Implementation and Evaluation of a WLAN IEEE 802.11ay Model in Network Simulator ns-3

Hany Assasa*†
IMDEA Networks Institute
Madrid, Spain
hany.assasa@imdea.org

Nina Grosheva*
IMDEA Networks Institute and
Universidad Carlos III de Madrid
Madrid, Spain
nina.grosheva@imdea.org

Tanguy Ropitault
Associate, National Institute of
Standards and Technology
Prometheus Computing LLC
Sylva, NC, USA
tanguy.ropitault@nist.gov

Steve Blandino
Associate, National Institute of
Standards and Technology
Gaithersburg, MD, USA
steve.blandino@nist.gov

Nada Golmie
National Institute of Standards and
Technology
Gaithersburg, MD, USA
nada.golmie@nist.gov

Joerg Widmer
IMDEA Networks Institute
Madrid, Spain
joerg.widmer@imdea.org

ABSTRACT

The IEEE Task Group ay has recently defined new physical and medium access control specifications to design the next-generation 60 GHz wireless standard IEEE 802.11ay. Built upon the predecessor IEEE 802.11ad, IEEE 802.11ay introduces various technological advancements such as Multiple-Input and Multiple-Output (MIMO) communication, channel bonding/aggregation, and new beamforming techniques to offer unprecedented performance with 100 Gbit/s of throughput and ultra-low latency. Such performance paves the way for new emerging wireless applications such as millimeter-wave distribution networks, data center inter-rack connectivity, mobile offloading, augmented reality/virtual reality, and 8K video streaming. Studying and analyzing these new use-cases is of paramount importance and demands high fidelity network-level simulator due to the scarcity and cost of real IEEE 802.11ay test-beds.

In this paper, we present our implementation of the IEEE 802.11ay standard in the network simulator ns-3. Our implementation captures the specifics of IEEE 802.11ay operations such as the 802.11ay frame structure, channel bonding, new beamforming training procedures, quasi-deterministic MIMO channel support, and single-user MIMO and multi-user MIMO beamforming training. We also validate and demonstrate the performance of the aforementioned techniques by simulations. The code for our simulation model is publicly available.

CCS CONCEPTS

• **Networks** → **Network simulations**; **Wireless local area networks**.

*Both authors contributed equally to this work.
†The author is currently affiliated with Pharrowtech.

WNS3 2021, June 23–24, 2021, Virtual Event, USA

ACM ISBN 978-1-4503-9034-7/21/06…$15.00
https://doi.org/10.1145/3460797.3460799

KEYWORDS

Millimeter Wave, IEEE 802.11ay, 60 GHz, WiGig, MIMO, ns-3, Simulations

ACM Reference Format:
Hany Assasa, Nina Grosheva, Tanguy Ropitault, Steve Blandino, Nada Golmie, and Joerg Widmer. 2021. Implementation and Evaluation of a WLAN IEEE 802.11ay Model in Network Simulator ns-3. In *2021 Workshop on ns-3 (WNS3 2021), June 23–24, 2021, Virtual Event, USA.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3460797.3460799

1 INTRODUCTION

The Millimeter-Wave (mmWave) band has become immensely popular in the recent past. Many mobile network operators around the world started rolling out 5G mobile systems in the mmWave spectrum to alleviate the wireless capacity crunch. Besides, consumer-grade devices increasingly include mmWave support. The IEEE 802.11ad standard [8], introduced in 2012, was the first Wireless Local Area Network (WLAN) standard to provide Medium Access Control (MAC) and physical (PHY) layer specifications for wireless networking in the unlicensed 60 GHz band. Despite the technical achievement that IEEE 802.11ad represented at its release, this standard did not fully exploit the vast capacities of the 60 GHz band. Many emerging wireless applications such as mmWave distribution networks, uncompressed content streaming for augmented reality/virtual reality technologies, and dense network deployments cannot easily be addressed with IEEE 802.11ad. The main reasons lie in the fact that first, the standard was not designed for network scalability and second, it does not exploit advanced PHY layer technologies such as Multiple-Input and Multiple-Output (MIMO) and channel bonding that can boost the performance and reliability by several orders of magnitude. Implementing these PHY layer technologies is challenging due to the wide communication bandwidth in the mmWave band which exacerbates linear and non-linear impairments at the Radio Frequency (RF) devices. However, the recent advancements in the design and fabrication of mmWave electronics paved the way towards high performance, robust, low-power, and low-cost RF Integrated Circuits.

Figure 1: EDMG Waveform

This motivated the WiFi alliance to form the Task Group ay in 2015 to define the next-generation mmWave standard, named IEEE 802.11ay [9]. The following design factors were taken into account during the standardization phase: i) the standard must support a throughput of at least 20 Gbit/s, ii) it must maintain backward compatibility with IEEE 802.11ad, and iii) it must extend the set of possible use cases and scenarios by introducing novel solutions at the MAC and PHY layers. Most of these requirements are achieved thanks to the incorporation of advanced physical layer solutions that are predominant in wireless systems operating at sub-6 GHz. These solutions include MIMO, channel bonding and aggregation, fast beamforming training, and multi-user transmission. At the time of writing, no IEEE 802.11ay compliant Commercial Off-the-Shelf (COTS) devices or network-level simulators exist which hinders research progress and innovation. In this work, we fill this gap by introducing our IEEE 802.11ay implementation in the popular network simulator ns-3. The main contributions of our paper are as follows:

- We upgrade our ns-3 IEEE 802.11ad model [2–4] to support IEEE 802.11ay. This includes the 802.11ay frame structure, Modulation and Coding Schemes (MCSs), channelization, and error-model.
- We add support for all Enhanced Directional Multi-Gigabit (EDMG) Training (TRN) field variants.
- We extend our Quasi-Deterministic (Q-D) channel model to support MIMO communication.
- We introduce a MIMO analog beamforming training procedure for both Single-User MIMO (SU-MIMO) and Multi-User MIMO (MU-MIMO) cases. Additionally, we implement the SU-MIMO channel access procedure.
- Finally, we make our implementation publicly available.

2 BACKGROUND ON IEEE 802.11AY

In this section, we briefly present the major new features of the PHY and MAC layers of the IEEE 802.11ay standard.

2.1 EDMG Waveform

Figure 1 depicts the EDMG frame format. To maintain backward compatibility with IEEE 802.11ad, the EDMG frame reuses both the Directional Multi-Gigabit (DMG) preamble and DMG header fields. Thus, the EDMG frame is divided into two parts. The first part, referred to as the Non-EDMG portion, comprises a Legacy-Short Training Field (L-STF), Legacy-Channel Estimation Field (L-CEF), and legacy-header fields and is recognizable by DMG devices. The second part, which is known as the EDMG portion, contains all the fields that are recognized by EDMG Stations (STAs), including the EDMG STF and CEF fields and the new EDMG headers.

Similar to IEEE 802.11ad, IEEE 802.11ay supports three physical layer frame types: Control, Single Carrier (SC), and Orthogonal Frequency Division Multiplexing (OFDM). The Control PHY is dedicated to the transmission of management and control frames such as DMG beacons and Beamforming Training (BFT) frames. Thus, it is designed to be robust for communication under low Signal-to-Noise Ratio (SNR) conditions.

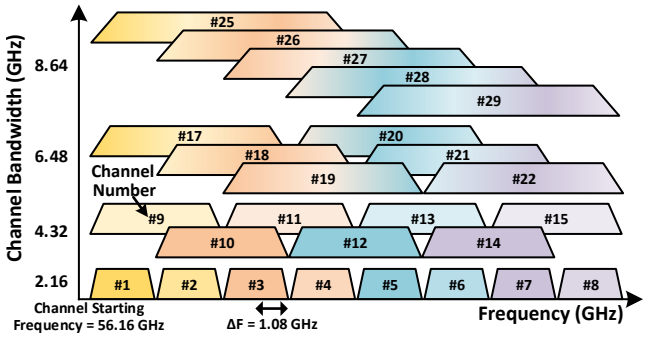

Figure 2: EDMG Channel Configurations

For data communication, either EDMG SC or EDMG OFDM can be used. The EDMG SC PHY defines an expanded set of MCSs (1 to 21) with a maximum PHY throughput of 8085 Mbit/s per spatial stream over a single channel. Likewise, EDMG OFDM specifies 20 EDMG MCSs with a maximum throughput of 8316 Mbit/s. The standard mandates the support of EDMG SC mode MCSs 1 to 5 and 7 to 10 with a single spatial stream, while EDMG OFDM is optional.

2.2 Channel Configuration

In IEEE 802.11ad, the 60 GHz band covers operation from 57 GHz to 64 GHz divided into four channels of 2.16 GHz. Communication at this frequency range suffers from high oxygen absorption which limits the communication range. With the growing interest in Fixed Wireless Access (FWA) deployments and the adoption of the unlicensed mmWave band for backhauling and fronthauling, the Federal Communications Commission decided to double the available bandwidth to cover 57 GHz to 71 GHz, providing a total of 14 GHz of unlicensed spectrum. The new frequency range between 64 GHz and 71 GHz does not suffer from high oxygen absorption which makes it suitable for backhaul applications where long-range communication is needed.

Figure 2 shows the possible channel configurations for IEEE 802.11ay. IEEE 802.11ay supports operation in eight 2.16 GHz channels. To increase the data rate further, IEEE 802.11ay allows bonding a contiguous set of channels. A maximum of four channels can be bonded which results in a channel width of 8.64 GHz. The standard mandates the support of two bonded channels (4.32 GHz).

2.3 Beam Refinement Protocol

IEEE 802.11ad introduced the Beam Refinement Protocol (BRP) to refine the beams obtained from the BFT in the Sector Level Sweep (SLS) phase. The BRP appends a special element, called the TRN field, at the end of the packet to perform fast beam switching across multiple narrow beam patterns within the same packet. IEEE 802.11ad mandates that any signal transients that occur due to the change of a beam pattern must settle within 36 ns. Building an RF Integrated Circuit with such specifications is challenging and requires an optimized analog and digital architecture. Due to these constraints, many COTS devices either omit BRP support or implement a proprietary version with a relaxed switching time. To

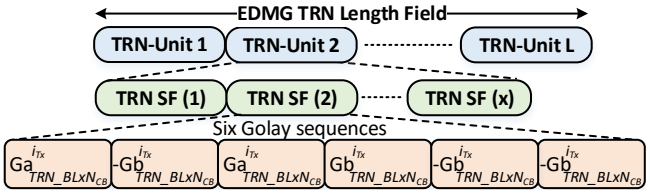

Figure 3: EDMG TRN Field Structure

address this, IEEE 802.11ay redesigned the TRN field to cope with end-devices with heterogeneous hardware.

Figure 3 shows the EDMG TRN field structure. A TRN field is composed of a variable number of TRN-Units. Each TRN Unit in turn contains multiple TRN subfields where a single TRN subfield contains six Golay sequences. IEEE 802.11ay introduces a variable size of the Golay sequence that can be configured by the user and additionally, in the case of channel bonding, depends on the number of continuous channels. Golay sequences have very robust correlation properties which make them suitable for channel estimation. IEEE 802.11ay defines a unique orthogonal set of Golay sequences for each space-time stream (i_{Tx}) to facilitate channel estimation for MIMO communication.

2.4 MIMO Communication

In IEEE 802.11ad, even though a DMG STA can have multiple Phased Antenna Arrays (PAAs) connected to its RF chain, only a single PAA can be used at a time which results in a single stream transmission. Instead, IEEE 802.11ay supports MIMO for a multi-fold increase in throughput. IEEE 802.11ay supports concurrent transmission and reception of up to eight spatial streams at the same time and over the same frequency. The standard mandates the support of analog RF precoding for MIMO communication. In this mode, PAAs can synthesize a narrow beam pattern to create a spatial channel for each stream. However, depending on the quality of the phase shifters and the geometry of the PAA, generating a pencil beam patterns with low inter-stream interference is not always feasible. To this end, IEEE 802.11ay also supports a hybrid analog and digital beamforming protocol to compensate for the deficiencies of analog beamforming through digital precoding, and achieve higher MIMO gains.

IEEE 802.11ay implements two MIMO variants: Single-User (SU)-MIMO allows transmitting and receiving multiple spatial streams (up to eight) between two devices, whereas with downlink Multi-User (MU)-MIMO, an Access Point (AP) can transmit different spatial streams to multiple users (up to 8) at the same time.

3 IMPLEMENTATION

We now present the design and the implementation details of our IEEE 802.11ay model in ns-3. It is publicly available on GitHub [1].

3.1 IEEE 802.11ay Framing

As presented in Section 2.1, IEEE 802.11ay introduces a new set of MCSs for both EDMG SC and EDMG OFDM with the addition of new coding rates. Our implementation supports all of these new MCSs. Besides, we provide a detailed PHY layer model for transmitting and receiving different fields in the EDMG Physical Layer Convergence Protocol (PLCP) frame. To ensure accurate simulations, we integrate IEEE 802.11ay SNR to Bit Error Rate (BER) lookup tables (LUTs) generated by the IEEE 802.11ay link-level simulator described in [10].

3.2 EDMG TRN Field

We implemented the flexible and configurable TRN field structure presented in Section 2.3. Additionally, we incorporated the corresponding state machines for transmitting and receiving all variants such as EDMG BRP-TX, EDMG BRP-RX, and EDMG BRP-RX/TX. The EDMG BRP-RX/TX frame is used for transmit and receive beamforming training in the same packet. This TRN structure is newly introduced in IEEE 802.11ay and is used for both Single-Input and Single-Output (SISO) and MIMO BFT. Due to space constraints, in Figure 4 we show only the state-machine for transmitting EDMG BRP-TX and EDMG BRP-RX frames, where during the transmission of EDMG BRP-RX frames the grey blocks are omitted and number of training subfields in a Unit M is set to 10.

As seen on Figure 4, the EDMG TRN field is composed of L TRN Units. In the case of BRP-RX frames, each Unit is composed of 10 subfields used for receive training. Otherwise, each Unit includes P subfields transmitted with the same beampattern as the preamble (that can be used for synchronization or channel estimation) and M subfields used for beamtraining. IEEE 802.11ay allows for N consecutive subfields to be transmitted with the same beampattern. The complete structure of the different types of BRP frames is explained in [7].

3.3 MIMO Q-D Channel Generation

In [4], we presented the Q-D channel model of our IEEE 802.11ad implementation. The channel realizations were generated by the National Institute of Standards and Technology (NIST) Q-D Channel Realization Software [5], which is a full 3D ray-tracing model that captures the geometrical properties of the channel for each point-to-point pair. The software generates a 3-D multi-point to multi-point double directional channel Impulse Response (CIR) providing the magnitude, phase, and time of arrival, Direction of Departure (DOD), and Direction of Arrival (DOA) of individual propagation paths between multiple points in space. For MIMO channels, we augmented the NIST Q-D Channel Realization Software to generate the point-to-point CIR not only between device pairs, but also between the devices' PAA pairs.

3.4 MIMO Operation

We extended the `QdPropagationEngine` class to include a MIMO engine that handles the calculation of the received signal power whenever a transmission is initiated with more than one active PAA. Our approach avoids the scheduling of multiple events for the different streams transmitted to guarantee the same simulation scalability as SISO. On the transmitter side, a single transmission event is scheduled and the transmit power is allocated equally between the transmit PAAs. On the receiver side, the MIMO engine uses the MIMO Q-D channel realizations provided by the NIST Q-D Channel Realization Software to calculate the received signal power for each pair of active transmit and receive PAAs. The `DmgWifiPhy` class then receives a list of RX signal powers and handles the event

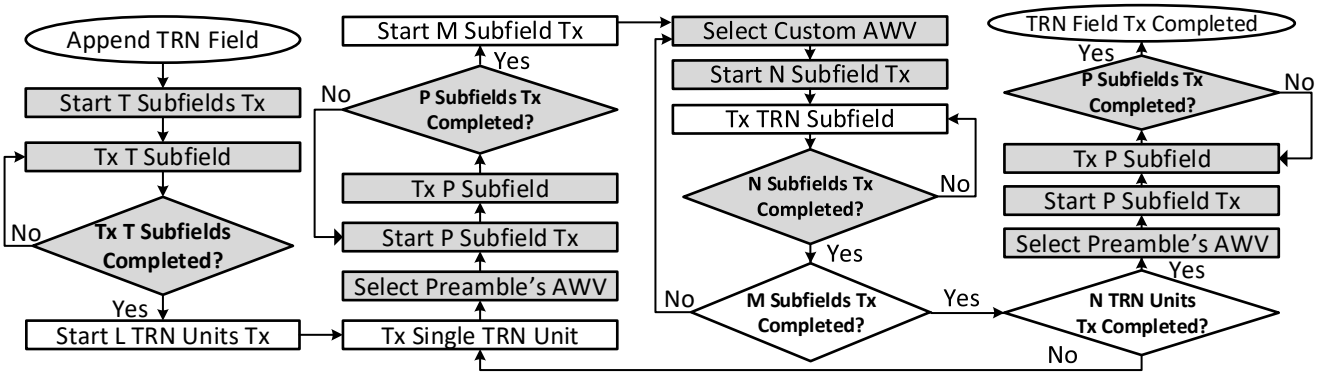

Figure 4: EDMG BRP-TX & EDMG BRP-RX Transmit State Machine Implementation

reception according to the type of MIMO transmission (e.g., data, beamforming training, etc.).

In the case of SU-MIMO data communication, a packet decoding operation is scheduled as explained in Section 3.6. However, for BRP packets transmitted during the MIMO BFT procedures, a different approach is necessary. The standard specifies that these packets are transmitted using spatial expansion, i.e., a single space-time stream is mapped to all active transmit chains with a relative cyclic shift between the different chains. This allows the receiver to separate signals coming from the different transmit PAAs and removes unintended beamforming effects. For simplicity, in our implementation the effect of spatial expansion is modeled by only decoding the stream with the highest received power and we assume that the cyclic shift diversity is sufficient to remove the interference from the other received streams. The decoding of the packet then follows the standard SISO procedure. The TRN field of the BRP packets is also transmitted in MIMO mode and is composed of orthogonal waveforms. This orthogonal design allows to train multiple transmit and receive antennas simultaneously by extracting the TRN subfield of each stream without any interference. Therefore, for MIMO TRN subfields, we can calculate the SNR of each received stream. These values are calculated without taking into account any inter-stream interference and are equivalent to SISO transmissions. Additionally, we add the possibility to calculate the Signal-to-Interference-plus-Noise Ratio (SINR) values of each TRN subfield. These values are calculated by adding the received power from the other active TX antennas as inter-stream interference. We use the SNR values in the SISO phase of SU-MIMO BFT in order to get accurate measurements for the SISO performance, and we use the SINR later in the MIMO phase of SU-MIMO and MU-MIMO BFT to evaluate the effects of inter-stream interference.

3.5 MIMO Beamforming Training

MIMO communication involves using multiple transmit and receive PAAs to transmit data in several spatial streams. To successfully establish independent streams, it is crucial to minimize the inter-stream interference to achieve sufficient per-stream SINR for data decoding. To this end, IEEE 802.11ay introduces MIMO BFT. MIMO BFT is a very challenging task since an exhaustive evaluation of all the possible PAA stream configuration combinations is not viable in real-world MIMO implementations. For example, a small codebook with 27 predefined sectors in a 2x2 MIMO setup would already require testing over half a million combinations.

IEEE 802.11ay decided to decouple MIMO BFT in two phases to overcome this problem: the SISO phase and the MIMO phase. The SISO phase aims to find the optimal SISO BFT for every SISO transmit/receive PAA pair of the MIMO communication. Even though these results do not provide an estimation of the inter-stream interference, they can be used to identify a promising subset of candidates to evaluate in the MIMO phase. In the subsequent MIMO phase, the different transmit and receive MIMO candidate combinations are tested and the actual MIMO performance with inter-stream interference is measured.

The selection of candidates to test in the MIMO phase is implementation specific and not defined by IEEE 802.11ay. Thus, for the transmit training, we developed a custom approach based on [6], which suggests assigning a joint-beam score to different beam pattern combinations and selecting the top K combinations as MIMO phase candidates. In our implementation, the joint-beam score is the sum of the individual transmit beam pattern SNRs obtained in the SISO phase. The implementation can be easily extended to other selection algorithms. The list of transmit candidates is trained in the MIMO phase. Each candidate is comprised of a TX configuration for each PAA involved in the MIMO training.

At the receiver side, the measurements at one RX PAA are independent of the configuration of the other RX PAAs. Therefore, instead of testing specific RX combinations, it is possible to just test each RX sector once and then, in post processing, determine the performance of different combinations by combining the measurements taken at the different PAAs. We thus implement a simultaneous sweeping with all PAAs across all sectors for the receive training in the MIMO phase. This greatly improves the scalability as the overhead of the receive training is determined by the number of predefined sectors in the codebook and does not increase with the number of PAAs being trained.

Additionally, in the MIMO phase, we implement an option to refine the beam selection by testing different Antenna Weight Vectors (AWVs) for each sector. As accurate estimation of the inter-stream interference is crucial to this phase, if this option is activated, all possible combinations of transmit AWVs are tested. The number of combinations increases exponentially with the number of active PAAs and therefore this option improves the accuracy of the chosen beams but reduces the scalability of the MIMO phase training.

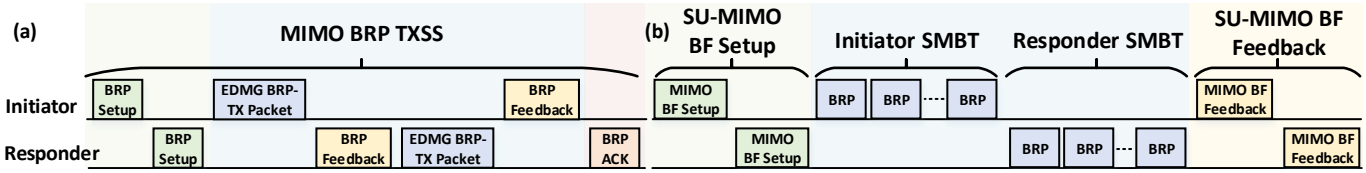

Figure 5: SU-MIMO Beamforming Training Phases: (a) SISO Phase; (b) MIMO Non-reciprocal Phase

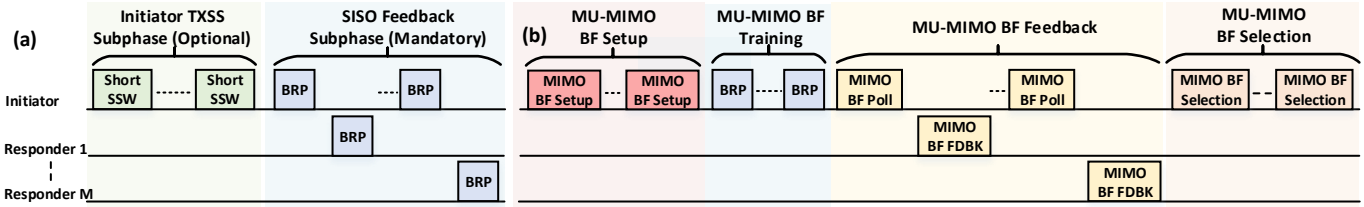

Figure 6: MU-MIMO Beamforming Training Phases: (a) SISO Phase; (b) MIMO Non-reciprocal Phase

After the MIMO phase is completed, it is necessary to rank the performance of the different combinations tested and determine the optimal MIMO configuration. To this end, we choose the combinations that maximize the minimum per stream SINR as it maximizes the probability that multiple spatial streams can be established.

It is important to note that in our implementation, we make no assumptions about the transmit and receive PAA pairs that establish the streams. Instead, all possible pairs are tested and the optimal combination is selected. Additionally, we added traces to allow the user to obtain the full set of SISO and MIMO phase measurements, as well as the chosen lists of TX candidates by our selection algorithm. In this way, the user can gain insights into the MIMO performance and evaluate the MIMO BFT algorithms.

We implemented standard-compliant SU-MIMO and MU-MIMO BFT algorithms. IEEE 802.11ay specifies that the SISO Feedback can be obtained from a previous SISO BFT or an optional new SISO Transmit Sector Sweep (TXSS) can be performed. In both algorithms, we choose to support the SISO TXSS subphases to guarantee the most-up-to-date SISO Feedback, as in this case the training is executed just before the MIMO phase. Additionally, the MIMO phase can be non-reciprocal or reciprocal, depending on whether the STAs involved in the training support antenna pattern reciprocity, i.e., the transmit antenna configurations are the same as the receive antenna configurations. For now, we support the non-reciprocal MIMO phase as it must be supported by all MIMO capable STAs and can also be used in reciprocal scenarios. Below we discuss the specifics of the SU-MIMO and MU-MIMO algorithms we implemented.

3.5.1 SU-MIMO Beamforming Training. The SU-MIMO BFT algorithm enables training between two SU-MIMO capable devices. It includes training of the transmit and corresponding receive antenna configurations for both devices involved, which means that after the conclusion of the BFT SU-MIMO communication can be established in both directions.

Figure 5 shows our SU-MIMO BFT algorithm implementation. As explained above, it includes both the full SISO phase with the training subphases and the non-reciprocal MIMO phase.

In the SISO phase, only transmit training is performed using BRP packets with Transmit Training (TRN-T) subfields transmitted and received with multiple active PAAs. As explained in Section 3.4, the orthogonal design of the MIMO TRN field in these packets allows us to determine the SNR values of each transmit chain without considering any inter-stream interference. In this way, multiple PAAs can be simultaneously trained which significantly reduces the training duration and increases the scalability as the number of PAAs being trained increases.

The MIMO phase, on the other hand, involves both transmit and receive training of MIMO combinations. This is done with BRP packets with TRN-R/T subfields, which enable simultaneous transmit and receive training. The same transmit configuration is kept for as many TRN Units as the Responder has requested for receive training. During the reception of these Units, the Responder switches the RX configuration at the start of each TRN subfield. As we explained in Section 3.4, in this phase we record the calculated SINR values that allow us to estimate the inter-stream interference.

3.5.2 MU-MIMO Beamforming Training. The MU-MIMO protocol, shown in Figure 6, is conceptually very similar to the SU-MIMO BFT protocol , with two main differences. First, during the MU-MIMO BFT an Initiator trains with multiple Responders from a MU group, requiring a modification of the Feedback phases to a poll and response format. Second, IEEE 802.11ay only defines MU-MIMO transmissions in the downlink direction and performs only transmit training for the Initiator and receive training for the Responders.

Additionally, the transmit training in the SISO phase is performed with Short Sector Sweep (SSW) packets transmitted and received in SISO mode, instead of MIMO TRN-T subfields. This is because the Initiator is training with multiple Responders and it is not possible to guarantee that all of them will be able to receive the BRP packets. In order to reduce the training time, the new short SSW frames are used, instead of legacy SSW frames. The short SSW frame is a PHY layer frame that is 6 bytes long compared to 26 bytes for the legacy SSW which results in a 31% reduction in the transmission time. We add support for these frames by enabling the transmission of WiFi packets without a MAC header.

The MU-MIMO training is performed using TRN-R/T subfields, similar to SU-MIMO. However, it requires an additional subphase called MU-MIMO BF Selection, where the Initiator informs the MU group of the Responders and optimal MIMO configurations

Table 1: Simulations Parameters

Parameter Name	Parameter Value
Application Type	OnOffApplication
Payload Size	1472 Bytes
Transport Protocol	UDP
Aggregation Type	A-MSDU and A-MPDU
A-MSDU Max. Size	7935 Bytes
A-MPDU Max. Size	4 194 303 Bytes
Block ACK Size	1024 Frames
Number of Transmit Sectors	27 Sectors
Sector Azimuth Steering Angles	-80°:20°:80°
Sector Elevation Steering Angles	-45°, 0°, 45°
Transmit Power	10 dBm

that have been selected for MU-MIMO communication. This allows the Responders to use the correct receive configuration when MU-MIMO transmissions take place.

3.6 SU-MIMO Channel Access Procedure and Data Transmission

IEEE 802.11ay defines various methods for MIMO channel access before data transmission.As MU-MIMO data transmission is left for future work, we only discuss the SU-MIMO implementation. We implement a Ready-to-Send (RTS)/DMG Clear-to-Send (CTS) mechanism where a control trailer is added to the RTS and DMG CTS frames. The control trailer contains signaling regarding the SU-MIMO configuration used for data transmission, allowing the STAs to set up the transmit and corresponding receive antenna configurations previously trained.

Moreover, for the data transmission, we extend the `DmgWifiMac`, `MacLow`, `DmgWifiPhy` and `InterferenceHelper` classes to support transmission and decoding of MIMO packets. In the `Interference Helper`, we calculate the per stream SINR values that take into account the inter-stream interference and use this to determine the per-stream packet success rate. Analogous to the calculation of the chunk success rate, the success rate for the packet is equivalent to the multiplication of the per-stream Packet Success Rates (PSRs).

4 EVALUATION

In this section, we evaluate and validate our IEEE 802.11ay implementation in ns-3. All our simulation scenarios utilize the Q-D channel model. Simulation parameters are summarized in Table 1. All the devices in the network use a 2x8 element Uniform Rectangular Array (URA) PAA which yields a narrow beam in the azimuth plane, and a wide beam in the elevation plane. We use the full Aggregate MAC Service Data Unit (A-MSDU) and Aggregate MAC Protocol Data Unit (A-MPDU) aggregation defined by IEEE 802.11ay. In order to support the expanded A-MPDU aggregation we implement the EDMG Compressed Block Acknowledgement that allows to acknowledge the reception of up to 1024 MPDUs.

4.1 Achievable Throughput

In this simulation, we evaluate the maximum achievable throughput for the IEEE 802.11ay protocol for all the EDMG MCSs with various channel widths. Our scenario consists of two IEEE 802.11ay devices with a Line-Of-Sight (LOS) link with a distance of one meter. We configure the two devices to use the optimal beam pattern thus ensuring a high SNR value that prevents any packet loss. To eliminate beamforming training overhead, we install `DmgAdhocWifiMac` which is an experimental MAC layer implementation that facilitates studying PHY layer features without adding the complexity of the full MAC protocol. This MAC implementation allocates the whole Beacon Interval (BI) for data transmission.

Figure 7 depicts our simulation results for EDMG SC and EDMG OFDM PHYs. To exclude the overhead of each layer in the protocol stack, we measure the throughput at the application layer. We observe that the maximum achievable throughput with four bonded channels is around 29.6 Gbit/s for EDMG SC and 31.25 Gbit/s for EDMG OFDM. We notice a degradation in the throughput for EDMG-MCS-17. This is because EDMG-MCS-17 uses a 64-QAM modulation scheme with a coding rate of 1/2, which results in fewer data bits per SC block compared to EDMG-MCS-16. It is worth mentioning that this might cause issues with Rate Adaptation Algorithms (RAAs) as they would expect a monotonic increase in throughput when increasing the MCS.

The throughput obtained in this simulation considers an ideal scenario where we have neither collision on the wireless medium nor packet loss. In a real network, the throughput will be lower due to i) the overhead imposed by different channel access periods in the BI, ii) the usage of the RTS/CTS handshake protocol, and iii) frequent link maintenance through BFT in the Data Transmission Interval (DTI) access period. The impact of the latter depends mainly on the size of the codebook and the number of PAAs.

4.2 SU-MIMO Beamforming Training Validation

The scenario to validate our SU-MIMO implementation consists of one AP and one STA, each equipped with two PAAs separated by 3cm along the x-axis, deployed in a 5m × 10m × 3m room as depicted in Figure 9. Each PAA is connected to a separate transmit chain which allows for a maximum of two spatial streams.

Figure 8 depicts the results from the different phases of our SU-MIMO BFT algorithm between the AP (TX) and the STA (RX). The SISO phase measurements in Figure 8 (a) show the SNR of the different transmit sectors from both TX PAAs measured at both RX PAAs. Since the PAAs separation distance is small, we can observe that the SNRs from the same transmit sector at both receiver's PAAs are very similar in most cases. The SISO results then serve as input to our selection algorithm that selects the top K combinations as shown in Figure 8 (b). The list of K candidates is tested in the MIMO phase shown in Figure 8 (c), which results in a set of SINR measurements. For this scenario, we use the top K=85 combinations tested, as we observed that this value offers a good compromise between scalability and accurate SU-MIMO configuration. In Figure 8 (d) we present a heatmap of the minimum per stream SINR for each tested candidate. On the x-axis, we show the different TX candidates according to their ranking by the selection algorithm, the first column representing the candidate with the highest joint SNR. On the y-axis, we present the different receive combinations tested. As explained in Section 3.5, we can determine the SINR for all possible receive combinations and we present them sequentially

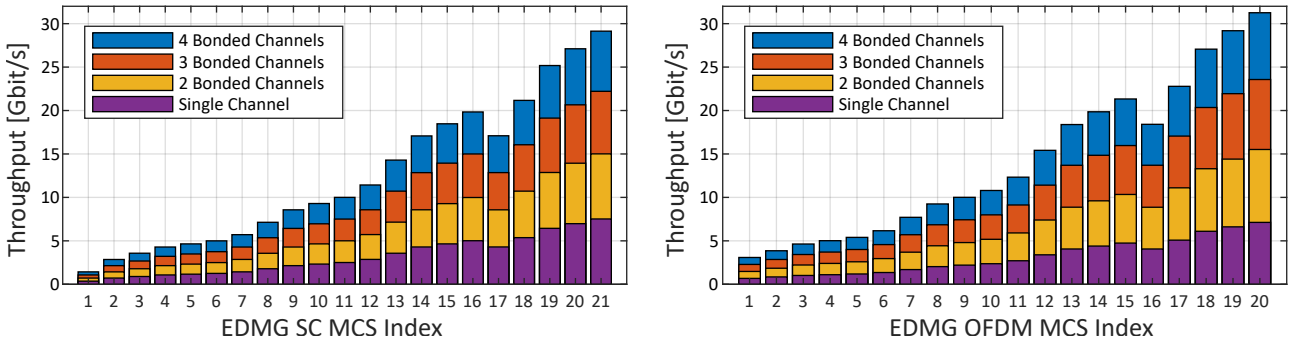

Figure 7: EDMG MCSs Throughput for Different Channel Sizes

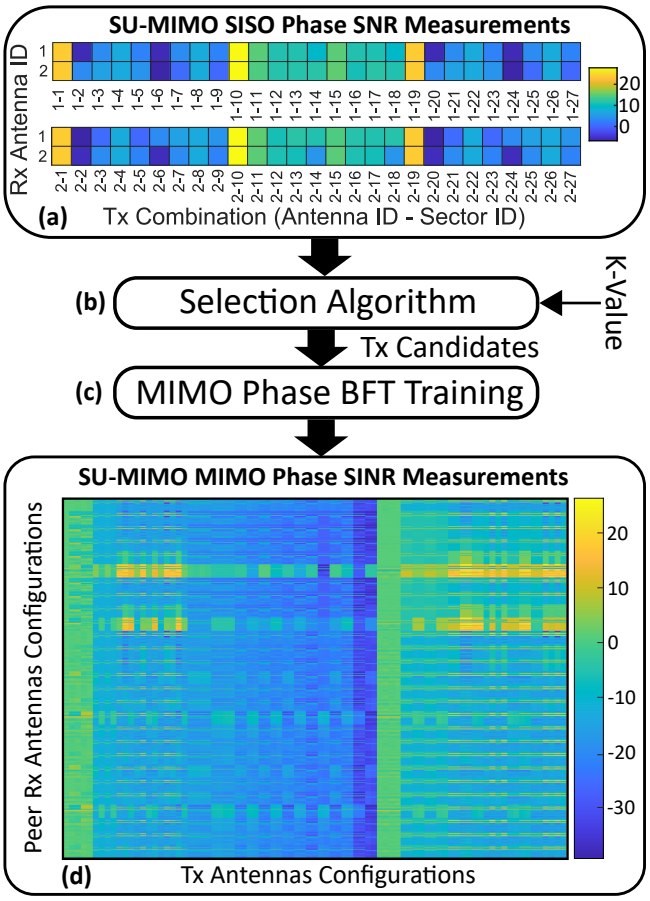

Figure 8: MIMO Beamforming Training Results

(the bottom row representing (RX PAA 1 - Sector 1, RX PAA 2 - Sector 1) and the top row representing (RX PAA 1 - Sector 27, RX PAA 2 - Sector 27). We can see that the highest-ranked candidates (leftmost columns) experience low SINR. For these candidates, both PAAs have beam patterns that utilize the LOS path as it gives the highest SISO SNR. However, when used for MIMO communication, such a combination results in high inter-stream interference due to the small PAA separation. The candidates with the lowest measured SINR similarly suffer from high inter-stream interference, but additionally have lower received signal strength as they utilize reflected paths. This shows the significance of the MIMO phase, as the optimal SISO configurations can sometimes result in poor MIMO performance. Additionally, we can observe no obvious pattern in the SINR measurements for the different configurations tested. This implies that it can be extremely challenging to predict the MIMO performance from the SISO Feedback and that the selection of good candidates for the MIMO phase is crucial to the overall functioning of the MIMO BFT algorithms. As mentioned in Section 3.5, our implementation was designed to be able to evaluate the effect of different selection algorithms and can therefore be of crucial interest to study mmWave MIMO behavior. Finally, we observe two high SINR areas. The first area, located in the top left half allows for SINRs of around 15 dB. However, by testing a higher number of candidates we discover a second high SINR area in the top right half of the map with more optimal antenna configurations that can achieve SINRs above 20 dB.

Figure 9 shows a visualization of the best SU-MIMO configuration chosen by our BFT algorithm. We can clearly see that the first stream established, shown in Figure 9 (a), utilizes the reflections from the front and back walls and has very low gain for the LOS path and the reflections from the side-walls and the ceiling/ground. The second stream, shown in Figure 9 (b), utilizes precisely those links and receives very low interference from the front and back wall reflections. The resulting combination shown in Figure 9 (c) has very high per stream SINR of 23.52 dB and 39.25 dB respectively, validating that our BFT algorithm can successfully determine good antenna configurations for MIMO communication.

Finally, after the BFT is completed, we validate our SU-MIMO data transmission implementation using the output of the MIMO Phase BFT to setup transmit and receive antennas. The large SINR experienced by the two streams enables the use of EDMG-SC MCS-21 (8 Gbit/s). We observe an aggregate throughput of around 14 Gbit/s, validating the multi-stream transmission implementation.

4.3 MU-MIMO Beamforming Training Validation

In this scenario, we deploy one EDMG AP and two STAs in the same room as depicted in Figure 10. The AP is equipped with two RF chains, each connected to a separate PAA, while the two STAs have a single PAA. As a result, the AP can transmit two spatial streams, allowing communication with two users at the same time.

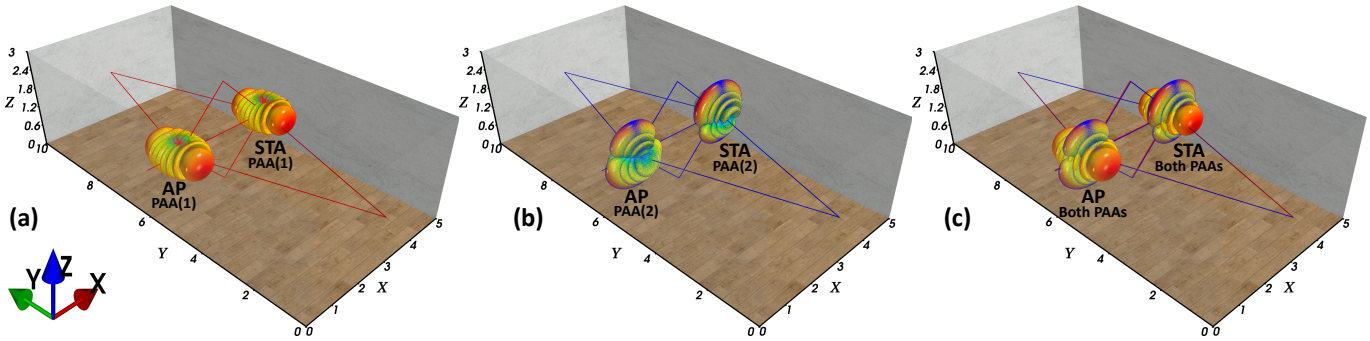

Figure 9: SU-MIMO Beamforming Training Qualitative Results: (a) PAAs Beam Patterns Corresponding to Stream 1; (b) PAAs Beam Patterns Corresponding to Stream 2; (c) Combined PAAs Beam Patterns for Streams 1 and 2

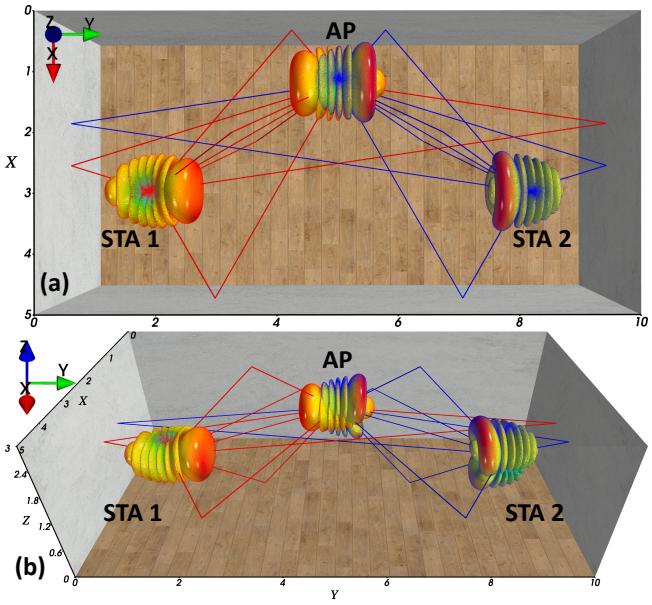

Figure 10: MU-MIMO Beamforming Training Qualitative Results (a) Top View; (b) Side View

Due to space constraints, we show only the optimal MU-MIMO configuration chosen by our algorithm in Figure 10. We can see that the high spatial separation between the STAs allows us to have two streams that utilize different multi-path components. The resulting per stream SINRs of 33.8 dB and 33.3 dB are very high and will be sufficient for MU-MIMO communication with high data rates.

5 CONCLUSIONS AND FUTURE WORK

In this paper, we presented our implementation of the IEEE 802.11ay standard in network simulator ns-3. We implemented a diverse set of MAC and PHY features including IEEE 802.11ay framing, channel bonding, EDMG BRP variants, SU-MIMO beamforming training with data transmission, and MU-MIMO beamforming training. We demonstrated the maximum achievable throughput per spatial stream for each EDMG MCS for different channel configurations. Besides, we illustrated some qualitative results for SU/MU-MIMO beamforming training and beam selection algorithm. We plan to continue improving the robustness and fidelity of our IEEE 802.11ay module. Additionally, we are working on the following features: multi-channel scheduling, MU-MIMO channel access procedure, TDD protocol, and polarization support.

ACKNOWLEDGMENTS

This work has been funded by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 861222 (MINTS) and NIST (USA) under the Measurement Science and Engineering (MSE) Research Program, grant n° 60NANB19D086.

REFERENCES

[1] Hany Assasa and Nina Grosheva. 2021. *A Collection of Open-source Tools to Simulate IEEE 802.11ad/ay WLAN Networks in Network Simulator ns-3.* https://github.com/wigig-tools/ns3-802.11ad

[2] Hany Assasa and Joerg Widmer. 2016. Implementation and Evaluation of a WLAN IEEE 802.11ad Model in ns-3. In *Proceedings of the 2016 Workshop on ns-3.* Seattle, WA, USA.

[3] Hany Assasa and Joerg Widmer. 2017. Extending the IEEE 802.11ad Model: Scheduled Access, Spatial Reuse, Clustering, and Relaying. In *Proceedings of the Workshop on ns-3.* Porto, Portugal.

[4] Hany Assasa, Joerg Widmer, Tanguy Ropitault, and Nada Golmie. 2019. Enhancing the ns-3 IEEE 802.11ad Model Fidelity: Beam Codebooks, Multi-antenna Beamforming Training, and Quasi-deterministic mmWave Channel. In *Proceedings of the Workshop on ns-3.* Florence, Italy.

[5] Anuraag Bodi, Steve Blandino, Neeraj Varshney, Jiayi Zhang, Tanguy Ropitault, Mattia Lecci, Paolo Testolina, Jian Wang, Chiehping Lai, and Camillo Gentile. 2021. *The NIST Q-D Channel Realization Software.* https://github.com/wigig-tools/qd-realization

[6] Felix Fellhauer, Nabil Loghin, Dana Ciochina, Thomas Handte, and Stephan ten Brink. 2017. Low Complexity Beamforming Training Method for mmWave Communications. In *2017 IEEE 18th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC).* 1–5. https://doi.org/10.1109/SPAWC.2017.8227683

[7] Yasaman Ghasempour, Claudio R. C. M. da Silva, Carlos Cordeiro, and Edward W. Knightly. 2017. IEEE 802.11ay: Next-Generation 60 GHz Communication for 100 Gb/s Wi-Fi. *IEEE Communications Magazine* (2017), 186–192. https://doi.org/10.1109/MCOM.2017.1700393

[8] IEEE. 2016. 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. *IEEE Std 802.11-2016 (Revision of IEEE Std 802.11-2012)* (2016).

[9] IEEE. 2020. IEEE Draft 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 2: Enhanced Throughput for Operation in License-Exempt Bands Above 45 GHz. *IEEE P802.11ay/D7.0, December 2020* (2020).

[10] Neeraj Varshney, Jiayi Zhang, Jian Wang, Anuraag Bodi, and Nada Golmie. 2020. Link-Level Abstraction of IEEE 802.11ay based on Quasi-Deterministic Channel Model from Measurements. In *2020 IEEE 92nd Vehicle Technology Conference (VTC 2020-Fall).* Victoria, B.C, Canada.