12

802.11ax for 5G

Richard van Nee

Qualcomm, the Netherlands

12.1 Introduction

It is expected that 5G is made up of several radio access technologies as there is probably no single technology capable of meeting all requirements. One likely candidate for being part of 5G is the new 802.11ax standard. This wireless LAN standard offers data rates up to 10 Gbps in channels up to 160 MHz. It improves capacity for large number of users by employing OFDMA which reduces overhead relative to the TDD approach of existing wireless LAN standards. At the same time, the use of uplink OFDMA provides a significant range benefit for clients that have less transmit power than access points, a situation that is pretty common for handheld devices. Further capacity enhancements are provided by the use of uplink MU-MIMO in addition to downlink MU-MIMO that was already present in the 802.11ac standard.

Peak throughputs of wireless LAN have grown exponentially for quite some time as shown in Figure 12.1. From 2 Mbps for early 802.11 products in the nineties, peak rates have steadily climbed to several Gbps for current 802.11ac products and expect to reach approximately 10 Gbps with 802.11ax [1]. While early 802.11 standards primarily focused on single link rates, the attention shifted towards network throughput since 802.11ac with the introduction of MU-MIMO [2]. Improving network throughput in dense networks is the primary goal of 802.11ax. This is important to deal with increased levels of inter-network interference and to provide more capacity for networks with large numbers of clients, for instance in cases where WiFi is used for data offload to relieve the pressure on LTE spectrum. The latter use case is one of the main drivers to make WiFi part of 5G.

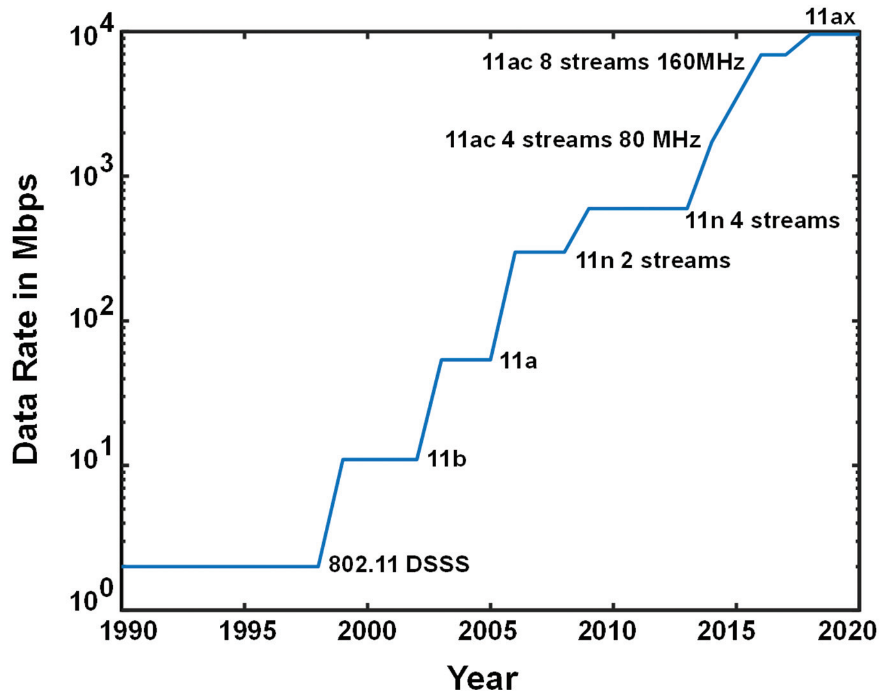

Figure 12.1 WLAN data rate growth.

12.2 802.11ax Features

The main new features in 802.11ax are:

1. Longer symbol durations with 4 times smaller subcarrier spacing
2. OFDMA
3. Uplink MU-MIMO
4. 1024-QAM
5. Range extension with DCM
6. Dynamic CCA

Longer symbols with 4 times smaller subcarrier spacing are used for several reasons. First, it provides a larger data rate by reducing the relative duration of the guard interval and by a slight increase of the band filling. Second, it allows the use of larger guard intervals which provides more delay spread robustness. Guard intervals of 800 ns, 1600 ns and 3200 ns are defined in 802.11ax together with a 12.8 microseconds FFT interval whereas 802.11ac uses a guard interval of 800 ns or 400 ns with a 3.2 microseconds FFT interval.

The additional delay spread robustness is especially useful for outdoor use where RMS delay spreads in the order of a microsecond may occur. A third reason for using longer symbol durations is that it provides a more relaxed synchronization requirement for uplink OFDMA and uplink MU-MIMO.

To enable channel estimation, an 11ax preamble contains a number of HE-LTF symbols equal to or larger than the total number of spatial steams, similar to 11ac. One disadvantage of using larger symbol durations is that the overhead of this channel estimation would be 4 times larger than for 11ac if all HE-LTF symbols used the longest symbol interval of 16 microseconds. To reduce this overhead, 11ax specifies 4 different modes of operation with symbol durations of 4 μs (including 800 ns guard interval), 7.2 μs (including 800 ns guard interval), 8 μs (including 1.6 μs guard interval), and 16 μs (including 3.2 μs guard interval). All modes beside the longest 16 μs require tone interpolation to get the channel estimates for all data tones.

For non-OFDMA mode, 11ax uses 242 tones in 20 MHz including 8 pilots, giving a peak data rate of 1147 Mbps for 8 streams, 1024-QAM, coding rate 5/6, with 800 ns guard interval. For 40 MHz, 484 tones are used including 16 pilots, giving a peak rate of 2294 Mbps. 80 MHz uses 996 tones with 16 pilots, which are also used by 160 or $80 + 80$ mode in each 80 MHz segment, giving a top rate for $80 + 80$ of 9607.8 Mbps for 8 spatial streams and 1201 Mbps for a single spatial stream.

12.3 Interoperability and Mode Detection

The new 802.11ax standard will be used both in the 2.4 GHz and 5 GHz band. Devices that implement 802.11ax have to be interoperable with existing 802.11 products which mean they need to support all existing 802.11 waveforms for either the 2.4 or 5 GHz band. At the same time, the new 802.11ax waveforms should not cause performance degradation for existing 802.11 devices that did not implement 802.11ax. To achieve the latter, all 802.11ax waveforms start with a 802.11a-like preamble, similar to 802.11n and 802.11ac. The signaled rate in the 802.11a preamble is set to 6 Mbps and the length field is set to match the duration of the 802.11ax packet. This ensures that all devices that can receive the 802.11a part will properly defer for the entire packet even though those devices may not be able to receive any 802.11ax data.

The existing preamble structures for 802.11a, 11n, and 11ac are depicted in Figure 12.2. It is important that a device can quickly and reliably detect what preamble type is being received. For 11a, 11an and 11ac, this detection

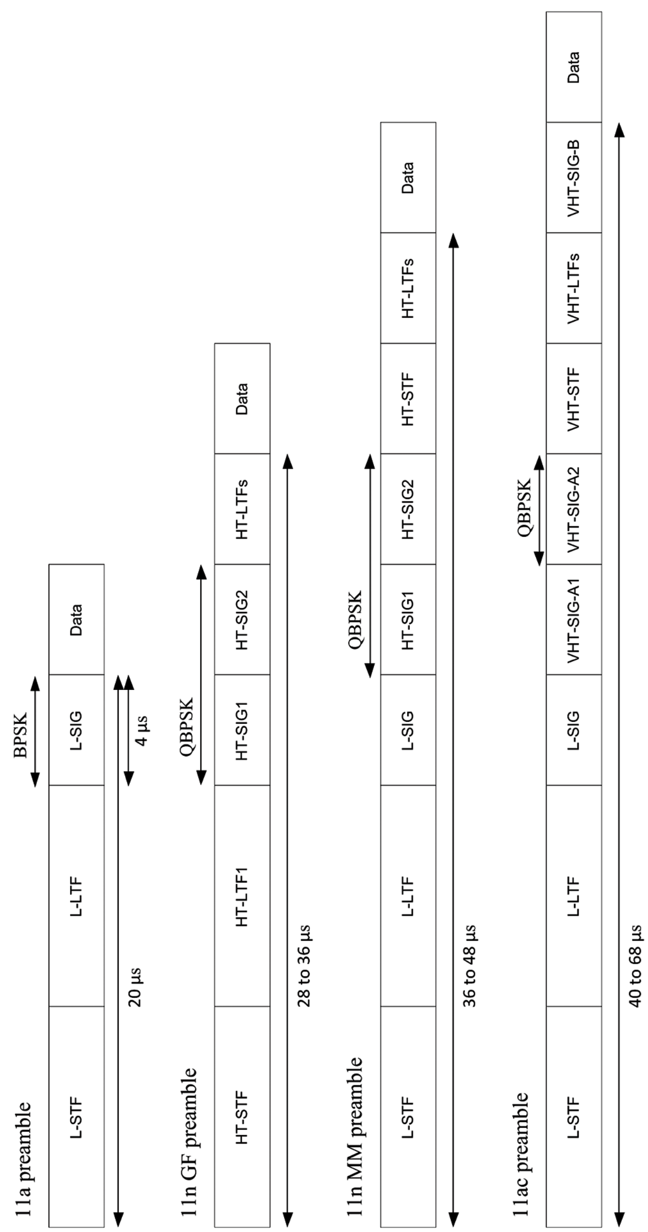

Figure 12.2 Preamble structures for 11a, 11n and 11ac.

is based on the use of QBPSK at different parts of the preamble. QBPSK is equal to BPSK rotated by 90 degrees. An 11n greenfield preamble uses QBPSK at the first symbol after the LTF symbol. An 11a packet has a BPSK L-SIG symbol at this same location. Hence, a receiver can detect whether a packet is 11a or 11n-GF by estimating whether the constellation of this first symbol after LTF is BPSK or QBPSK. An 11n mixed-mode (MM) preamble has a BPSK L-SIG followed by two QBPSK HT-SIG symbols. The L-SIG symbol is a valid 11a-type signal field indicating the use of 6 Mps BPSK and with a length field that covers the entire 11n-MM packet. After decoding this L-SIG symbol, a receiver does not know yet whether the packet is 11a or 11n-MM, it only knows it cannot be 11n-GF based on the QBPSK check on L-SIG indicating BPSK. By doing a second QBPSK check on the symbol after L-SIG, it can detect whether the packet is 11n-MM versus 11a. For 11ac packets, a third QBPSK check is required at the second symbol after L-SIG. The packet is detected as 11ac based on the L-SIG being BPSK and signaling 6 Mbps, the first symbol after L-SIG also being BPSK, and the second symbol after L-SIG being QBPSK. Devices that did not implement 11ac will detect the 11ac packet as 11a based on the valid L-SIG symbol and the absence of QBPSK in the first two symbols after the LTF symbol.

Figure 12.3 shows the additional preambles defined by 11ax. There are 4 new preambles for Single User (SU) mode, Multi User (MU) modes being downlink MU-MIMO and downlink OFDMA, Trigger modes being uplink MU-MIMO and uplink OFDMA, and extended range mode. For existing devices that did not implement 11ax, all 11ax packet types are detected as 11a such that a proper defer will be done based on the L-SIG length field that is set to cover the entire duration of the 11ax packet. An 11ax device can discriminate 11ax versus 11n and 11ac based on the fact that the L-SIG length field modulo 3 is 1 or 2 for 11ax versus 0 for 11n and 11ac. The L-SIG length field indicates a number of bytes, but 11n, 11ac and 11ax only use this length field to indicate the total duration of the packet and not the byte length. Since the spoofed 11a rate of 6 Mbps gives 3 bytes per symbol, the same packet duration can be signaled by 3 different L-SIG length values with 1, 2, or 3 bytes in the last symbol. If the L-SIG length modulo 3 is 1 or 2, then a receiver knows that the packet cannot be 11n or 11ac, but it does not know yet whether it is 11a or 11ax. To detect this, there are 2 different checks that can be made. First, the L-SIG is repeated in 11a packets. This repetition can be used for mode detection. At the same time, it provides a range advantage in the extended range 11ax mode where all fields use half the lowest 11a rate. Second, the L-SIG, repeated L-SIG, HE-SIG-A,

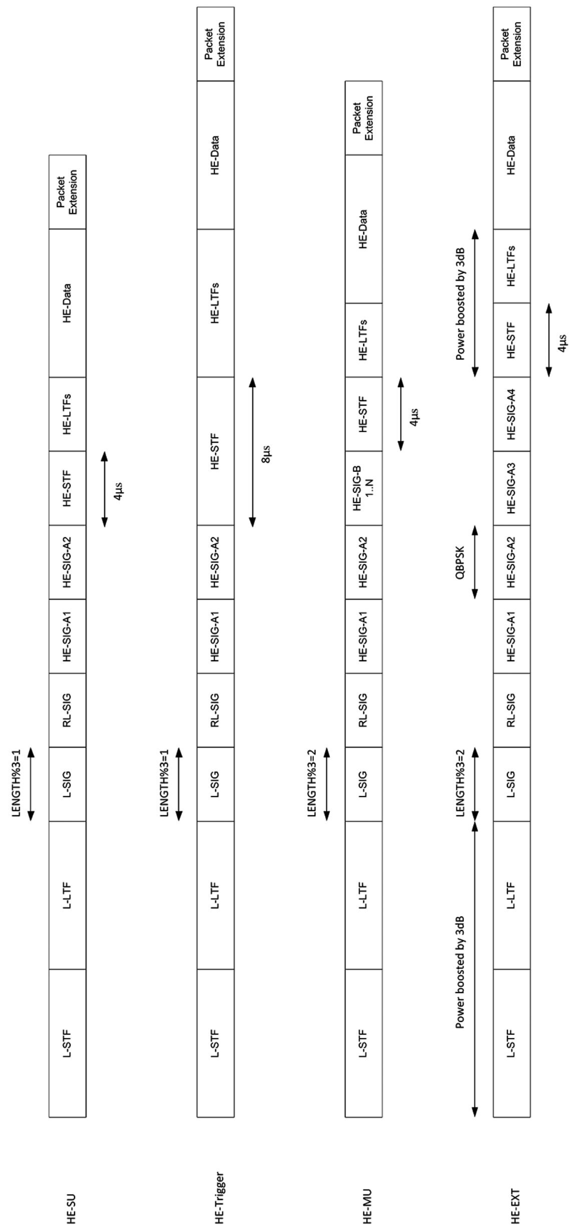

Figure 12.3 Preamble structures for 11ax.

and HE-SIG-B symbols all have 4 edge tones more than 11a symbols. These extra tones contain known BPSK values and they can be used both for mode detection as well as for improving channel estimation on the HE-data.

Once a receiver detected that a packet is 11ax and finds out the specific 11ax packet type, it starts to decode signaling data in HE-SIG-A. This field contains data that is needed to decode the rest of the packet such as the MCS (BPSK up to 1024-QAM, coding rate ½ up to 5/6), number of spatial streams, guard interval and LTF size, bandwidth, SU/MU, STBC, beamforming, and several other parameters. For downlink MU packets, there is a lot more signaling as every downlink client can have a different MCS, number of streams, and coding type. To accommodate this extra signaling, MU packets have a variable number of HE-SIG-B symbols. The length and MCS of the HE-SIG-B field is indicated in HE-SIG-A.

12.4 OFDMA

Downlink OFDMA is used in 11ax to minimize the overhead of PHY preambles and MAC backoff when there are a significant number of clients and relatively small packets. The overhead reduction is obtained by combining multiple small packets into one large packet. Without OFDMA, large PHY rates tend to give a high efficiency loss when the packet size per client is relatively small as preamble and backoff times become much larger than the data portion of each packet. The smallest 11ax preamble is 40 μs, so a packet should preferably be at least 400 μs to make the preamble overhead no more than 10%. At a PHY rate of 10 Gbps, however, that means a single packet contains about 500 KB. For a single client, this requires a significant amount of aggregation, which may not be possible if there are latency constraints on parts of the traffic for that client. OFDMA helps in this case as the aggregation can be done over a group of clients.

Figure 12.4 shows the resource unit (RU) allocation for 80 MHz OFDMA. There can be up to 37 simultaneous clients with 26 tones, or 16 clients with

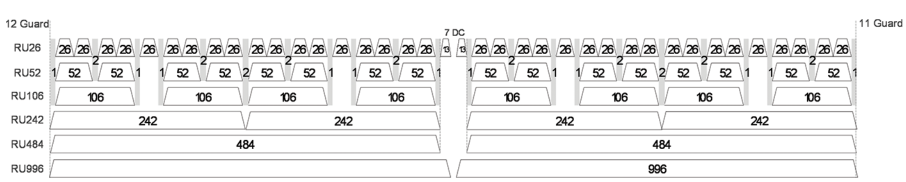

Figure 12.4 80 MHz OFDMA resource unit allocations.

52 tones, 8 clients with 106 tones, 4 clients with 242 tones, 2 clients with 484 tones, or a single client with 996 tones. The same RU sizes are used in other bandwidth modes. In 20 MHz mode, there can be up to 9 simultaneous clients with 26 tones each.

Figure 12.5 shows the packet strucure including signaling per client for an 80 MHz downlink OFDMA packet. The first part of the preamble up to HE-SIG-A is identical for all clients and is repeated in each 20 MHz subchannel to make sure that overlapping networks are able to decode the length of the packet that is encoded in the L-SIG symbol. The HE-SIG-B field is present only in MU packets, either MU-MIMO or OFDMA. It consists of a signaling block that is common to all MU clients and a per-client signaling part. The common part has a separate CRC and tail bits. Each 20 MHz is encoded separately such that a client only needs to decode one 20 MHz subchannel. The per client part has a CRC and tail bits for each group of 2 clients, such that a client only needs to decode a small part of the entire HE-SIG-B correctly in order to decode his signaling info. In case of an odd number of clients, the last group just has one client followed by CRC and tail bits.

Uplink OFDMA is specified in 11ax for the same reason as downlink OFDMA to reduce both PHY and MAC overhead. There is an additional power accumulation gain benefit though; if all uplink clients transmit at the same maximum power in a 26-tone RU than in a full 20 MHz RU in non-OFDMA mode, then the maximum received power at the AP is almost 10 dB more if 9 uplink clients transmit simultaneously in a 20 MHz channel relatively to a single client. For the clients, this means they can transmit at a significantly higher data rate relatively to non-OFDMA mode. This provides a network throughput gain on top of the gain caused by reduced PHY and MAC overhead. Notice that in practice, the power accumulation gain may be limited because of regulatory constraints.

One major difference between uplink and downlink OFDMA is that uplink OFDMA requires time and frequency synchronization of uplink clients, as well as some level of power control. The same is true for uplink MU-MIMO. In an uplink frame, the packets from different clients should arrive at the AP with timing differences that are small relative to the guard interval. To achieve this, clients are required to have an accurate SIFS response time with an error not exceeding $+/-400$ ns. The clients are allowed to transmit uplink frames only in response to a trigger frame by the AP. They have to synchronize their carrier frequency and sampling clock to that of the AP. The accuracy of the frequency synchronization has to better than 1% of a subcarrier spacing to prevent a significant degradation caused by inter-carrier interference.

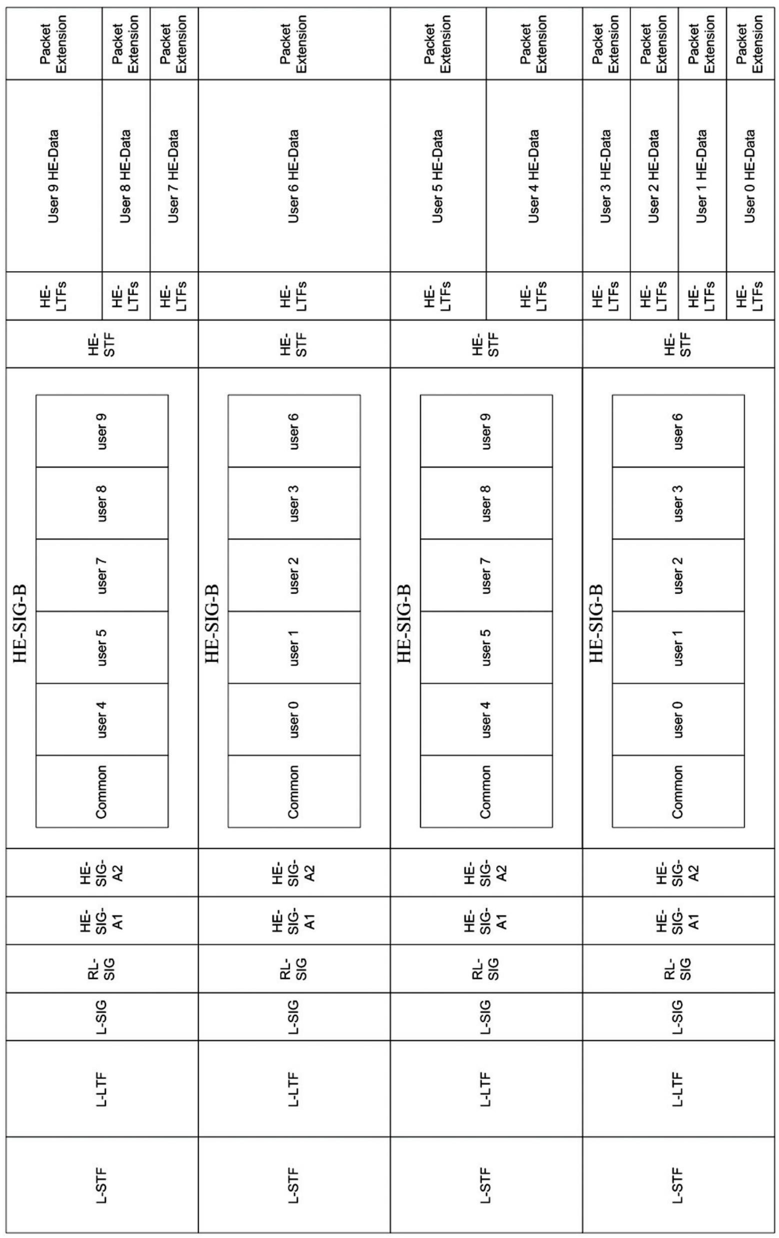

Figure 12.5 11ax packet structure for 80 MHz OFDMA.

Since the subcarrier spacing is 78.125 kHz, the accuracy requirement is about +/−780 Hz. Uplink MU-MIMO and uplink OFDMA clients also have to perform power control with an accuracy of +/− 3 dB. To enable power control, the AP signals its transmit power and a target received power level per client in the trigger frame. An uplink client can estimate the path loss from the received power level of the trigger frame and the signaled AP transmit power. By assuming reciprocity, it can then set its transmit power to achieve the desired received power at the AP as requested in the trigger frame.

12.5 Uplink MU-MIMO

In 802.11ac, downlink MU-MIMO was introduced as a way to increase the throughput in cases where the AP had more antennas than the clients. Since then, there has been a demand to increase uplink throughput as well, as the amount of uplink traffic has significantly increased over the past years with smartphone users producing large amounts of video data that is shared across social media. The downlink MU-MIMO throughput advantage is now extended to the uplink in 802.11ax by introducing uplink MU-MIMO. Without MU-MIMO, clients have to compete for the medium and transmit packets sequentially like depicted in Figure 12.6. With uplink MU-MIMO, clients can transmit simultaneously like shown in Figure 12.7, increasing the uplink throughput up to a factor of 4 for 4 uplink clients.

In addition to the throughput multiplication factor, uplink MU-MIMO has the same power accumulation benefit as uplink OFDMA. By having 4 simultaneous uplink MU clients, the received power at the AP is 6 dB more relative to having a single client. This means each uplink MU client can transmit at a higher MCS than in SU mode. Another benefit of uplink MU-MIMO that it benefits downlink MU-MIMO. Downlink MU-MIMO according to 11ac requires compressed channel feedback per client to be transmitted sequentially in the uplink as depicted in Figure 12.8, creating a significant amount of overhead. In 11ax, it is possible – but not required – to transmit the compressed channel feedback in an uplink MU-MIMO mode as shown in

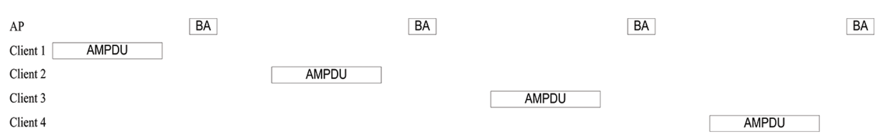

Figure 12.6 Multiple clients sending a packet without use of uplink MU.

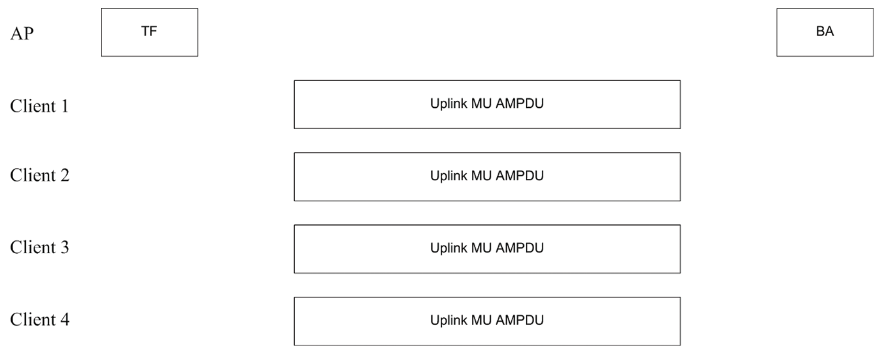

Figure 12.7 Multiple clients sending a packet using 11ax uplink MU.

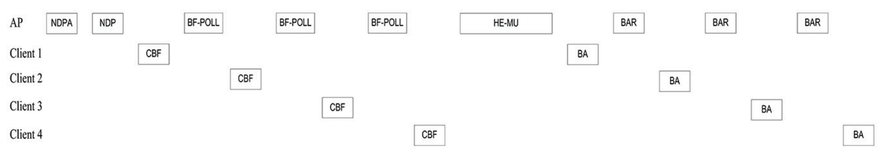

Figure 12.8 Downlink MU TXOP for 11ac.

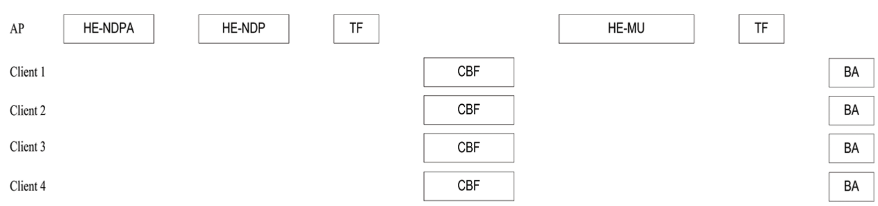

Figure 12.9 Downlink and uplink MU for 11ax.

Figure 12.9. This reduces the overhead and hence improves the throughput of downlink MU-MIMO.

12.6 Range Extension

If current WiFi clients are at the edge of the AP coverage range, they often suffer from the problem that they while they can hear the AP, they cannot associate as the AP cannot hear them. This uplink problem is caused by the fact that the AP generally has more transmit power and hence more range than typical client products. To solve this issue, 11ax defined a range extension mode with a special preamble shown in Figure 12.3. The STF and LTF fields are boosted by 3 dB for the extend range preamble to make sure that the extended range performance is not limited by preamble detection. The preamble signaling symbols are repeated to provide 3 dB of gain in the signaling part. For the data part, Dual Carrier Modulation (DCM) can be used to provide a lowest data rate that is half of the lowest BPSK rate without DCM. DCM basically duplicates half of the tones. To reduce the peak-to-average power ratio caused by the duplication of tone values, for BPSK the odd copied tones are inverted while for QPSK and 16-QAM, the copied tones are conjugated. DCM provides lower rates with twice the diversity order and 3 dB gain relative to no DCM with the same constellation size.

12.7 Dynamic CCA

The current 802.11 standard specifies a few power levels for which any device should defer. There are different power levels depending on bandwidth with higher defer level for larger bandwidth modes. Also, there is a separate level for the case that a valid preamble is detected versus a larger energy detect threshold for cases where the preamble had not been detected, which may happen for instance because the device was receiving some other packet that partially overlapped with another transmission. There are two mechanisms in 802.11ax to improve the defer behavior with the goal of increasing network capacity in dense environments. First, the HE-SIG-A field includes a BSS (Basic Service Set) color which is basically a shortened version of the AP address. After decoding the BSS color, a station knows whether the packet is from its own network or from some overlapping network. If the packet belongs to the same network, it is better to always defer regardless of received power. There is no advantage trying to transmit 2 simultaneous packets in this case because the AP can only receive one packet at a time. It may also be that the received packet is from the AP itself, in which case the AP is not able to receive until it is finished with its transmission. If the BSS color indicates that the packet is from another network, then it may be possible to transmit on top of it. Rather than using a fixed CCA threshold like in current 802.11, 802.11ax allows the use of a dynamic defer threshold where the threshold depends on the transmit power. The lower the transmit power, the higher the defer threshold. The reasoning is that a station with lower transmit power causes interference over a smaller range, hence it can increase its defer threshold such that it defers only for relatively nearby other stations.

12.8 Conclusions

The 802.11ax standard increases the maximum single user and multi-user data rate to about 10 Gbps. Several enhancements are made to increase the throughput in dense networks with many clients, such as the introduction of OFDMA, uplink MU-MIMO, and dynamic CCA. These improvements make 802.11ax attractive for data offloading of crowded LTE networks and also make it a suitable candidate for inclusion in 5G.

References

[1] IEEE P802.11ax/D0.1 draft standard, March 2016.
[2] IEEE Std 802.11ac-2013, amendment to IEEE Std 802.11-2012.

About the Author

Richard van Nee received the M.Sc. degree in Electrical Engineering from Twente University in Enschede, the Netherlands, in 1990, followed by a Ph.D. degree from Delft University of Technology in 1995 (both cum laude). From 1995 to 2000, he worked for Lucent Technologies Bell Labs. In 1996, he developed an OFDM based packet transmission system for the 'Magic WAND' project, which was the first demonstration of OFDM for wireless LAN. Together with NTT, he made an OFDM based proposal that got adopted by the IEEE 802.11a standard. He was also co-author of the CCK-code proposal that got adopted by the IEEE 802.11b standard. The 802.11b CCK codes were originally published in a 1996 paper by Richard van Nee as a way to reduce the peak power in OFDM, but it also turned out to be a useful code set to boost the WLAN data rates to 11 Mbps in 802.11b, while keeping good delay spread robustness and conform to the FCC spreading rules that were still in place at that time. The 802.11b standard brought WLAN data rates on par with existing wired LAN at the time, which contributed to its rapid market adoption and to the formation of the WiFi alliance. In 2001, Richard van Nee co-founded Airgo Networks – acquired by Qualcomm in 2006 – that developed the first MIMO-OFDM modem for wireless LAN and which techniques formed the basis of the IEEE 802.11n standard. He was responsible for the algorithm design of the groundbreaking Airgo technology with many innovations like MIMO-OFDM preamble designs and near-maximum likelihood MIMO decoding techniques. In 2012, he presented the first spec framework of the 802.11ac standard on behalf of a large group of authors. This was the first standard to use MU-MIMO. Together with Ramjee Prasad, he wrote a book on OFDM, entitled 'OFDM for Mobile Multimedia Communications.' This was the first book to describe OFDM for wireless communications, including the first book desciption of the 802.11a standard. It has been used by many engineers worldwide to learn about OFDM for wireless

LAN. He also wrote chapters on wireless LAN in other books, including the Encyclopedia of Telecommunications edited by Proakis, as well as many papers. In 2002, he received the Dutch Veder award for his contributions to standardization of wireless LAN. He holds more than 60 patents related to various WiFi standards and served as an expert witness in several WiFi related lawsuits. He is currently a Senior Director at Qualcomm where he is responsible for WiFi algorithm design and for developing new 802.11 standards.