Gigabit Wireless LANs: an overview of IEEE 802.11ac and 802.11ad

Eldad Perahia, Michelle X. Gong
{eldad.perahia, michelle.x.gong}@intel.com

Intel Corporation

This paper gives an overview of the upcoming IEEE Gigabit Wireless LAN amendments, i.e. IEEE 802.11ac and 802.11ad. Both standard amendments advance wireless networking throughput beyond gigabit rates. 802.11ac adds multi-user access techniques in the form of downlink multi-user (DL MU) multiple input multiple output (MIMO)and 80 and 160 MHz channels in the 5 GHz band for applications such as multiple simultaneous video streams throughout the home. 802.11ad takes advantage of the large swath of available spectrum in the 60 GHz band and defines protocols to enable throughput intensive applications such as wireless I/O or uncompressed video. New waveforms for 60 GHz include single carrier and orthogonal frequency division multiplex (OFDM). Enhancements beyond the new 60 GHz PHY include Personal Basic Service Set (PBSS) operation, directional medium access, and beamforming. We describe 802.11ac channelization, PHY design, MAC modifications, and DL MU MIMO. For 802.11ad, the new PHY layer, MAC enhancements, and beamforming are presented.

I. Introduction

As the IEEE 802.11n (High Throughput) standard amendment development matured and associated products became popular in the market, IEEE 802.11 initiated a new study group in May 2007 to investigate “Very High Throughput (VHT)” technologies. Initially the study group was started to address IMT-Advanced operation, in the hopes of such new spectrum in bands < 6 GHz would be allocated to unlicensed usage. However, this initial objective was dropped and the focus of the study group shifted to enhancing 802.11n in the 5 GHz band. Later in Nov 2007, discussions on operation in the 60 GHz band began. The rationale being that the millimeter wave band would provide for much wider band channels than in the microwave band, enabling single link throughputs greater than 1 Gbps.

The Wi-Fi Alliance was solicited to provide usage models to help develop requirements [1]. The general categories of the usage models included wireless display, distribution of high definition TV, rapid upload / download, backhaul, outdoor campus, auditorium, and manufacturing floor. Specific usages that will be most prevalent in the market place include compressed video streaming around a house, rapid sync-and-go, and wireless I/O. With streaming around the home, it is envisioned that TVs and DVRs around the home will have wireless capability and 100+ Mbps aggregate of videos from a DVR can be displayed wirelessly on TVs in different rooms. With rapid sync-and-go, users can quickly sync movies or pictures between mobile devices such as a phone, a laptop, or a tablet. With a 1 Gbps radio link, a 1 GB video file will take much less than a minute to transfer between devices. Data rates exceeding 1 Gbps will provide the capability for a wireless desktop, with wireless connections between a computer and peripherals such as monitors, printers, and storage devices.

With this input, the Very High Throughput study group developed two Project Authorization Requests (PARs), one for the 5 GHz band (802.11ac) and one for the 60 GHz band (802.11ad). The scope for 802.11ac includes: 1) single link throughput supporting at least 500 Mbps, 2) multi-station throughput of at least 1 Gbps, 3) exclusion of 2.4 GHz band, and 4) backward compatibility and coexistence with legacy 802.11 devices in the 5 GHz band. The PAR was approved in September 2008 and the 802.11ac task group began in November 2008.

The task group initially developed a specification framework document [2], a functional requirements and evaluation methodology document [3], an amendment to 802.11n channel model document [4], and a usage model document [5]. Based on this, an initial draft 0.1 was developed and approved by the task group in January 2011. This draft went through an internal task group comment and review cycle, and in May 2011 draft 1.0 was released to the 802.11 working group for the letter ballot process. The task group is currently in the process of addressing comments on draft 1.0, with draft 2.0 expected to be complete in November 2011. It is expected that after several working group letter ballots, the sponsor ballot will begin in Jan 2013. Final approval of the 802.11ac standard amendment is expected in December 2013. However, as was the case with 802.11n, initial products with basic

802.11ac features will emerge on the market well before final approval based on an earlier draft. (More detailed information on the process to amend the 802.11 standard is given in [7].)

As opposed to 802.11ac, which is an evolution of 802.11n, 802.11ad was formed in January 2009 to address the newly developing 60 GHz market. The initial applications of 60 GHz technology included HDMI cable replacement by transmitting uncompressed HDMI, as targeted by the WiHD industry consortium. Additional applications for PC and handheld industries covering the previously described sync-and-go and wireless docking usages were then addressed by the Wireless Gigabit Alliance.

There were also two previously completed 60 GHz standards. ECMA-387 provides a 60 GHz PHY, building upon the WiMedia UWB MAC. IEEE 802.15.3c adds a 60 GHz PHY to the 802.15.3 WPAN MAC.

To differentiate 802.11ad from the other standards, the PAR for 802.11ad includes two specific requirements, 1) enable fast session transfer between 802.11 PHYs, and 2) maintain the 802.11 user experience. Fast session transfer provides seamless rate fall back (and rate rise) between VHT in 60 GHz and 2.4/5 GHz PHYs, such as 802.11n, for multi-band devices. It also enables combo 60 GHz + 2.4/5 GHz devices to deliver expected WLAN coverage . However the PAR does not mandate that all devices have multi-band capability.

As an amendment to 802.11, VHT in 60 GHz maintains the 802.11 user experience by maintaining the network architecture of the 802.11 system, e.g. infrastructure basic service set, extended service set, access point, and station. Also the 802.11 user experience is maintained by reusing and maintaining backward compatibility to the 802.11 management plane, e.g. association, authentication, security, measurement, capability exchange, and the Management Information Base (MIB). In addition, the PAR requires that the standard support a mode of operation that enables a throughput of at least 1 Gbps.

802.11ad leverages the Wireless Gigabit Alliance specification. As such it will provide a 60 GHz installment in the successful 802.11 / Wi-Fi family. The key advantage of IEEE 802.11ad over the other standardization activities in the 60 GHz arena is that it builds on the already existing strong market presence of Wi-Fi in the 2.4/5 GHz bands.

The 802.11ad task group also initially developed a functional requirements document [8], evaluation methodology document [9], and channel model document [10]. Afterwards a complete proposal meeting these requirements was approved in May 2010 to become draft 0.1. Draft 1.0 was completed in September 2010 after an internal task group comment and review period. This and subsequent drafts were reviewed and commented on by the 802.11 working group in letter ballots. Currently 802.11ad has completed draft 5.0, which is expected to be submitted for sponsor ballot in December 2011. Final approval for the 802.11ad standard amendment is expected in December 2012.

Table 1 provides a summary of the basic technology parameters for 802.11ac and 802.11ad. Figure 1 illustrates the historical PHY data rate improved with these recent amendments.

Table 1: Basic Technology Parameters for 802.11ac and 802.11ad

	802.11ac	802.11ad
Access Technology	Multi-user + Spatial Division Multiplexing / OFDM	Single-user, one spatial stream / Single Carrier or OFDM
Frequency Band	5 GHz	60 GHz
Channel Bandwidth (MHz)	20, 40, 80, 160	2160
Maximum Data Rate (Mbps)	80 MHz, 4 spatial streams: 1733 160 MHz, 4 spatial streams: 3466 160 MHz, 8 spatial streams: 6933	Single Carrier: 4620 OFDM: 6756

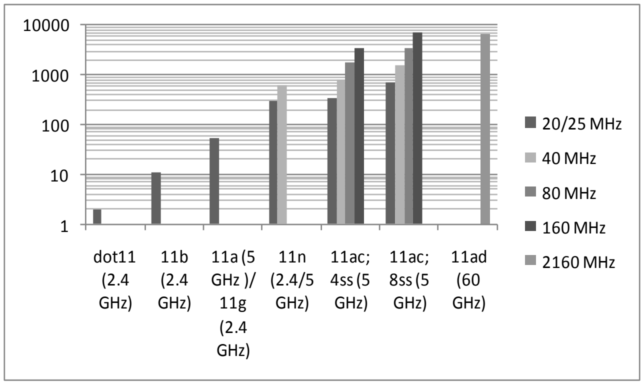

Figure 1: Historical 802.11 PHY Data Rate Improvement

In Section II, channelization, PHY design, MAC enhancements and downlink multi-user MIMO are discussed for 802.11ac. In Section III, the PHY layer, the MAC layer, and beamforming mechanisms for 802.11ad are reviewed. Concluding remarks are given in Section IV.

II. VHT in 5 GHz

802.11n provided new features such as MIMO, 40 MHz, and packet aggregation to significantly increase the data rates over 802.11a/b/g.

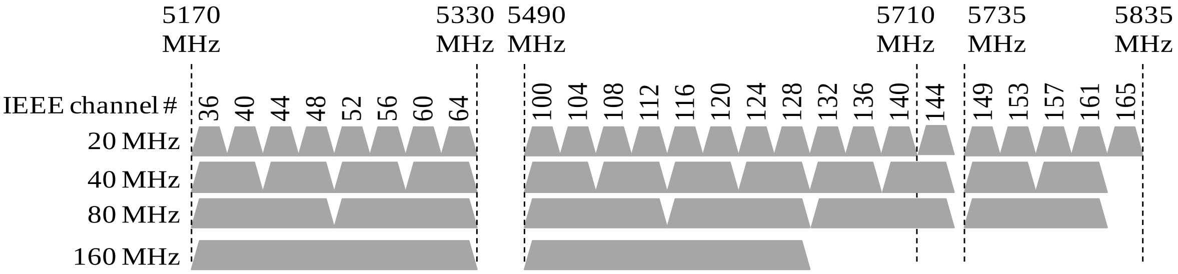

Figure 2: 802.11ac Channelization

As an evolution to 802.11n, 802.11ac adds 80 MHz, 160 MHz and non-contiguous 160 MHz (80 + 80 MHz) channel bandwidths. The other major throughput enhancement feature is multi-user capability in the form of downlink multi-user MIMO (DL MU-MIMO).

Furthermore, 802.11ac increases the modulation constellation size from 64 QAM to 256 QAM. The number of spatial streams is increased to 8 to better support DL MU-MIMO. The packet aggregation size limits are also increased to better support the higher data rates. Enhancements to the coexistence mechanisms are also provided to support the wider channel bandwidths.

Some of the 802.11n features are modified for new 802.11ac devices to simplify the mechanisms. Primarily, transmit beamforming in 802.11n had many options making interoperability between different manufacturers difficult. In 802.11ac transmit beamforming is limited to the explicit feedback mechanism (no implicit feedback). Furthermore the only type of feedback is compressed-V feedback (no uncompressed-V, no CSI). Channel sounding for transmit beamforming is limited to NDP (no staggered sounding). In addition the modulation/coding schemes (MCS) are limited to the same MCS on each stream (no unequal modulation).

The following sections will go into detail on channelization, PHY layer design, MAC features, and MU-MIMO.

II.A. Channelization

802.11n provides 20 and 40 MHz channels, with 40 MHz comprised of two adjacent 20 MHz channels. In addition, there are no partially overlapped 20 MHz channels, and there are no partially overlapped 40 MHz channels. Because partially overlapped channels introduce significant in-band interference, extremely complex coexistence schemes would have to be defined to mitigate such interference. To avoid such an in-band interference problem and to simplify protocol design, 802.11ac continues the 802.11n's design philosophy, i.e. only define non-overlapping channels. 80 MHz channels are comprised of adjacent 40 MHz channels, with no partially overlapped 80 MHz channels. And 160 MHz channels are comprised of adjacent 80 MHz channels, with no partially overlapped160 MHz channels.

Also, channel 144 has been added, which was not included in 802.11n. With this addition, there is a maximum of six 80 MHz channels possible, where regulatory bodies permit. There are only two 160 MHz channels, which is the primary reason for the inclusion of non-contiguous 160 MHz operation. Non-contiguous 160 MHz (80 + 80 MHz) channels are comprised of any two valid, non-adjacent 80 MHz channels. With non-contiguous operation many combinations of 80 + 80 MHz channelization is possible. Figure 2 illustrates the 802.11ac channelization.

II.B. PHY Design

The 802.11ac PHY design philosophy follows closely that of 802.11n. (For more information on the 802.11n PHY refer to [11].) The preamble of the packet is comprised of the following fields in the order listed:

- Legacy short training field (STF): start of packet detection, AGC setting, initial frequency and timing synchronization
- Legacy long training field (LTF): channel estimation, fine frequency and timing synchronization
- Legacy signal field (L-SIG): spoofs legacy devices, indicates VHT payload symbol length
- VHT-SIG-A: replaces 802.11n HT-SIG; contains VHT PHY single user and some MU parameters
- VHT-STF: similar to 802.11n HT-STF; allows readjustment of AGC

- VHT-LTF: similar to 802.11n HT-LTF; used for channel estimation
- VHT-SIG-B: new VHT field; contains additional per-user parameters

Like 802.11n, the fields up to and including VHT-STF are comprised of a 20 MHz waveform. This is replicated in each adjacent sub-channel for wider channel bandwidths.

Following the preamble is the data field. The first 16 bits of the data field is the SERVICE field. In 802.11ac, this has been modified to include a CRC for VHT-SIG-B. In addition, PHY padding comes after the data followed by tail bits. This is different from 802.11n where tail bits preceded pad bits; the change is due to adding MU and the rise in the maximum number of bytes per packet which in turn meant that the packet length had to be signaled differently. The data is scrambled, encoded, and the interleaved. This is followed by the constellation mapper and then the spatial mapper.

The 80 MHz waveform is based on a 256 point FFT. There are 234 data sub-carriers, 8 pilot sub-carriers, and 14 null sub-carriers, three of which are at DC. This is more than double the number of data sub-carriers as the 40 MHz waveform (108 data tones), so 80 MHz data rates are more than double the 40 MHz data rates.

However, the 160 MHz sub-carrier design is an exact replication of two 80 MHz segments. This allows for the same sub-carrier design for contiguous 160 MHz and non-contiguous 160 MHz (80 + 80 MHz). Furthermore, the phase of the local oscillator is not required to be correlated between lower and upper portions of the signal at the transmitter for contiguous 160 MHz and non-contiguous 160 MHz (80 + 80 MHz). Again, this allows for additional commonality between contiguous 160 MHz and non-contiguous 160 MHz (80 + 80 MHz).

II.C. MAC Enhancements in 802.11ac

802.11ac essentially modifies the 802.11n MAC to address coexistence and medium access with the wider channels. In addition minor modifications are made to the 802.11n aggregation mechanism to improve efficiency at gigabit per second data rates.

With the numerous 20 and 40 MHz channels in the 5 GHz band in 802.11n, overlapping channels between BSS's are easy to avoid by choosing a different channel. In the worst case if an overlap between neighbors using 40 MHz is unavoidable, the primary 20 MHz sub-channels are chosen to match to maximize coexistence capability. With much wider channels in 802.11ac, it becomes much harder to avoid overlap between neighboring BSS's. In addition it becomes harder to choose a primary channel common to all overlapping networks. To address this problem, 802.11ac improves co-channel operation with three enhancements: enhanced secondary channel Clear Channel Assessment (CCA), improved dynamic channel width operation, and a new operating mode notification frame.

In 802.11 the CCA mechanism is employed to detect other signals and defer transmission appropriately. The basic requirement for an OFDM-based device is to receive a valid signal at a level of at -82 dBm. It must also detect any other signal at a level of -62 dBm, termed Energy Detect (ED). When 802.11n added the 40 MHz channel comprised of a primary 20 MHz channel and a secondary 20 MHz channel, only ED was required on the secondary channel due to the added complexity of detecting a valid 802.11 signal on the secondary channel. This meant that other systems occupying the secondary channel of another 40 MHz BSS would be disadvantaged by 20 dB. In 802.11ac, valid signal detect on the secondary channels was added at a level of -72 dBm or -69 dBm according to bandwidth, to improve CCA performance on the secondary channels. In addition, it is required that a device detect a valid packet on the secondary channels not just based on the preamble of a packet, but also in the middle of the packet.

The basic Request-to-Send (RTS) and Clear-to-Send (CTS) mechanism of 802.11 is modified to improve dynamic channel width operation. Consider an interference scenario illustrated in Figure 3, whereby STA2 is transmitting to AP2 and AP1 is communicating with STA1. AP2 is occupying overlapping channels of the secondary 40 MHz channel of AP1. STA1 and STA2 can interfere with each other, but the interference is not heard by the two APs. To address this situation, bandwidth signaling is added to the RTS and CTS frames. As illustrated in Figure 4, AP1 sends an RTS with the bandwidth of the intended transmission, which is 80 MHz comprised of channels 36, 40 44, and 48 in this example. Before STA1 replies with a CTS frame, it senses the medium on all secondary channels for PIFS. If the secondary 40MHz channel is not free, STA1 sends a CTS response with the BW of the clear channels, i.e. 40 MHz comprised of channels 36 and 40 in this example. Then AP1 sends data to STA1 only on the clear channels and STA1 replies with Block Ack (BA) frames that are duplicated over the clear channels.

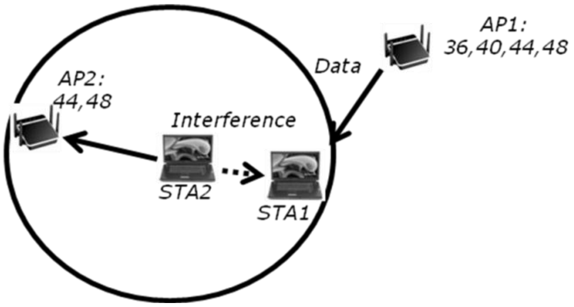

Figure 3: Interference Scenario

However, if the interference in this example is frequent, another new mechanism may be employed. In such a case, STA1 can send a Operating Mode Notification frame to AP1 to tell the AP that the client is changing its bandwidth on which it operates. For example, STA1 can change its operating bandwidth from 80 MHz to 40 MHz with the constraint that the client still needs to use the same primary channel as the AP. Subsequently the AP will only send data frames at this reduced bandwidth.

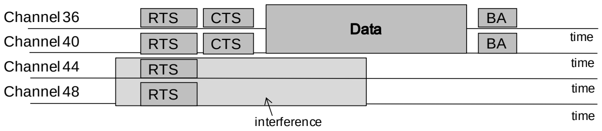

Figure 4: MAC protection for dynamic bandwidth operation

802.11n added two forms of aggregation, A-MSDU and A-MPDU. For further information on A-MSDU and A-MPDU refer to [11]. Most 802.11n devices only implemented A-MPDU. Implementing both A-MSDU and A-MPDU, while permitted, had little benefit at 802.11n data rates. However, now with the much higher data rates of 802.11ac, the combination of A-MSDU and A-MPDU aggregation is necessary to maintain good efficiency.

MSDUs that are typically less than 1500 bytes in length are aggregated to form an A-MSDU. The maximum A-MSDU size has been increased in 802.11ac to 11454 bytes. The A-MSDU is encapsulated in an MPDU. MPDUs are then aggregated to form an A-MPDU. The maximum size of an A-MPDU has been increased in 802.11ac to 1 MByte.

As an example, with a PHY data rate of 1.6 Gbps, MAC efficiency increases from 40% with A-MPDU only to 75% with the combined effect of A-MSDU, A-MPDU and increased aggregate sizes.

Another difference from 802.11n is that all 802.11ac packets are required to be A-MPDUs. This is because the PHY signal field no longer conveys the length in bytes, just in OFDM symbols. Furthermore, an MPDU contains only packet duration but not length information. Because the delimiter in an A-MPDU contains MPDU length information, 802.11ac requires that even a packet with a single MPDU is transmitted in an A-MPDU format to provide packet length information in bytes.

II.D. Downlink Multi-User MIMO

With MIMO based on spatial division multiplexing (SDM) in 802.11n, one device transmits multiple data streams to another device. In 802.11ac DL MU-MIMO, an access point (AP) simultaneously transmits data streams to multiple client devices. For example, consider an AP with 6 antennas, a handheld client device with one antenna (STA1), a laptop client device with two antennas (STA2), and a TV set top box client device with two antennas (STA3). An AP can simultaneously transmit one data stream to STA1, two data streams to STA2, and two data streams to STA3. This is illustrated in Figure 5.

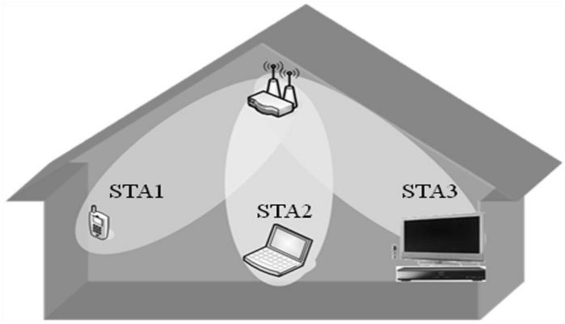

Figure 5: Example of Downlink Multi-User MIMO

The primary advantage of DL MU-MIMO is that client devices with limited capability (few or one antenna) do not degrade the network capacity by occupying too much time on air due to their lower data rates. With DL MU-MIMO, network capacity is based on the aggregate of the clients of the simultaneous transmission. However, this benefit comes with increased cost and complexity.

From a PHY perspective, the AP should have more antennas than total number of spatial streams for diversity gain. In addition, the AP requires channel state information from each of the clients participating in the DL MU-MIMO transmission in order to form the antenna weights. With DL MU-MIMO, the antenna weights are much more sensitive to changes in the channel. In the case of transmit beamforming, if the antenna weights are stale, the system performance degrades to the case without transmit beamforming. However with DL MU-MIMO, if the antenna weights do not accurately match the channel, the streams to one client introduce interference to the other clients, leading to negative (in dB) signal-to-interference-plus-noise ratio. Therefore channel state information must be higher resolution and more frequently updated. To constrain the dimensions of the system to a manageable size, 802.11ac defines that the maximum number of users in a transmission is four, the maximum number of spatial streams per user is four and the maximum total number of spatial streams (summed over the users) is eight.

As designed, a MU packet has the same preamble structure as a single user packet. However, beginning with the VHT-STF, the remaining fields in the preamble are directionally transmitted to recipient clients, simultaneously in time and frequency. The parameter information conveyed in VHT-SIG-B and the SERVICE field is specific for each client. In addition, MAC padding is required to fill the MAC frames to the last byte to make them equal in time for each client. The PHY fills in the last few bits for each client to ensure that each has the same number of symbols.

From a MAC perspective, since with DL MU-MIMO multiple packets are transmitted simultaneously to different clients, a mechanism is needed to receive acknowledgements from these clients [6]. The approach used for 802.11ac multi-user acknowledgements builds on the 802.11n implicit block acknowledgement feature.

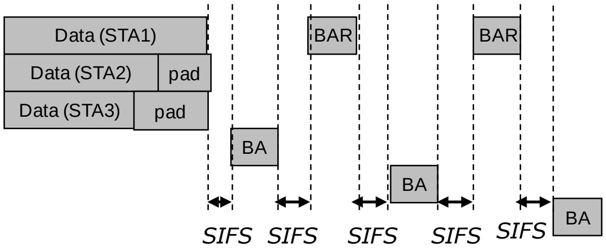

Figure 6: DL MU MIMO Response Mechanism

As illustrated in Figure 6, after transmitting a DL MU MIMO data burst, the AP uses an implicit block acknowledgement request for the first client, meaning that the first client replies immediately to the MU transmission with a Block ACK (BA). The AP subsequently polls the second client with a Block ACK Request (BAR) and the second client responds with a BA. This continues until all the clients in the original transmission are polled.

This procedure assumes that the clients know that they are part of the MU transmission and their order. This is achieved with the Group ID in VHT-SIG-A. Prior to the MU transmission, the group definition information is conveyed by the AP to all the DL MU-MIMO capable clients in the Basic Service Set (BSS). Based on the definition of the Group ID, there are 62 different possible groupings for different combination of client devices. In addition to the Group ID, VHT-SIG-A also contains a table indicating how many data streams are being transmitted to each client in the transmission.

The sounding and feedback protocol builds upon the 802.11n null data packet (NDP) mechanism (see [11] for more information on NDP). The sounding feedback sequence starts with the AP sending an NDP announcement (NDPA) frame immediately followed by an NDP. The NDPA identifies which client will be the first responder after the NDP and may identify other clients that will be polled subsequently. The client identified as first by the NDPA replies with a sounding feedback frame after the NDP. Then the AP polls all the remaining clients. Such a sequence requires a recovery mechanism in case a response is not received from a client. In this case a feedback poll can be resent to the client.

In the case of single user transmit beamforming, the same sounding and feedback protocol is used, but the sequence stops after the single client responds with feedback.

III. VHT in 60 GHz

The enormous amount of unlicensed spectrum worldwide at 60 GHz is the primary motivation for development in this band. Figure 7 illustrates the spectrum availability for major geographic areas (or regulatory domains) worldwide. Also included in the figure is the 802.11ad channelization, which defines a channel width of 2160 MHz. This enables four non-overlapping channels in Europe and three non-overlapping channels in USA, Canada, Korea, and Japan. China and Australia have fewer channels with two and one, respectively.

The primary challenge with millimeter wave operation is the much poorer pathloss than in the microwave band. For example, the free space path loss at 1 m for 60 GHz is 68 dB compared to 47 dB at 5 GHz. In addition, obstruction losses are also much higher. Brick wall attenuation has been measured to be 20 dB, a composite wall with studs in the path can result in 35 dB attenuation, and path loss due to concrete has been found to be as high at 70 dB. Even attenuation due to person obstructing the path ranges between 10 – 15 dB [15].

We further compare a 60 GHz link budget to a 5 GHz 802.11n link budget, as follows. There is a 21 dB increase in pathloss from 5 GHz to 60 GHz. In addition current technology for low cost consumer devices in 60 GHz will reduce transmit power in 60 GHz by 6 dB as compared to 5 GHz. Based on the channel bandwidth increase of 50x, the increase in noise bandwidth results in another 17 dB loss. And again current 60 GHz technology will initially result in 3 dB worse noise figure than 5 GHz technology.

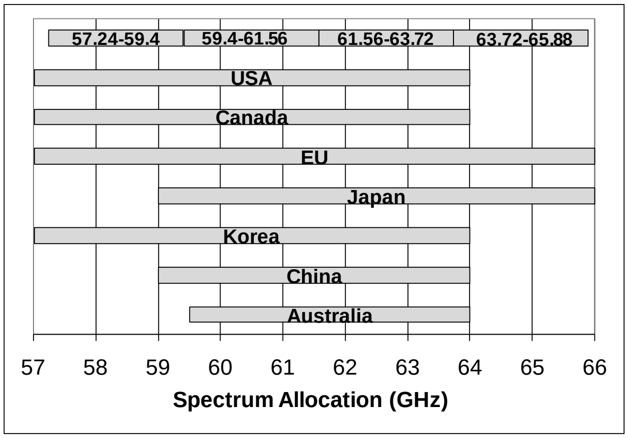

Figure 7: Worldwide Unlicensed Spectrum Allocation in 60 GHz

However, the target applications and use cases are shorter range for this technology (leaning towards personal area networking), so 10 dB can be added to the link budget margin for 60 GHz compared to longer range 5 GHz system. In addition, typically 60 GHz system design will be target a lower order modulation than an 802.11n system due to the large channel bandwidth, which could result in a 5 dB gain. In summary there could be a 32 dB loss (or larger) in link margin from 5 GHz to 60 GHz, as shown in Table 2. This leads us to the need for antenna gain at the transmitter and receiver [16].

Table 2: 60 GHz Link Budget Compared to 5 GHz 802.11n

Link budget Parameter	60 GHz compared to 5 GHz (dB)
Free space pathloss	-21
Tx power	-6
Rx Noise bandwidth	-17
Rx Noise figure	-3
Maximum range	+10
Modulation required SNR	+5
Total relative link margin	-32

The other advantage of 60 GHz besides available bandwidth is that the much smaller wavelength results in very small antenna array size even with many elements. Antenna arrays are typically designed with half wavelength spacing. At 60 GHz this would be roughly 2.5 mm. Therefore a 16 element array in a 4x4 configuration would only occupy a 1 cm x 1cm area, as shown in Figure 8. Theoretically, such an antenna array would result in 12 dBi of antenna gain. As the technology matures and antenna arrays become larger, the combination of a transmit and receive antenna array on either end of the link will make up for the previously described link margin loss.

Furthermore, the use of antenna arrays leads to another advantage. With larger number of antenna elements, the beam of the antenna array becomes more and more directional. This narrow beam pattern enables spatial reuse, a property that can be used to increase network capacity by increasing the deployment density with reduced interference between nearby devices. Analysis on spatial reuse network capacity gain is given in [12].

In addition, it has been shown in [13] that reflected paths may still have high enough signal-to-noise ratio to support a 60 GHz link. Therefore if there is an obstruction of the line-of-sight path, it may be possible with a steerable antenna to point the antenna beam pattern in the direction of a reflected path and still maintain the 60 GHz link. For example, a more advanced device could dynamically steer the antenna beam pattern if a person walked between the link of a set top box and TV.

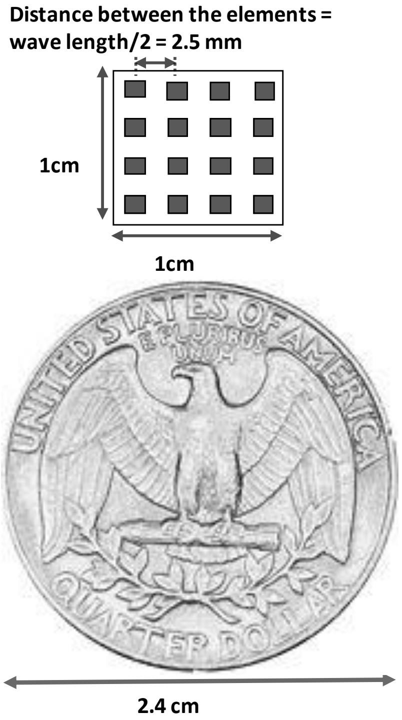

Figure 8: Example Antenna Array Size

The 802.11ad draft standard amendment includes many mechanisms to address the 60 GHz environment, device technology, and usage models. A new BSS type is defined, termed personal BSS (PBSS), to better support personal area networking. The medium access control mechanisms include both scheduled and contention access to support both streaming and bursty traffic. A unified and flexible beamforming scheme is defined to support transmit and/or receiver beam steering to enable robust communication at and beyond ranges of 10m.

As required in the PAR, a fast session transfer mechanism is defined to enable multi-band operation across 2.4/5/60 GHz. At the PHY layer, a single carrier waveform is specified for simpler devices up to 4.6 Gbps, and an OFDM waveform is defined for more advanced devices supporting data rates up to 6.8 Gbps. The following sections will go into further detail on the PHY layer, MAC layer, and beamforming.

III.A. PHY Layer

The 802.11ad PHY layer includes three different packet structures. The first is termed the Control PHY, which is designed for low SNR operation prior to beamforming. The second is the Single Carrier (SC) PHY, which enables low power / low complexity transceivers. A derivative of the SC PHY is the Low Power SC PHY, which provides additional support for further reduction in implementation processing power with simpler coding and shorter symbol structure. The third is the OFDM PHY, which provides high performance in frequency selective channels and achieves the maximum data rates by using up to 64 QAM.

To ensure interoperability between all 802.11ad devices, the Control PHY and the SC PHY are mandatory for all devices. Furthermore, to simplify the design, the three PHY types are created with common properties. Each PHY packet structure is composed of a short training field, channel estimation field, header, data, and optionally any beamforming training fields, as shown in Figure 9. In addition, the same Golay sequences are used for preamble training. Furthermore, with the exception of the Low Power SC, all PHY types use a common LDPC structure for coding. The packet structure also has embedded support for beamforming.

Figure 9: 802.11ad PHY Packet Structure

The short training field (STF) is used for packet detection, AGC, frequency offset estimation, and synchronization. It is composed of multiple repetitions of a 128 sample Golay sequence, transmitted using π/2-BPSK at the SC sample rate of 1760 MHz. Thirty eight repetitions (2.91 µsec), using a mixture of the Ga128 and Gb128 Golay sequences, are used for the Control PHY. The SC and OFDM PHY use 17 Ga128 Golay sequences (1.09 µsec). The complementary sequences are used to differentiate between the Control PHY and the SC / OFDM PHYs.

The channel estimation field (CEF) is used for channel estimation and for a receiver to differentiate SC PHY and OFDM PHY. The CEF is also created with the same Golay sequences as the STF. Nine Ga128 and Gb128 Golay sequences are used (655 nsec), with complementary sequences between the SC/Control and OFDM PHY to differentiate the SC PHY and the OFDM PHY. The Golay sequence provides the good autocorrelation properties necessary for channel estimation.

The header contains the parameters indicating the format of the packet, e.g. modulation/coding scheme and length, among others parameters, for SC and OFDM. Figure 10 illustrates the transmission block diagram for all three PHY types.

As can be seen in Figure 10, the Control PHY employs a spreading sequence on the header and data fields. A 32 sample Golay sequence enables operation in very low SNR to close links prior to beamforming. In addition, a rate ½ code is used, shortened from the common rate ¾ LDPC code. In addition to better coding gain, the shorter LDPC code is more efficient for short packets expected with the Control PHY. Furthermore π/2-differential BPSK modulation is employed for more robust phase noise performance. The resulting mandatory data rate is 27.5 Mbps.

The SC PHY is built on a 512 symbol block size. There are 448 data symbols and 64 guard interval symbols. As previously described, the SC symbol rate is 1760 MHz. In addition, π/2 rotation is applied to all modulations. This reduces the peak-to-average power ratio of the transmission for BPSK. In addition, since only BPSK is a mandatory modulation, a device could implement the modulator as GMSK. The mandatory SC data rates range from 385 Mbps to 1155 Mbps. By implementing QPSK and 16 QAM, the maximum data rate increases to 4620 Mbps.

With the Low Power SC PHY, the LDPC code is replaced by a Reed Solomon outer code and an inner block code. The rationale is that the LDPC code is one of the larger contributors to the power consumption of the SC mode.

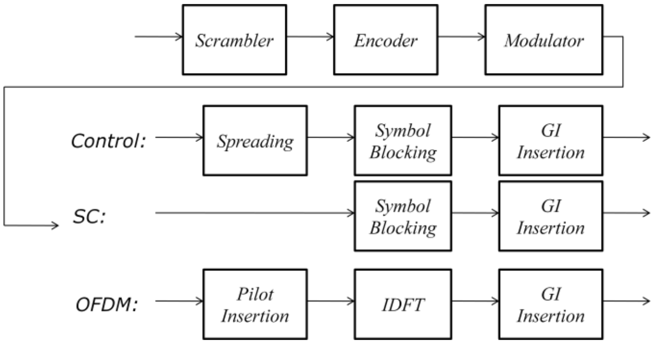

Figure 10: 802.11ad Header and Data Field Transmission Flow

The OFDM PHY has a sampling rate of 2640 MHz, exactly 1.5 times the SC symbol rate. The design is based on a 512 point FFT, with 336 data sub-carriers, 16 pilot sub-carriers, and 3 null sub-carriers at DC. The guard interval contains 128 samples. Spread QPSK is used for the two lowest data rates, and a symbol interleaver is used for the higher 16 QAM and 64 QAM constellation sizes. The OFDM data rates range from 693 Mbps to a maximum of 6756.75 Mbps.

III.B. MAC Layer

The primary challenge for the VHT 60 GHz MAC is how to handle directional transmissions used to mitigate the high propagation loss of millimeter waves [14]. Prior to 802.11ad, the 802.11 MAC relies on omni-directional transmissions for management and control frames (e.g. beacons, RTS/CTS) to ensure they can be heard by everyone in the BSS. With directionality, device discovery becomes more complicated. Furthermore, devices need to find the best direction for communication, which necessitates support for a beamforming protocol. Also the basic 802.11 concept of CSMA/CA has limitations in the presence of directional transmissions.

Several new MAC mechanisms are introduced in 802.11ad to address these issues. A new network architecture termed personal BSS (PBSS) is created, while retaining existing infrastructure and independent BSS network architectures. Channel access has been enhanced to support directionality and spatial reuse, including both random and scheduled access.

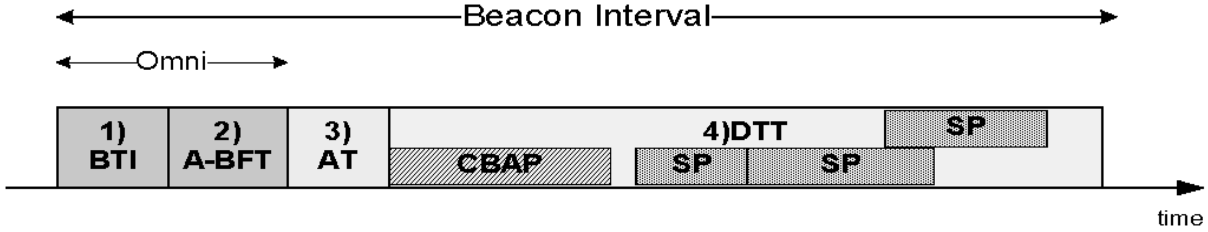

Figure 11: Beacon Interval Structure

The PBSS is similar to an independent BSS used for ad hoc networking, with the following differences. One station assumes the role of the PBSS Control Point (PCP). Only the PCP transmits beacon frames. The beacon interval (BI) structure is illustrated in Figure 11. The first phase is the beacon transmission interval (BTI), where discovery of new stations occurs. In the BTI, an AP or PCP performs one or more beacon transmissions potentially in different directions. The second phase is association beamforming training (A-BFT), where beamforming occurs between the AP or PCP and the stations (STA). The third phase is the announcement time (AT), used to convey control and management information between the AP/PCP and the stations, e.g. association, schedule, etc. The forth and last phase is the data transfer time (DTT), which contains the contention-based access periods (CBAPs) and service periods (SPs). The prescribed stations access the channel during SPs, based on negotiation with the AP/PCP. Or, an SP can be dynamically allocated to a STA. Any station can access the channel during a CBAP. Directional medium access rules for CBAPs, which are based on 802.11 EDCA rules, have been defined in 11ad. The 802.11 Clear-To-Send (CTS) and Contention-Free End (CF-End) frames are no longer valid for 11ad STAs. Instead, 802.11ad defines a Directional Band CTS (DBandCTS) frame, a DBand Denial To Send (DBandDTS) frame, and a DBandCF-End frame. All three frames are transmitted using directional mode. The frame format of DBandCTS is similar to CTS, except that DBandCTS contains an extra field, i.e. the Transmitter Address (TA) field. This is to convey the address information of both the transmitter and the receiver to STAs that did not receive the RTS frame. In 802.11ad, instead of maintaining one overall Network Allocation Vector (NAV) counter, a STA maintains one NAV timer for each source and destination pair. Upon receiving a RTS frame addressed to itself, if at the start of the RTS reception a STA has a non-zero NAV timer, the STA does not return a DBandCTS frame. Instead, it may transmit a DBandDTS frame directionally to the transmitter of the RTS frame. The source STA of a Service Period (SP) can transmit a DBandCF-End frame to the destination STA of the SP and to the PCP/AP to truncate a SP.

Channel access in DTT is coordinated using a schedule, which is delivered by the PCP/AP to the non-PCP/non-AP stations. Access during SPs is reserved to specific stations, as announced in the schedule or granted by the PCP/AP. With such scheduled access, 90% MAC efficiency can be achieved for payload sizes larger than 8 Kbytes for 1 Gbps PHY data rate, payload sizes larger than 16 Kbytes for 2 Gbps PHY data rate, and payload sizes larger than 64 Kbytes for 4 Gbps PHY data rate.

The fast session transfer (FST) introduced in 802.11ad enables seamless integration of 60 GHz with 802.11a/b/g/n/ac for multi-band devices. FST allows for the transfer of communication from any band or channel to any other band or channel. The mechanism also supports both simultaneous and non-simultaneous operation. In addition, FST is flexible for both transparent and non-transparent implementations. With transparent operation, the MAC address is the same on both bands or channels. With non-transparent operation, the MAC addresses are different.

III.C. Beamforming

A high antenna gain is necessary to compensate for the propagation loss at 60 GHz. As the antenna gain increases, the antenna beamwidth becomes narrower (e.g. a 13 dB gain antenna will have approximately 45 degree antenna beamwidth). Antenna beam pointing must be performed automatically to find the best path and potentially to avoid obstructions. In addition, different techniques may be used to implement 60 GHz antenna systems, e.g. switched beams, phased arrays, fully adaptive arrays, and multiple arrays. To accommodate different antenna techniques and systems, 802.11ad specifies a unified and flexible beamforming scheme that can be tuned to simple, low power devices as well as complex devices.

The beamforming training protocol is performed in three phases. The first phase is the sector level sweep (SLS) to select the best transmit and optionally receive antenna sector. The second phase is the beam refinement phase to train the transmit and receive antenna arrays. The last phase is beam tracking during data transmission to adjust for channel changes.

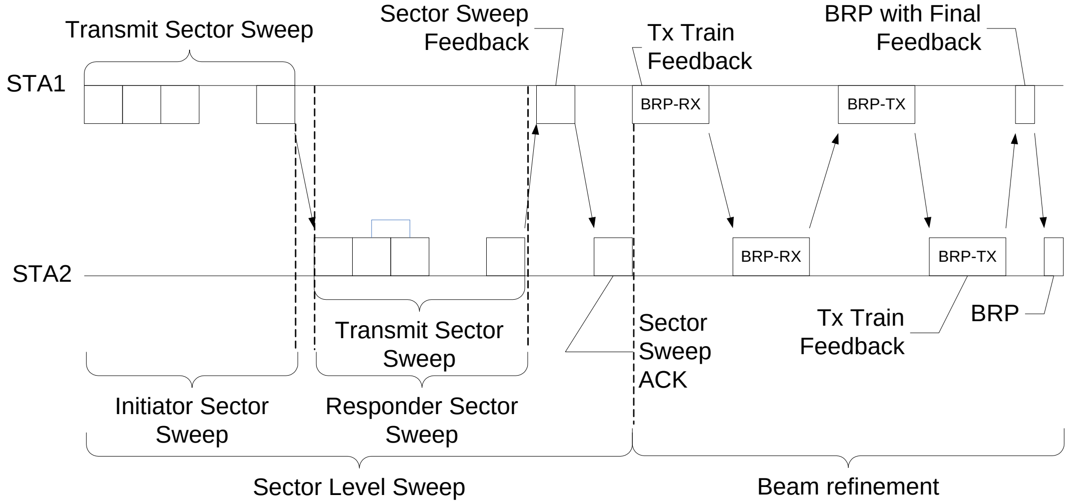

Figure 12: Beamforming Training

During SLS, the initiator of the beamforming exchange (STA1 in Figure 12) transmits training frames for each transmit antenna sector it wishes to train. The responder (STA2 in Figure 12) configures its antenna to have a quasi-omni directional pattern. This is followed by the responder sending a set of training frames for each transmit antenna sector it wishes to train, in which STA1 configures it's antenna to have a quasi-omni directional pattern. Sector sweep feedback information is then exchanged between the two devices. At this point in the process, both the initiator and the responder each possess their best transmit sector. If one of the devices chooses to use only one transmit antenna (e.g. a handset device with a single antenna), receive training of the other device can be performed as part of the SLS.

In the BRP, the receive antenna may be adapted. At both the initiator and the responder, multiple sets of antenna weights are tested against quasi-omni directional transmit antenna. This reverses the roles from the SLS phase. Furthermore, in pair-wise combination, small sets of transmit and receive antenna weights can be further tested, in case the quasi-omni directional patterns initially used were imperfect.

Beam tracking is performed by appending training fields to data packets, as illustrated in Figure 9. The training fields are used to make measurements, which are fed back in subsequent packets.

IV. Conclusion

Wi-Fi products based on next generation gigabit per second 802.11 technology will be emerging on the market soon to address use cases that demand higher throughput. 802.11ac will evolve in the 5 GHz bands with wider channels and multi-user capability to address broader coverage use cases typical of Wi-Fi devices, such as higher resolution video coverage around the home. 802.11ad will address personal area networking use cases new to 802.11 such as wireless docking with multi-gigabit per second links based on large amount of available spectrum in the 60 GHz band. 802.11ad will make use of directional antennas and beamforming to enhance link quality, and modifies channel access to address directionality and spatial reuse.

References

[1] A. Myles and R. de Vegt, "Wi-Fi Alliance (WFA) VHT Study Group Usage Models", IEEE 802.11-07/2988r4, March 19, 2008, https://mentor.ieee.org/802.11/dcn/07/11-07-2988-04-0000-liaison-from-wi-fi-alliance-to-802-11-regarding-wfa-vht-study-group-consolidation-of-usage-models.ppt.

[2] R. Stacey, E. Perahia, A. Stephens, et. al., "Specification Framework for TGac", IEEE 802.11-09/0992r21, Jan 19, 2011, https://mentor.ieee.org/802.11/dcn/09/11-09-0992-21-00ac-proposed-specification-framework-for-tgac.doc.

[3] P. Loc and M. Cheong, "TGac Functional Requirements and Evaluation Methodology", IEEE 802.11-09/0451r16, January 19, 2011, https://mentor.ieee.org/802.11/dcn/09/11-09-0451-16-00ac-tgac-functional-requirements-and-evaluation-methodology.doc.

[4] G. Breit, H. Sampath, S. Vermani, et. al., "TGac Channel Model Addendum", IEEE 802.11-09/0308r12, March 18, 2010, https://mentor.ieee.org/802.11/dcn/09/11-09-0308-12-00ac-tgac-channel-model-addendum-document.doc.

[5] R. De Vegt, "802.11ac Usage Models Document", IEEE 802.11-09/0161r2, March 9, 2009, https://mentor.ieee.org/802.11/dcn/09/11-09-0161-02-00ac-802-11ac-usage-model-document.ppt.

[6] M. Gong, E. Perahia, R. Stacey, R. Want, and S. Mao, "A CSMA/CA MAC Protocol for Multi-User MIMO Wireless LANs," IEEE Globecom 2010, Dec. 2010.

[7] E. Perahia, "IEEE 802.11n Development: History, Process, and Technology", IEEE Communications Magazine, Vol. 46, Issue 7, pg 48-55, 2008.

[8] E. Perahia, "TGad Functional Requirements", IEEE 802.11-09/0228r4, July 16, 2009, https://mentor.ieee.org/802.11/dcn/09/11-09-0228-05-00ad-functional-requirements.doc.

[9] E. Perahia, "TGad Evaluation Methodology", IEEE 802.11-09/0296r16, January 20, 2010, https://mentor.ieee.org/802.11/dcn/09/11-09-0296-16-00ad-evaluation-methodology.doc.

[10] A. Maltsev, V. Erceg, E. Perahia, et. al., "Channel Models for 60 GHz WLAN Systems", IEEE 802.11-09/0334r8, May 20, 2010, https://mentor.ieee.org/802.11/dcn/09/11-09-0334-08-00ad-channel-models-for-60-ghz-wlan-systems.doc.

[11] E. Perahia and R. Stacey, "Next Generation Wireless LANs: Throughput, Robustness, and Reliability in 802.11n", Cambridge University Press, 2008.

[12] M. Park, C. Cordeiro, E. Perahia, et. al., "Millimeter-wave multi-Gigabit WLAN: Challenges and feasibility", IEEE 19th International Symposium on Personal, Indoor and Mobile Radio Communications, September 2008.

[13] A. Maltsev, E. Perahia, R. Maslennikov, et. al., "60 GHz WLAN Experimental Investigations", IEEE 802.11-08/1044r0, September 9, 2008, https://mentor.ieee.org/802.11/dcn/08/11-08-1044-00-0vht-60-ghz-wlan-experimental-investigations.ppt.

[14] C. Cordeiro and S. Shankar, "Next Generation Multi-Gbps Wireless LANs and PANs", IEEE Globecom 2010, Dec 2010.

[15] J. Schontheir, "WP3-Study, The 60 GHz Channel and its Modelling", IST-2001-32686 BROADWAY, v1.0, May 30, 2003.

[16] C. Hansen, et. al., "NT-11 Beamforming Introduction", IEEE 802.11-10/0430r1, May 16, 2010, https://mentor.ieee.org/802.11/dcn/10/11-10-0430-01-00ad-nt-11.ppt.