Beamforming Report Poll
BSRP	Buffer Status Report Poll
CP	Cyclic Prefix
CSI	Channel State Information
CSMA/CA	Channel Sense Multiple Access with Collision Avoidance
CTS	Clear To Send
DCM	Dual sub-Carrier Modulation
DFS	Dynamic Frequency Selection
DL	Downlink
EIRP	Effective Isotropic Radiated Power

ER	Extended Range
FCC	Federal Communications Commission
FDD	Frequency Division Duplex
FFT	Fast Fourier Transform
GCR	Group Cast with Retries
GI	Guard Interval
HE	High Efficiency
HT	High Throughput
HEW	High Efficiency Wireless
LTF	Long Training Field
MIMO	Multiple-Input, Multiple-Output
NDP	Null Data Packet
OFDM	Orthogonal Frequency-Division Multiplexing
OFDMA	Orthogonal Frequency-Division Multiple Access
L-	Legacy
LDPC	Low Density Parity Check
MAC	Medium Access Control
MCS	Modulation and Coding Scheme
MU	Multiple User

NAV	Network Allocation Vector	TDM	Time Division Multiplexing
NDP	Null Data Packet	TDMA	Time Division Multiple Access
NDPA	Null Data Packet Announcement	UL	Uplink
PE	Packet Extension	VHT	Very High Throughput
PLCP	Physical Layer Convergence Protocol	Wi-Fi	Not an abbreviation
PPDU	PLCP Protocol Data Unit	WISP	Wireless Internet Service Provider

RF	Radio Frequency
RSSI	Received Signal Strength Indication
RTS	Request To Send
RU	Resource Unit
SINR	Signal to Interference and Noise Ratio
SNR	Signal to Noise Ratio
STA	Station
STF	Short Training Field
STS	Space Time Stream
STBC	Space-Time Block Coding
SU	Single User
TCP/IP	Transmission Control Protocol / Internet Protocol
TDD	Time Division Duplex

Appendix—References

- IEEE '15/0132r15 Specification Framework for TGax', May 2016
- IEEE 'P802.11axTM/D2.0 Draft Standard for Information technology', November 2017
- Wi-Fi Alliance 'Marketing Requirements Document for Interoperability Testing of Wi-Fi ax', version 1.1, November 2017 (not publicly available)
- Wi-Fi Alliance 'Marketing Requirements Document for Interoperability Testing of Approved ax Products, Draft Appendix: ax Feature Descriptions', version 1.4, 2017 (not publicly available)