The IEEE 802.11ad Standard: Challenges and Design Adaptations

Mahdi Koubaa, Stephan Günther*, Jonas Andre*
*Chair of Network Architectures and Services, Department of Informatics
Technical University of Munich, Germany
Email: mahdi.koubaa@tum.de, guenther@tum.de, andre@net.in.tum.de

***Abstract*—After the introduction of IEEE 802.11ad in December 2012, a new multi-gigabit Wi-Fi connection was made possible allowing new applications to be performed. However, using a 60 GHz frequency introduced new problems which were fixed by implementing a new design. This paper describes the faced challenges along with the organization and features of IEEE 802.11ad.**

***Index Terms*—ieee 802.11ad standard, wigig, multi-gigabit wi-fi**

1. Introduction

The 802.11ad standard was introduced primarily to provide a multi-gigabit wireless solution. It uses the 60 GHz carrier frequency and can deliver a data rate of 7 Gbit/s. This standard is mainly used for wireless data transmission and wireless displays.

1.1. The Use of 60 GHz Frequency

The main reason for using a high frequency in 802.11ad is that in higher frequencies a wider bandwidth is achievable without creating interferences in contrast to lower frequencies where interferences are more likely to happen. In the case of 802.11ad, a frequency of 60 GHz allows us to define a single carrier with a bandwidth of 2.16 GHz which is approximately 14 times wider than a single carrier bandwidth in a 5 GHz legacy Wi-Fi frequency [1]. The data rate should increase if we increase the bandwidth while maintaining similar network characteristics, such as the number of users and the modulation and channel encoding scheme. This property of channels has been very useful for building the Multi-Gigabit Wi-Fi.

1.2. Problems Occurring When Using 60 GHz

Due to the small wavelength in higher frequencies, an 802.11ad signal cannot propagate through walls and concrete objects, which means strong signals are originating either from a line-of-sight (LOS) path or from first order reflections on highly reflective materials. Additionally, since the Oxygen absorption of waves peaks at 60 GHz, signal attenuation is strong and signal range will be limited to approximately 10 m [1]. 802.11ad addresses this issue by implementing directional communications through beamforming antenna arrays. These antenna arrays can be weighted to concentrate signal focus in the intended direction and gain wider signal range [1]. This process will be described later and is called beamforming. To facilitate beamforming, the antenna space is partitioned into sectors representing the multiple directions that can be selected.

2. Physical Layer

In a physical layer, data should be encoded by channel encoding to assert a level of transmission failure detection and correction. Additionally, the 802.11ad physical packet is constructed. A few changes have been applied to the structure of the physical packet to adapt to the concept of directional communication. In a later phase, the signal should be modulated on a 60 GHz carrier signal. The IEEE 802.11ad standard supports diverse types of physical layer (PHY) which execute these phases differently.

2.1. Structure of The Physical Packet

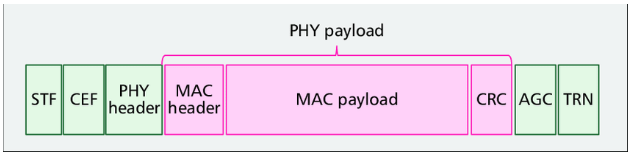

Figure 1: IEEE 802.11ad packet-structure [1]

The IEEE 802.11ad packet begins with a preamble containing a short training field (STF) and a channel estimation field (CEF) which are used in the detection of the implemented PHY and in the reconstruction of the original signal in case of weak channel conditions [2]. The CEF helps the receiver restore the original signal after distortion. The preamble is followed by a PHY Header containing essential information like the modulation and channel encoding scheme (MCS) used to transmit the data. Additionally, it contains the size of the transmitted data. The third part is the data i.e. PHY-Payload generated by the MAC-Layer [2]. An optional automatic gain control field (AGC) follows the data. This field carries information which helps in equalizing the signal amplitude. IEEE 802.11ad introduces a new optional training field (TRN). This field is newly introduced in the 802.11ad standard and is used by stations to train their antennas through beamforming [2].

2.2. Several Types of Physical Layers

To adapt to different use cases, different PHY layers were introduced. The control PHY was intended to work with low signal-to-noise ratio (SNR) operations preceding and during beamforming. Consequently, the control PHY

doi: 10.2313/NET-2023-06-1_09

Packet contains a longer STF field providing a better resistance to unwanted channel effects, compared to the remaining PHY types. Additionally, the control PHY uses MCS0 which implements a BPSK Modulation for better noise resistance and a robust channel encoding scheme with encoding rate $\frac{1}{2}$ to withstand data transmission failures. Since the control PHY is mainly applied to transmit control information between two pairing stations, the packets exchanged contain small data fields with a limit of 1023 B [2]. Due to the use of a binary modulation scheme, the limited data field size, and the low channel encoding rate, the data rate is limited to 27.5 Mbit/s when using this PHY Layer.

To achieve higher transmission rates after beamforming, 802.11ad presents the single carrier PHY (SC PHY) and orthogonal frequency division multiplexing PHY (OFDM). Packets in these PHY layers contain data fields reaching 262 143 B [2]. In these PHY layers the data rate achieved depends strongly on the MCS used.

In the single carrier PHY (SC PHY) data is modulated on a single carrier signal with a 1760 MHz bandwidth [2]. When first introduced, this PHY implemented 12 different MCS allowing different data rates scaling from 385 Mbit/s up to 4620 Mbit/s [2]. Recently this PHY has been extended with new MCS variants to support 8085 Mbit/s [2]. To reduce power consumption in mobile devices, a low-power SC PHY was introduced allowing 5 MCS methods. A trade-off is that the data rate in this PHY is limited to 2503 Mbit/s [2].

The OFDM PHY applies a frequency multiplexing method that modulates data on multiple subcarriers. The low data transmission rates achieved in each subcarrier are added together to result in a high data rate which can reach 6756 Mbit/s on the complete 1830.47 MHz band [2]. Using the OFDM PHY results in high energy costs and is no longer the fastest alternative since the extension of the SC PHY. Therefore, this PHY type is obsolete. The 802.11ad hardware is not required to implement all MCS methods which creates differences regarding performance between 802.11ad supporting devices.

3. Personal Basic Service Set

To make use of the directional communication, a new architecture concept “Personal basic Service Set” (PBSS) was introduced allowing peer-to-peer connections between stations [1]. Thanks to directionality, multiple peer-to-peer connections are allowed to coexist without resulting in interferences allowing spatial sharing. However, medium access control (MAC) in a PBSS network must be centralized in one node called "PBSS contol point" (PCP) [1]. The centralization of MAC is necessary for some of the MAC mechanisms described in the upcoming sections. If two stations intend to start a P2P communication in the absence of a PCP, the PCP role must be taken temporarily by one of them. Centralization of PCP can cause the workload to not be distributed equally through the network causing power management issues. Allowing PCPs to hand over the PCP role to other stations would be a good approach to handle this problem [1]. A PCP breakdown causing the whole network to become dysfunctional can potentially represent a vulnerability. This issue can be fixed by implementing an implicit PCP handover procedure which chooses the best alternative PCP when the former PCP is unreachable [1]. PBSS meets the requirements of many applications where ad-hoc-like communications are intended like wireless displays or wireless data storage devices [1].

4. Beacon Interval

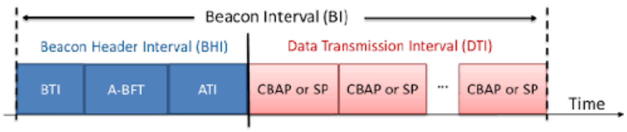

Figure 2: 802.11ad Beacon Interval [1]

The medium access control in a 802.11ad network architecture is managed through periodically recurring beacon intervals (BI) consisting of a beacon header interval (BHI) followed by a data transmission interval (DTI) [1]. The BHI replaces the beacon frame used in legacy Wi-Fi architectures. It includes a beacon transmission interval (BTI). During BTI the PCP/AP performs a sector-level sweeping to transmit beacon frames in all directions. The transmitted beacon frames are used for network announcements as well as for training the PCP/AP transmitter antennas [1]. An explanation of the sector-level sweeping and the beamforming training process is included in the last section of this paper. BTI is followed by an association beamforming training (A-BFT) which is used by stations to train their transmitter antennas and by the PCP/AP to complete its beamforming training. After associating an A-BFT trained station with the PCP/AP, directional communication can be initiated. Since BTI and A-BFT implement the control PHY for a robust association, data rates in these subintervals are too low and result in overheads due to the recurring nature of beacon intervals which constitutes a pertinent problem for some real-time applications such as wireless displays. Solving this problem included the outsourcing of information transmissions from BTI to a new subinterval [1]. As a result, beacon frames transmitted during BTI were restricted to the necessary information to minimize transmission overhead. To reallocate the outsourced transmissions, a new subinterval “announcement transmission interval” (ATI) was defined. During ATI, management information is exchanged with associated stations [1]. This information is necessary for MAC mechanisms used in the DTI. Since associated stations have trained antennas, the control PHY is no longer used in ATI and high data rate transmissions can be performed. The further data transmissions are performed during DTI. The DTI is partitioned into contention-based access periods (CBAP) and scheduled service periods (SP). During a CBAP, stations contend for medium access. The SP is a contention-free period reserved for P2P communication between two assigned stations [1].

5. Medium Access Control

IEEE 802.11ad uses a contention-based mechanism for medium access control. However, the exclusive use of contention for medium access in directional communications can cause problems. To better understand the issues that arise, we first review the contention concept in 802.11ad.

doi: 10.2313/NET-2023-06-1_09

We will then identify the problems. This is followed by an introduction to the adaptation techniques that are in use to address the issues.

5.1. Contention-based Medium Access

Performed in a CBAP, the contention based access in IEEE 802.11ad implements a carrier sense multiple access with collision avoidance (CSMA/CA) expanded with a request-to-send (RTS)/clear-to-send (CTS) exchange. In this medium access method, a node stores for each peer station a network allocation vectors (NAV) timer [1]. A NAV timer is used to know the time left for a peer station in its current communication. NAV timers are usually updated with greater duration field values from RTS/CTS frames received by overhearing communications between other peer stations. When contending for the carrier, a node performs a virtual carrier sensing which is done by evaluating the NAV entry of the destination node. If the NAV entry is non-zero, the next connection attempt will be scheduled after the timer expiration [3]. Otherwise, a physical carrier sensing is executed. The physical carrier sense defines the channel as idle if it does not become busy during a distributed coordination function interframe Space (DIFS) interval. If the channel is sensed idle, an RTS frame is sent to the intended receiver [3]. In case of receiving a CTS frame response, a p2p transmission is initiated. Time elapsing without receiving a CTS or physically sensing an occupied channel is identified as a collision and will cause the next attempt to be scheduled after a *binary exponential backoff* interval. A description of the *binary exponential backoff* can be found in [4].

5.2. Problems Occurring in Contention-Based Access

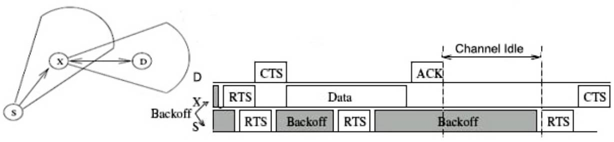

Figure 3: Example of deafness situation [5]

Using contention-based medium access exclusively can be problematic. While waiting for RTS frames, pairing stations do not know the direction of the next transmission. Therefore, nodes are forced to apply a quasi-omnidirectional antenna pattern to deal with such situations, which reduces the receive signal strength. An additional problem is deafness. Fig. 3 shows a situation where a deafness situation arises. We assume the existence of 3 misaligned stations S, D and X. To start a communication with X, S must beamform in the direction of X and contend for the channel. Let us further assume that X already succeeded in initiating a directional connection with D. Since while contending, the RTS/CTS frames are transmitted directionally, such transmissions may not be overheard by S. Consequently, S may not be able to update the NAV timer corresponding to X. As a result, the virtual carrier sensing fails to identify the channel as occupied. If the connection between X and D is long lasting, the station S will experience multiple collisions inducing series of backoff intervals. Due to the nature of binary exponential backoff, there is a high probability for S to be counting down a large backoff interval when X becomes available. Throughout this backoff interval, X might start a new communication with other stations, which could cause more delay for the communication between X and S to take place. This behavior can lead to a severe starvation for some network nodes and create unfairness. While NAV timers cannot always be helpful against deafness, implementing NAV timers remains fundamental to address the deafness problem for contending nodes.

5.3. Hybrid Medium Access Control

To deal with the inefficiencies of the contention-based access, 802.11ad implements a hybrid medium access control combining this scheme with two new mechanisms. The methods in question are Dynamic Channel Time Allocation and Time Division Multiple Access (TDMA).

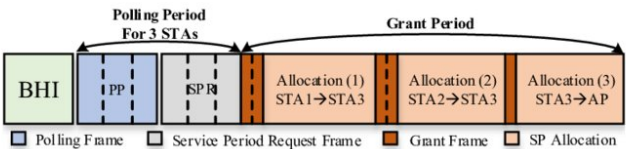

Figure 4: Dynamic Channel Time Allocation [6]

5.3.1. Dynamic Channel Time Allocation. Dynamic Channel Time Allocation is based on a polling phase followed by allocation periods. In the polling phase the PCP/AP sends polling frames to the associated stations. A polling frame schedules the receiving node to submit a service period request (SPR) frame to the PCP/AP to ask for channel time [1]. Afterwards, the PCP/AP allocates channel time according to the SPR frames received. Every allocation period is preceded by an associated grant period. During grant periods grant frames are sent to the individual communication peers of the following allocation period, allowing them channel access. If the PCP/AP takes part in the communication, only one grant frame is transmitted to the non-PCP/AP station [1]. In this medium access approach stations know the directions of the incoming frames since all the instructions come from the PCP/AP. Thus, no quasi-omnidirectional antenna patterns are needed, which improves the signal strength [1]. A Dynamic Channel Time Allocation can be executed in both SP and CBAP periods. When implemented in CBAP, there is a risk that the dynamic channel time allocation becomes disturbed by contending stations that try to acquire the channel. To minimize this risk, the PCP/AP applies a physical carrier sensing method using a point coordination function interframe space (PIFS) interval, which is shorter than the DIFS interval used in contention-based access [1]. This change allows the PCP/AP a more frequent and thus prioritized access to the channel in contrast to other stations. Additionally, extended direction fields are included in the polling and SPR frames to try protecting the polling phase by updating NAV timers of contending stations. The allocation periods are protected by direction fields in the preceding grant frames [1]. If used in CBAP, the PCP/AP can allow contention after performing the requested allocations, otherwise a new polling phase is started [1].

doi: 10.2313/NET-2023-06-1_09

5.3.2. Time Division Multiple Access. Stations using TDMA send resource requests during ATI to request channel time in the upcoming BI.In the next step, the PCP/AP schedules the requested allocations and associates them separately to SPs. The schedule of SPs will then be transmitted by the PCP/AP to all the associated stations in the next ATI [1]. As a result, stations not communicating during an SP enter sleep mode to save energy. This method is oriented to satisfy quality of service (QoS) requirements. Therefore, a resource request should include parameters like the allocation duration and isochronous or asynchronous traffic properties [1]. For duration of communication to be determined, the link in question should completely beam-trained and the transmission rate between the communicating parts should be known [1]. If the traffic stream is isochronous, channel allocations are adapted for meeting certain latency requirements in a constant-rate recurring payload. This type of payload is heavily implemented in wireless displays. In case of asynchronous traffic streams, channel allocations are optimized for non-recurring payload requirements used mostly in file downloads [1]. TDMA makes use of the directional nature of the occurring communications to achieve spatial sharing. Spatial sharing allows noninterfering communications to be initiated concurrently [7].

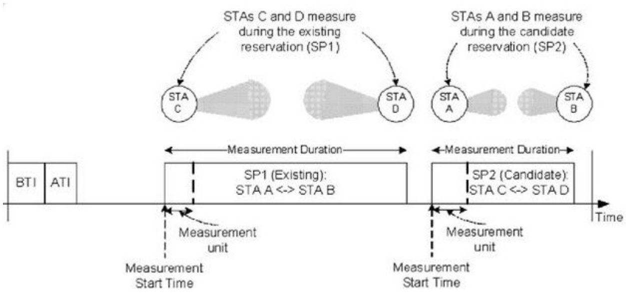

Figure 5: 802.11ad Spatial Sharing Assessment [7]

Fig. 5 describes how the spatial sharing is performed. We assume the existence of two pairs of communicating peers STA A and STA B, STA C and STA D with corresponding SPs, SP1 and SP2. At first the SPs are allocated by the PCP/AP in different time slots. To test the interference between the two communications, the PCP/AP requests STA C and STA D to beamform and measure average noise plus interference power indicator (ANIPI) during SP1. The same is requested from STA A and STA B during SP2. After completion, the obtained values are transmitted to the PCP/AP. The PCP/AP uses these values to decide if SP1 and SP2 should be allocated concurrently. To prevent interferences in case of sudden changes in the network, the PCP/AP periodically checks the existing spatial sharing configuration by requesting similar measurements to see if the SPs should remain concurrent.

6. Beamforming

To preform beamforming training, two phases are defined: sector-level sweep (SLS) and beam refinement protocol (BRP).

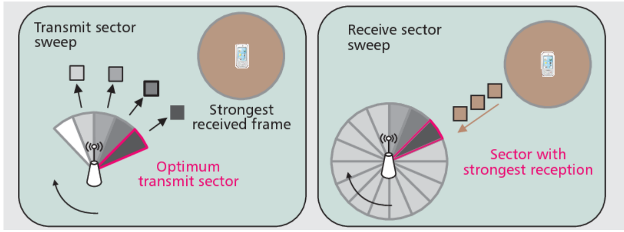

Figure 6: Transmit and Receive Sector Sweep [1]

6.1. Sector-Level Sweep Phase

When performing SLS on a pair of stations, both nodes receive training. The first station to be trained is the initiator, the second is the responder. To train transmitter antennas, a transmit sector sweep (TXSS) is performed. As shown in the left part of Fig. 6, in TXSS the training station sends sector sweep (SSW) frames using different sectors. Meanwhile, the pairing node uses a quasi-omnidirectional receive pattern to measure the SNR values of the received frames. Afterwards, the pairing station reports the optimum SNR and the sector identifier from the corresponding frame in an SSW Feedback [1]. For the receiver antenna training, a receive sector sweep (RXSS) is performed as shown in the right side of Fig. 6. In RXSS, SSW frames are transmitted omni-directionally to the training node which tries to measure the SNR using multiple receive sectors and pick the optimum sector. The SSW feedback for RXSS includes the optimum SNR [1]. During BTI and A-BFT, a sector-level sweep is implemented as follows. A TXSS for the PCP/AP is performed during BTI. Instead of SSW frames, Beacon frames are used to include network announcements.

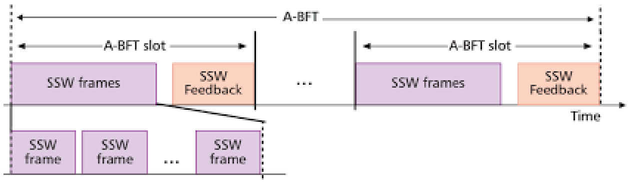

Figure 7: Sector-Level Sweeping in A-BFT [1]

The A-BFT generally includes TXSS for responder nodes. Since multiple responder nodes exist, the PCP/AP prepares multiple A-BFT timeslots. In each A-BFT slot a responder TXSS is performed. The feedback for the PCP TXSS is included in all the SSW frames transmitted during the responder TXSS. Afterwards, SSW feedback for the responder TXSS is sent by the PCP/AP. Stations must contend for accessing a timeslot. Since collisions can be detected if the PCP/AP does not respond with SSW feedback, carrier sensing is not used [1]. The PCP/AP can announce an A-BFT slot for a PCP RXSS [1], which is helpful if a responder node TXSS has already taken place in earlier BIs. Receiver antenna training can also be rescheduled to upcoming SLS or BRP phases in the DTI.

6.2. Beam Refinement Protocol Phase

The BRP procedure tries to refine the antenna directions adjusting the antenna array weights independently from the predefined antenna sectors by trying multiple antenna configurations [1]. This process is used to further improve the signal quality. Since BRP is following an SLS Phase, the pairing node could avoid using the quasi-omnidirectional pattern if a better directional pattern is

doi: 10.2313/NET-2023-06-1_09

known [1]. Instead of using multiple frames to test antenna patterns, multiple training fields are included together with parameters like the number of tested patterns inside the same frame. Consequently, using BRP reduces transmission overhead when compared to SLS [1]. This protocol is mainly used in DTI.

6.3. Beam-Training in DTI

As mentioned previously Beamforming training can be used in DTI. When using contention-based access training can be requested directly between stations without the help of a PCP/AP [1]. However, in dynamic channel time allocation and TDMA, training is requested respectively by SPRs and resource requests. The PCP/AP will then transfer the training parameters using grant frames and announcements [1].

7. Linux Support for IEEE 802.11ad

Support for WiGig has been included in the latest Linux kernels through the Wil6210 driver. This driver supports both AP mode, where the node operates as an access point, and station mode. However, the driver only supports Wilocity chips and does not support Intel adapters. In addition, it can be difficult to find a suitable adapter and beamforming antenna array.

8. Multi-Gigabit Throughput with Low Frequency Standards

An example of a multi-gigabit wireless standard that uses low frequency carriers is 802.11ax. This standard uses 1024-QAM modulation on a 160 MHz bandwidth and a 5GHz or 2.4 GHz Hz carrier. It can implement $\frac{5}{6}$ rate coding and use frequency multiplexing. In addition, this standard uses Multi-User Multiple-Input Multiple-Output (MU-MIMO) as a spatial multiplexing technique to transmit multiple streams simultaneously using the same carrier frequency. The throughput in this standard can reach 9602 Mbit/s. The use of low carrier frequencies has some advantages, such as long-range signal propagation and backward compatibility with older Wi-Fi standards. However, this approach uses complex modulation schemes that can be sensitive to noise.

9. Conclusion

To adapt to 60 GHz attenuated signal and LOS propagation, IEEE 802.11ad supports directionality. This paper describes the issues faced and the design adaptations introduced by 802.11ad to integrate this new concept. Different types of PHY are supported; Control, SC, and OFDM PHY, to provide more flexibility towards use cases. The PBSS was intoduced to benefit from antenna directionality and focus on P2P communications allowing new applications e.g. high throughput video streaming. 802.11ad implements a hybrid MAC combining Contention with TDMA and polling methods to overcome the deafness problem. A beamforming training mechanism was implemented to maintain directionality between pairing stations.The 802.11ad standard is partially supported by Linux via the Wil2160 driver. Finally, the IEEE 802.11ad standard has some alternatives in the 5GHz and 2.4GHz wireless standards, such as the 802.11ax standard.

References

[1] T. Nitsche, C. Cordeiro, A. B. Flores, E. W. Knightly, E. Perahia, and J. C. Widmer, "Ieee 802.11ad: directional 60 ghz communication for multi-gigabit-per-second wi-fi [invited paper]," *IEEE Communications Magazine*, vol. 52, no. 12, pp. 132–141, 2014.

[2] B. Schulz, "802.11ad - wlan at 60 ghz a technology introduction," 11 2017.

[3] S. S. N., D. Dash, H. E. Madi, and G. Gopalakrishnan, "Wigig and IEEE 802.11ad - for multi-gigabyte-per-second WPAN and WLAN," *CoRR*, vol. abs/1211.7356, 2012. [Online]. Available: http://arxiv.org/abs/1211.7356

[4] Wikipedia contributors, "Exponential backoff — Wikipedia, the free encyclopedia," 2022, [Online; accessed 10-December-2022]. [Online]. Available: https://en.wikipedia.org/wiki/Exponential_backoff

[5] R. Choudhury and N. Vaidya, "Deafness: a mac problem in ad hoc networks when using directional antennas," in *Proceedings of the 12th IEEE International Conference on Network Protocols, 2004. ICNP 2004.*, 2004, pp. 283–292.

[6] H. Assasa and J. Widmer, "Extending the ieee 802.11ad model: Scheduled access, spatial reuse, clustering, and relaying," 06 2017, pp. 39–46.

[7] "Ieee standard for information technology–telecommunications and information exchange between systems–local and metropolitan area networks–specific requirements-part 11: Wireless lan medium access control (mac) and physical layer (phy) specifications amendment 3: Enhancements for very high throughput in the 60 ghz band," *IEEE Std 802.11ad-2012 (Amendment to IEEE Std 802.11-2012, as amended by IEEE Std 802.11ae-2012 and IEEE Std 802.11aa-2012)*, pp. 1–628, 2012.