2020 vision on 802.11ax

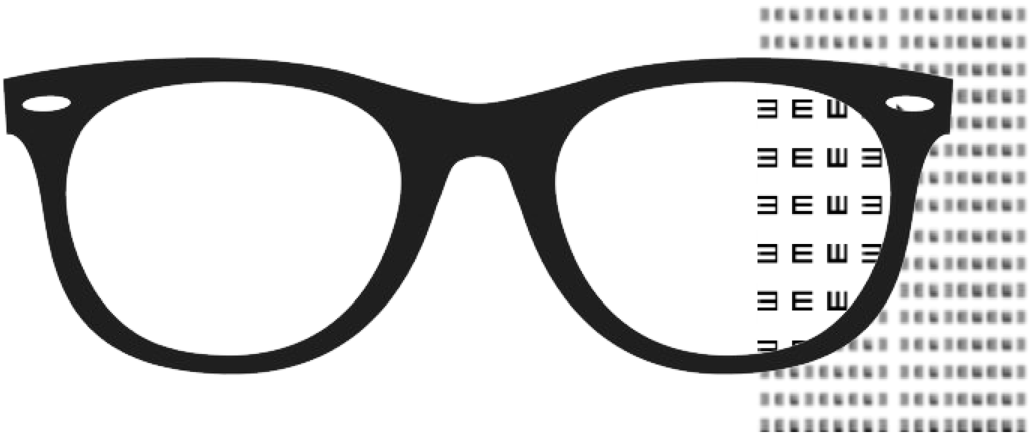

Abstract

802.11ac shows us how fast the network can be. 802.11ax will show us how to maximize all the available spectrum in the 2.4GHz and 5GHz ISM and UNII bands. 802.11ax will focus on improving the connected client experience by making the network more efficient.

Client and standards ecosystem support for 802.11ax is expected in the year 2020. With a clear view of this new technology, network architects will design the decade of the 2020s using 802.11ax. This paper discusses how 802.11ax will address the changing needs of Wi-Fi networks.

Daran Hermans, Product Management
Cambium Networks
2590 North First Street, suite 220. San Jose CA 95519
(888) 863-5250

More streaming content, More devices

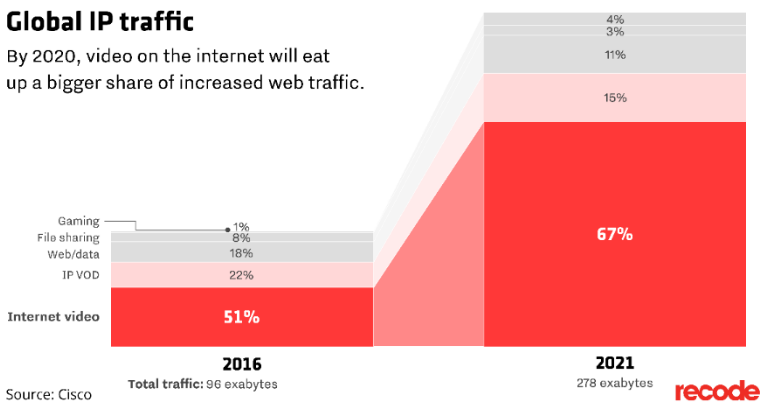

By 2021, streaming video over the Internet will be 82% of all Internet traffic. In terms of bytes transmitted, this represents a 70% increase over 2016.

Streaming video will not be limited to primetime in North America; it is global, all day, all the time. Live video, primarily sports but also news programs, will increase by 95% in the same time. And the quality of video in 2020 will also dramatically increase as more content providers push out 8k video.

Combined, hotels represent the largest “public access networks” in the world. No matter where you travel, you can find a hotel that offers free guest Internet access. Today hotel service providers report an average of 2.75 devices per guest, with 95% of guest devices being a smart phone or tablet. Streaming media is the number one application in total bitrate consumed by guests. Enterprise All Wireless Office, BYOD and online education are also driving higher device density year over year, and higher bitrates.

Voice over Wi-Fi calling offers tremendous advantages and cost savings to the client and the carrier. Calls can be placed and received using the phone dialer, not a separate app. Indoor cellular coverage issues go away, as the phone need only connect to a local 802.11 access point. Voice calls take less than 100kbps for the encapsulated voice packets and overhead. Yet, due to the way 802.11 OFDM operates, a small voice packet consumes one transmission opportunity (TxOP) and the full 20 MHz wide channel.

IOT devices are always-on, always-connected and wireless. In consumer and enterprise markets, IOT devices provide environmental monitoring, lighting and HVAC control, access control, motion alerts and more. As with the voice packets, IOT messages are small data exchanges between the IOT device and the application server. Today, many IOT devices run a custom MAC based on 802.15.4 (aka Zigbee). in the 2020s, low cost 802.11 radios and the concept of running everything in a single access domain may drive significant interest in the use of 802.11 radios for IOT devices.

802.11 networks get faster

802.11n was introduced in 2009; about the same time when smart phone sales were exploding, and tablets were introduced. Since early generation smart phones and tablets had a single antenna element, 802.11n provided a practical limit of 72 Mbps RF rate on a single 20 Mhz wide channel. When network packet overhead and half duplex transmission is factored in, the effective bitrate of the 802.11n networks is about 35Mbps.

802.11n cannot keep up with the device density and high bitrate changes experienced today.

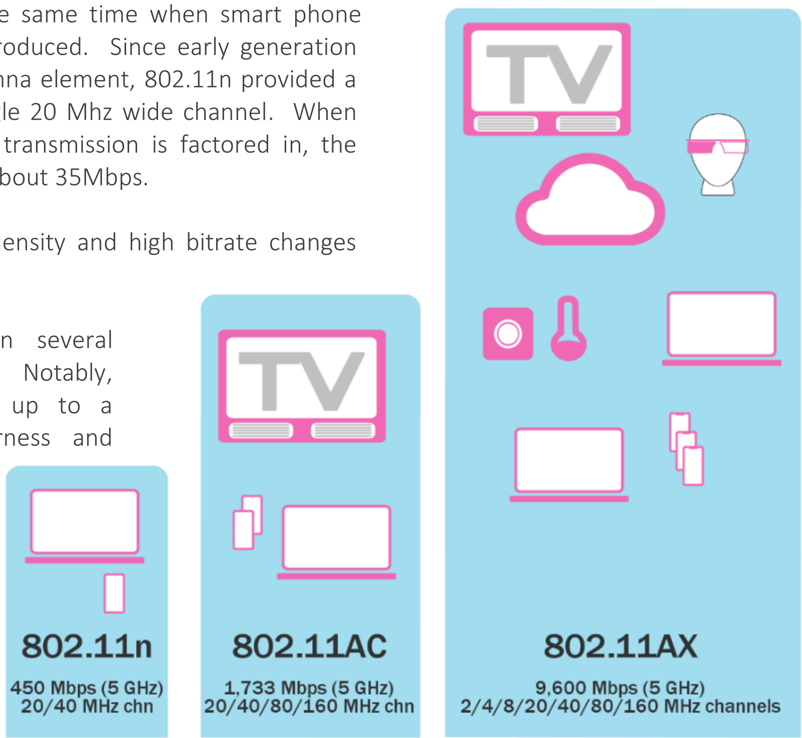

802.11ac introduced in 2013 builds on several successful technologies from 802.11n. Notably, 802.11ac allows wider channel bonding up to a practical limit of 80 MHz, airtime fairness and standardized beam-forming. Also, 802.11ac wave 2 technology introduced multi-user MIMO (DL MU-MIMO) capability to help bridge the gap between the access point capability and the client device capability. Access points are physically larger, support more antennas, are plugged into the AC mains for continuous power, and include high performance CPUs and DRAM. Client devices are designed for great cameras and screens, but do not have the CPU processing power to use mulitple streams and do not have space for optimal antenna design. Thus, MU-MIMO allows the access point to transmit concurrently to more than one client at a time. In 802.11ac, MU-MIMO is limited to downstream packet only, thus it is often written as DL MU-MIMO.

Speed = Capacity

A typical 802.11ac client device will support up to 80 MHz channel width and two spatial streams. This could be a tablet, or any top tier popular smart phone. With this configuration, the RF PHY rate will be 867 Mbps. A faster RF PHY rate leads to more contention-free time for other devices to access the RF medium, lower RF noise, and higher capacity.

The decade of the 2020s will bring widespread use of IOT devices in every home and enterprise and more personal devices. Additionally, 8k video will become commonplace and streaming media bitrates will rise again.

A new generation of Wi-Fi technology will be needed.

Higher channel efficiency

802.11ax is about efficient use of RF spectrum to increase network capacity

Whereas 802.11ac made the network faster, 802.11ax is designed for higher channel efficiency in a dense networking environment. The key technology behind 802.11ax, multiple access OFDMA, comes from the 3GPP LTE cellular technology and 802.16e WiMax. OFDMA, orthogonal frequency division multiple access, is a field proven technology to support higher density wireless networks.

802.11ax anticipates higher density in several ways; scheduled transmission time, dynamic frequency allocation in 2 MHz increments, spatial reuse and improved battery management mechanisms.

OFDMA enables efficient spectrum use

802.11a, 802.11g, 802.11n and 802.11ac use OFDM (orthogonal frequency division multiplexing). OFDM was a significant enhancement over 802.11b DSSS. OFDM splits the signal over multiple narrow band carriers. This results in reduced crosstalk interference and allowed higher RF rates.

In an 802.11ac network with OFDM, the smallest frequency unit assigned to a transmission is 20 Mhz. The entire 20 Mhz frequency is used to transmit to and from any single client device. If that device is sending a small 64 Byte IOT packet or a 1248 Byte video packet, it uses the same 20 Mhz. (the author acknowledges that narrowband industrial and IOT applications using 802.11 can be purpose built to use 5 MHz or 10 Mhz frequencies. This is not equivalent to OFDMA sub-channels).

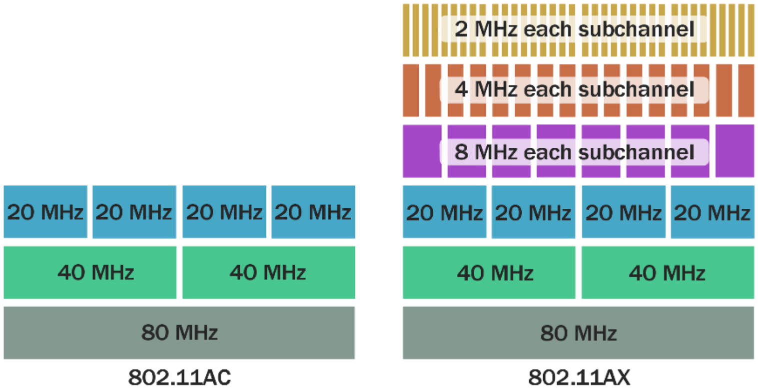

OFDMA divides up the 20 MHz channel into 256 subcarriers. The smallest unit is 26 subcarriers in 2 Mhz of frequency. These subcarriers can be grouped together into larger units; at 52 (4 Mhz), 106 (8 Mhz), and 242 (20 Mhz) to support higher bandwidth applications. Each grouping of subcarriers are called a resource unit (RU).

Any single client device is assigned one or more RU. Each RU can be transmitted at a unique QAM level. Additionally, as the client device moves farther away from the access point, the number of RU and the QAM level can be adjusted to maintain a strong signal level even as throughput degrades. Dynamic RU allocation is a direct benefit of OFDMA use in LTE networks.

By splitting the transmissions into smaller frequency units, and dynamically assigning RUs, the network can use the entire 20 Mhz wide channel to support multiple devices running at different speeds and with different throughput requirements.

Consider an enterprise hotel with IOT sensors and conference attendees in the ball room. IOT sensors can be allocated a single 2 Mhz RU to send sensor updates to a server, while attendees accessing the Internet can be assigned a 8 Mhz RU to support a high capacity and high density wireless network. The full 20 MHz channel will be utilized at each transmission opportunity.

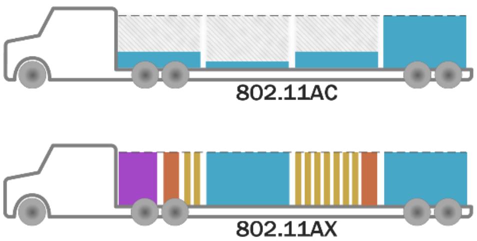

Thus far, the discussion has been limited to a 20 Mhz wide channel. 802.11ax uses OFDMA to subdivide 40, 80, and 160 Mhz wide channels into an incresingly dense number of subcarriers. Additionally, the access point will synchronize upstream OFDMA tramissions from the client to the access point. Thus OFDMA allows simultaneous, multiple client tranmissions in each direction.

MU-MIMO. It's back

Well, it never really went away. In 802.11ac wave 2 access points, DL MU-MIMO provides some network throughput improvement by scheduling transmissions to multiple clients at the same time. Refer to the Appendix for a longer discussion of DL MU-MIMO in 802.11ac. This technology is particularly well suited to support mulitple high bitrate clients such as video streaming that use a large frame size. As the 802.11ax technology evolves, MU-MIMO will work in both downstream and upstream directions. Another key advantage with 802.11ax is support for eight concurrent MU-MIMO transmissions concurrently, double the number supported by 802.11ac. One of the limiting factors for MU-MIMO in 802.11ac is the high management overhead required to communicate subcarrier information to the access point. In 802.11ax, enhancements to MU-MIMO wil include grouping multiple client updates together to improve protocol efficiency and remove excess management overhead. The capacity and high throughput advantages make MU-MIMO another key technology for 802.11ax. Note that MU-MIMO and MU-OFDMA are inherently incompatible. Thus the access point will choose when and how to use each technology for maximum performance and efficiency.

Spatial Reuse

At the outset, we talked about technology changes and use cases that are adding increased number of devices consuming more content every year. For a network architect, this is the classic definition of a "high density" network. The traditional way to design this network is to use 20 Mhz wide channels in the 5 GHz band, and to limit the number of overlapping wireless cells to prevent excessive co-channel and adjacent channel interference. That works, up to a point.

All 802.11 technologies are "listen before talk". This means they will measure the signal strength of other 802.11 frames on the wireless medium, and defer transmission if the signal strength is above a certain threshold. On 5 GHz networks, the theshold is -82dBm RSSI. This process is called Clear Channel Assessment (CCA), and it does just like it sounds: it determines if the channel is clear for transmission.

As the network density increases, there is an increased chance that any two devices cannot hear each other, but each can hear a third device. If the two devices determine that the medium is "clear", their transmissions may collide and an error occur. These two devices are called "hidden nodes" since they cannot hear each other and defer transmission. In 802.11 networks, this hidden node issue is managed by using CTS (clear to send) and RTS (request to send) packets to ensure that only one device has the transmission rights for a given time slot. RTS/CTS works well at reducing medium access contention, but slows down the network throughput and capacity by adding extra management overhead to the network. Hardly a good solution to support increased numbers of wireless devices.

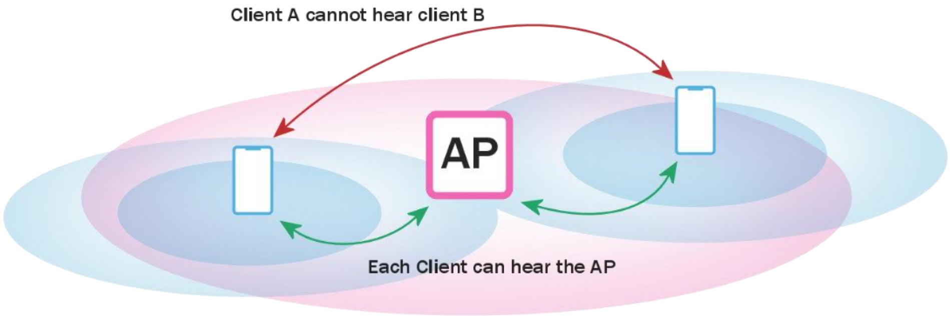

In 802.11ax, we now have a mechanism to mark the preamble of the wireless frame with a unique "color", called BSS Coloring, based on the BSS ID of the wireless cell. Using this technique, multiple networks can overlap. A single BSS is essentially one 802.11 access point and its connected wireless clients. Each wireless client transmitting a wireless frame will include the unique identifier in the wireless preamble. The receiving device, whether an access point or another 802.11ax client, will compare the bits received in the preamble against its own BSS color bits. If the bits matches, then the packet belongs to the same BSS and medium access proceeds according to established rules. If the bits do not match, the transmitting device can ignore the frame, discard it, and medium access is not affected.

Consider how this technology can dramatically increase throughput in a high density network like an education campus. When paired with load balanced band steering, wireless clients can be spread amongst multiple overlapping access points for ultra high density along with high throughput.

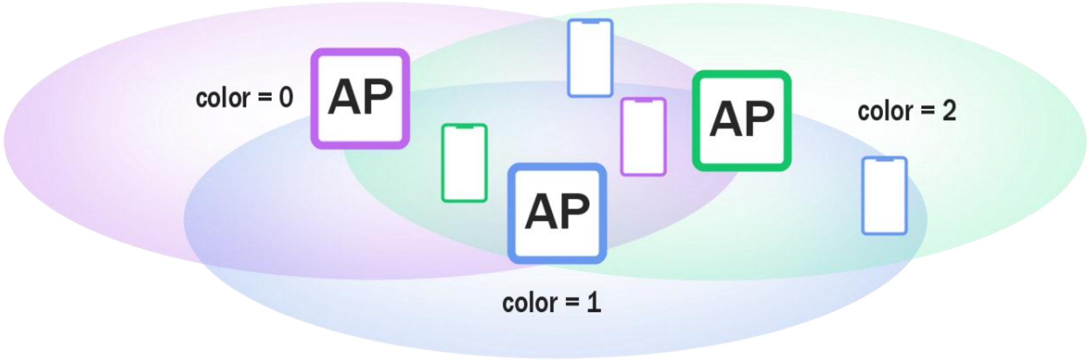

802.11ax supports 2.4 GHz

With 802.11ac, the 2.4 GHz band seemed like a long forgotton sock lost in the wash. We kinda know it's there, but we don't really want to spend the time to try and use it. 802.11ax will include support for the 2.4 GHz band and allow legacy industrial, medical and IOT devices to continue operation, while also extending OFDMA advantages to newer 2.4 GHz devices.

Other New Toys in 802.11ax

The new standard will offer 1024 QAM to increase the number of bits transmitted in each symbol. 1024 QAM will provide a 25% speed improvement over 802.11ac. However, 1024 QAM will not work with small RU of 2, 4 or 8 MHz. 1024 QAM can only be used with a full 20 MHz channel. Additionally, the higher bit density will require greater single to noise ratio, between 35dB and 40dB, to achieve the higher speeds. If you assume a noise floor of -90dB, that means the clients will have to receive the access point signal at a level of -55dB or lower. Designing a network at -50dB RSSI will require more access points – fortuneately 802.11ax includes BSS Coloring to facilitate higher density and overlapping BSSes.

With unabated growth in smart phones and IOT devices, 802.11ax will include features designed to preserve battery life. Target Wait Time (TWT) mechanism controls when and for how long, a client device can be in sleep mode before it wakes up to receive packets from the access point. Combined with uplink transmission scheduling; TWT enables longer sleep times and coordinated and deterministic communications to reduce network errors and preserve battery life.

When is the network upgrade needed?

When deciding when to consider a network upgrade; service providers need to calculate the cost/benefit of upgrading to 802.11ax. should you be an early adopter? Or wait until more ecosystem support is available? What applications and typical devices are in my network?

802.11ax standard will be ratified November 2020

Consider:

- Client and ecosystem support including standards ratification and interoperability support
- Density growth trends
- Backward compatibility with legacy devices and systems
- Available RF spectrum within a crowded RF neighborhood
- State of IOT applications if applicable
- Equipment cost

Consumer phone upgrade cycle has slowed in recent years to an average of 23 months. As a general rule of thumb, once the top two smart phone vendors introduce 802.11ax radios into their phones, assume that the network will be ready for an upgrade 12 months later. Samsung S10, iPhone 11 and iPhone SE all support 802.11ax, in addition to other brands and models.

Wrap Up

802.11ax focuses on metrics that will improve the client experience and increase cell throughput. Network capacity is expected to increase by a factor of 4x vs existing 802.11ac networks.

Cambium Networks radio and RF engineering is working with chipset vendors to develop 802.11ax access points that will integrate with the Cambium Wireless Fabric and enhance network visibility, control, security and more.

Talk with your Cambium Networks resellor or distributor to get on the Cambium notification list for 802.11ax product development updates.

Cambium Networks Wireless Fabric

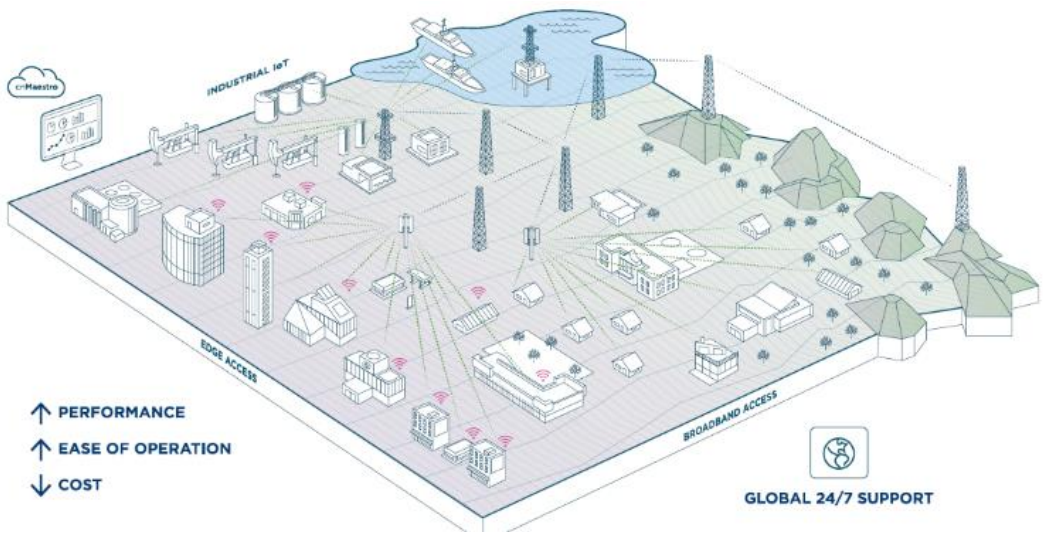

Appendix

The objective of the 802.11ac standard is to Increase network throughput. Make it faster. 802.11ac has seven features designed to increase the speed and capacity of the wireless network. A short review of these features will give us insights into where we can improve.

Feature	Objective	Relative Value over 11n
VHT Frames	VHT frames added data packet aggregation to increase the amount of data transmitted before an ackknowledgement is required. VHT frames also Include additional information in the header to enable other 802.11ac features discussed below.	
256 Modulation	802.11n uses a 64 QAM modulation. Basically, it describes how dense the data bits can be packed into each symbol. 802.11ac includes 256 QAM modulation. This increases the RF rate by 33%.	
Airtime Fairness	For any 802.11 wireless network, only one device can transmit at a time. Consider that an 802.11n client running at 72 Mbps RF rate on a 20 Mhz channel will consume much more time than a new 802.11ac client running at 433 Mbps on the same AP if both clients are streaming the same movie. Airtime fairness ensures the legacy device does not consume all the transmission time. Airtime fairness results in high client density and higher network capacity.	
Explicit TxBF	Explicit transmit beamforming is a standardized way to effectively increase the RF power to a single device. This benefits the network throughput by isolating the signal in a particular direction, and standardization ensures that clients can work with any access point that supports the standard. The term “explicit” describes the AP using client derived sub-channel information to create a stearing matrix unique to that client. Explicit TxBF is required to make MU-MIMO functional. No form of beamforming can project a pencil thin “beam of RF energy” to a generic wireless client. Thus TxBF works best when there is physical seperation between different clients connected to the same AP. Since the client must send sub-channel informatio to the AP before beamforming can occur, TxBF works best when the client device is stationary.	
Channel Bonding	802.11ac requires an AP to bond two or four 20 Mhz channels. Four channels * 20 Mhz wide will produce a single 80 MHz wide channel for 2x speed over 802.11n. The spec envisions bonding up to eight channels for a single 160 Mhz wide channel, or bonding two sets of four channels for 80 MHz + 80 MHz. When 160 MHz of frequencies are used for a single transmission, interference between access points is a significant problem. Depending on the regulatory domain, 160 MHz can provide from one to three usable channels in the 5 GHz	

	UNII band. This is just not enough to allow for multiple access points in the same enterprise. When designing a network for high density; such as an education or stadiums, network architects will always default to use multiple 20 MHz channels to minimize interference and maximize the number of client devices that can connect error-free. Particularly in enterprise, education and hospitality markets; channel bonding is not a significant improvement over 802.11n.	
MIMO streams	802.11n introduced the ability to transmit and receive multiple data streams simultaneously. When the streams are combined, the effective data rates are multiplied. With 802.11ac; two, four and eight MIMO streams can be multiplied. Each MIMO stream requires one antenna element. Thus, to transmit eight MIMO streams will require eight antenna elements. Smart phones, tablets, and even laptops lack the necessary space and power budget to accommodate eight antenna elements. The best price-performance trade off between the number of antenna elements, battery life, CPU load and network performance is two. Unfortunately, 802.11n also supported two MIMO streams so this feature has limited value relative to 802.11n.	
MU-MIMO	Multi-user MIMO is designed to accommodate the “gap” between the number of antenans possible in the access point and the number of antennas typically found on the client device. As discussed above, most client devices will have one or two antenna elements. Since the AP can have four or eight antenna elements, why not find a way to utilize the extra streams? Enter downlink multi-user MIMO, or DL MU-MIMO. To make it work, the access point queries the client device using a sounding protocol to obtain sub-channel information. Sub-channel information can also be sent by the client in unsolicited management frames. Using this information, the access point does three things: 1. Calculates a steering matrix that will be used for TxBF 2. Calculates a unique pre-coding to pre-pend to the client data frames 3. Allocates the client to a MU-MIMO group for the next available transmission opportunity (TXOP) So far, so good. MU-MIMO sounds like a great solution to leverage the extra antennas that can be built into an access point. In practice, MU-MIMO is not as effective as we would like. There are three reasons: 1. Extra signal processing consumes CPU cycles and reduces battery life. For smart phones, this one is a deal breaker. Few smart phone handset vendors support MU-MIMO at all. The most popular brand,	

Apple, does not support MU-MIMO in its 802.11ac phones and tablets.

2. MU-MIMO relies on transmit beamforming to carry multiple transmissions to multiple clients that support MU-MIMO. In many high density environments, there isn't enough physical space between each device to make a difference. Imagine a classroom with laptops side by side every 50cm (20in) Beamforming is not a laser beam, it's more of a suggestion.
3. Mobiles devices are mobile. When the client sends sub-channel information back to the AP, the data is valid for that location. The AP cannot accurate calculate a steering matrix for a device in motion. The extra management overhead to support the sounding frames for MU-MIMO are not offset by the benefit, particularly for a device in motion.

Many network architects find that a properly designed MIMO network will outperform a poorly supported MU-MIMO network. This feature, in 802.11ac, is disappointing.

Glossary

IOT	The Internet of Things. As used in this paper, this term describes the wireless devices and applications used in consumer homes, industrial sites, medical, hotels, and enterprise networks to automate tasks, monitor, measure, and report status, temperature, location, etc. IOT devices often use proprietary medium access control based on BLE, 802.15.4, or standardized protocols. With 802.11ax low latency, small subcarriers, and extended battery life – we may see IOT devices adopt 802.11ax transceivers soon.
MIMO	Multiple Input, Multiple Output. First introduced with 802.11n, MIMO uses multiple antenna elements to transmit and receive RF signals. Since indoor RF signals are subject to multipath reflections, MIMO leverages multipath propagation to increase signal gain and multiple the effective datarate. The amount of throughput increase is limted by the receiving device. For example, if an access point has four antennas and a client has one, there is no MIMO benefit to the client. Since each antenna element requires operating power, processing power and physical space; it logically follows that access points can have many more antennas than small battery powered client devices.
MU-MIMO	As noted above, access points can more easily support more antenna elements than a battery powered, low CPU speed smart phone. Multi-user MIMO is a techincal introduced in IEEE 802.11ac that allows an access point to transmit distinguishable signals to more than one client at a time.
Wave 2	AKA IEEE 802.11ac phase two. The term "wave 2" is not part of the standard, but is an industry agreed upon term used to describe capabilities delivered in the second phase of 802.11ac radio development. Generally speaking, the industry definition of wave 2 is support for MU-MIMO.
TWT	Target Wait Time. A mechanism in 802.11ax to control how long, and when a client radio will go into power save mode. TWT preserves battery life.

http://www.ni.com/white-paper/53150/en/