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
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 (TGay) has recently defined new physical and medium access control specifications to design the next-generation wireless standard in the 60 GHz band, the so-called IEEE 802.11ay. Build upon its 802.11ad predecessor, IEEE 802.11ay aims to offer unprecedented performance (100 Gbps throughput, ultra-low latency) by introducing various technological advancements such as multiple-input and multiple-output (MIMO) communication, channel bonding/aggregation, and new beamforming techniques. Such performance paves the way to new emerging wireless applications such as millimeter-wave distribution networks, data center inter-rack connectivity, mobile offloading, augmented reality (AR)/virtual reality (VR), 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/costs 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 11ay frame structure, channel bonding, new beamforming training procedures, quasi-deterministic MIMO channel support, and Single-User (SU)-MIMO (SU-MIMO)/ Multi-User (MU)-MIMO (MU-MIMO) beamforming training. We validate and demonstrate by simulations the performance of the aforementioned techniques. The code for our simulation model is publicly available.

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

WNS3 2021, June 23, 2021, Virtual

ACM ISBN 978-1-4503-XXXX-X/18/06…$15.00
https://doi.org/10.1145/1122445.1122456

CCS CONCEPTS

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

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, 2021, Virtual.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/1122445.1122456

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 the mmWave spectrum in their 5G mobile systems to alleviate the current wireless capacity crunch. Besides, consumer-grade devices are increasingly including mmWave support. The IEEE 802.11ad standard [9], introduced in 2012, was the first wireless local area networks (WLAN) standard to provide medium access control (MAC) and physical (PHY) specifications for wireless networking in the unlicensed 60 GHz band. Despite the technical achievement that IEEE 802.11ad represented at its release (around 6.72 Gbps throughput), this standard never fully took benefit of the vast capacities of the 60 GHz band. Many emerging wireless applications such as mmWave distribution networks, uncompressed content streaming for VR/AR technologies, and dense network deployments proved to be hardly attainable with IEEE 802.11ad. The main reasons lie in the fact that first, the standard was not designed with network scalability in mind. Then, it did not exploit advanced PHY layer technologies such as MIMO and channel bonding that can boost its performance/reliability by order of magnitudes. 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 consumption, and low-cost radio-frequency integrated circuits (RFICs).

This motivated the WiFi alliance to form the TGay in 2015 to define the next-generation mmWave standard, the so-called IEEE 802.11ay [1]. The following design factors were taken into account during the standardization phase: i) The standard must support a throughput of at least 20 Gbps. ii) It must maintain backward compatibility with IEEE 802.11ad. 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 implementation for the IEEE 802.11ay 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 [3–5] to support IEEE 802.11ay. This includes defining 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 MIMO analog beamforming training procedure for both SU-MIMO and MU-MIMO cases. Additionally, we implement SU-MIMO channel access procedure.
- Finally, we make our implementation publicly available for the research community.

The paper is structured as follows. In Section 2, we provide background on the new technologies introduced in IEEE 802.11ay with a special focus on the differences with its predecessor IEEE 802.11ad. Section 3 presents our ns-3 IEEE 802.11ay implementation, and Section 4 highlights our evaluation campaign for the proposed model. Finally, Section 5 concludes the paper.

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 is referred to as the Non-EDMG portion 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).

Similar to IEEE 802.11ad, IEEE 802.11ay supports three types of physical layers technologies: Control, Single Carrier (SC), and

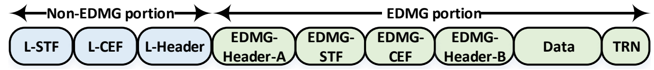

Figure 1: EDMG Waveform.

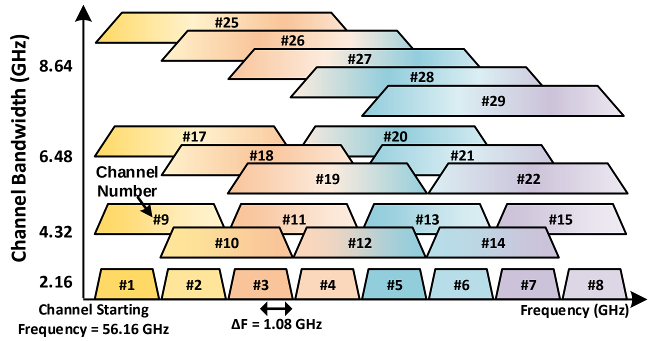

Figure 2: EDMG Channel Configurations.

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 frames. Thus, it is designed to be robust for communication in low Signal-to-Noise Ratio (SNR) conditions. All frames transmitted in this mode can be recognized and decoded by legacy DMG devices.

For data communication, either EDMG SC or EDMG OFDM can be used. The standard mandates the support of EDMG SC mode MCSs 1 - 5 and 7 - 10 with a single spatial stream. The SC PHY has an extended set of MCSs (1 to 21) with a maximum PHY throughput of 8085 Mbps per spatial stream over a single channel for a normal guard interval (GI). On the other hand, EDMG OFDM defines 20 unique MCSs with a maximum throughput of 8316 Mbps. The support of EDMG OFDM is optional.

2.2 Channel Configuration

In IEEE 802.11ad, the 60 GHz band covered 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) deployment and the adoption of the unlicensed mmWave band for backhauling and fronthauling, the FCC decided to double the bandwidth to cover from 57 GHz to 71 GHz, allowing 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 backhauling applications where long-range communication is needed.

Figure 2 shows all 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 to obtain a larger channel. A maximum of four channels can be bonded which results in 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 beamforming training (BFT) in

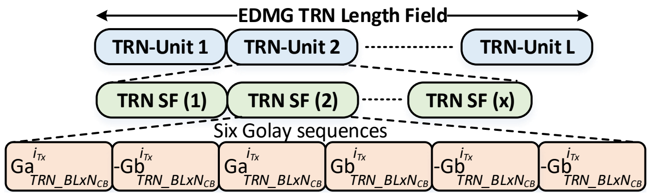

Figure 3: EDMG TRN Field Structure

the Sector Level Sweep (SLS) phase. The beam refinement protocol (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 transient that occurs due to the change of a beam pattern must settle within 36 ns. Building an RFIC with such specifications is challenging and requires an optimized analog and digital architecture. Due to these constraints, many COTS devices either omit the BRP support or implement a proprietary version with a relaxed switching time. To tackle this issue, IEEE 802.11ay performed a major redesign of the TRN field to cope with heterogeneous hardware capable end-devices.

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 SF 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. As a result, IEEE 802.11ay defines a unique and orthogonal set of Golay sequences for each space-time stream (i_{Tx}) to facilitate channel estimation during 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 utilized at a time which results in a single stream transmission. This motivated IEEE 802.11ay to adopt MIMO support as a way to increase its throughput by multi-fold. IEEE 802.11ay supports concurrent transmission and reception of 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, PAA can synthesize a narrow beam pattern to create multiple orthogonal spatial channels for each stream. However, depending on the quality of the phase shifters and the geometry of the PAA, generating a pencil beam pattern is not always feasible. Thus, IEEE 802.11ay proposes a hybrid analog and digital beamforming protocol to compensate for non-idealities in the analog domain and achieve the maximum gain of a MIMO system.

IEEE 802.11ay implements two variants of MIMO transmission. The first variant is known as SU-MIMO which allows transmitting and receiving multiple spatial streams (up to eight) between two devices. The second type is known as downlink MU-MIMO. In this type, an access point (AP) can transmit multiple spatial streams to multiples users (up to 8) at the same time.

3 IMPLEMENTATION

In the following section, we present the design and the implementation details of our IEEE 802.11ay model in ns-3. Our implementation is publicly available on GitHub [2].

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 64-QAM, a high order modulation scheme. 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 NIST 802.11ay link-level simulator [10].

3.2 EDMG TRN Field

We implement the flexible and configurable TRN field structure presented in section 2.3. Additionally, we incorporate the corresponding state machines for transmitting and receiving all variants such as EDMG BRP-TX, EDMG BRP-RX, and EDMG BRP-RX/TX. Interested readers can refer to [8] for further details on the different variants.

Figure 4 shows the various states for transmitting EDMG BRP-TX and EDMG BRP-RX frames. During the transmission of EDMG BRP-RX frame, the grey blocks are omitted and M is set to 10. The EDMG BRP-RX/TX variant is used for transmit and receive beamforming training at the same time. 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, we show only the state-machine for transmitting EDMG BRP-TX and EDMG BRP-RX.

3.3 MIMO Q-D Channel Generation

In [5], we presented the Q-D channel model which was added to our IEEE 802.11ad implementation. The channel realizations were generated by the NIST Q-D Channel Realization Software [6], 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 details of 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. To enable MIMO channel realization, the NIST Q-D channel Realization software has been augmented to allow the generation of the point-to-point CIR not only between each device pair, but also between each device's PAAs pair.

3.4 MIMO Operation

We extend 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

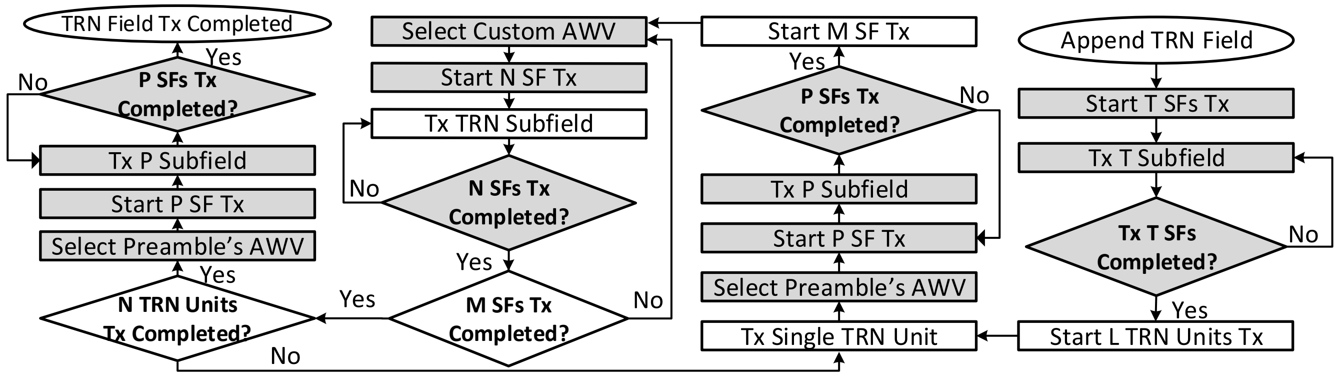

Figure 4: EDMG BRP-TX & EDMG BRP-TX Transmit State Machine Implementation.

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 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 transmit chains active 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. In our implementation, the effect of spatial expansion is modeled by only attempting to decode the stream with the highest received power, considering that the cyclic shift diversity will be sufficient to remove the interference from the other received streams. The decoding of the packet afterward 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 us to train multiple transmit and receive antennas simultaneously by extracting the TRN-SF of each stream without any interference. Therefore, for MIMO TRN SFs, 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 SF. These values are calculated by adding the received power from the other TX antennas active 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 be able to successfully establish independent streams, it is crucial to minimize the inter-stream interference to enable sufficient per-stream SINR for data decoding. To this end, IEEE 802.11ay introduced 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, using a small codebook with 27 predefined sectors in a 2x2 MIMO setup would require testing over half a million different 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 obtain 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/select a promising subset of candidates to evaluate in the MIMO phase, enabling scalability. In the following MIMO phase, the different transmit and receive MIMO candidate combinations are tested and the MIMO performance, including the inter-stream interference effect, 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 [7], which suggests assigning to different beam pattern combinations a joint-beam score and to select the MIMO phase candidates from the top K combinations. In our implementation, the joint-beam score is the sum of the individual transmit beam patterns SNRs obtained in the SISO phase. The implementation can be easily extended to other selection algorithms. The list of transmit candidates given by our algorithm is trained in the MIMO phase. Each candidate is comprised of a TX configuration for each PAA involved in the MIMO training.

For the receive training, our implementation uses a different approach. This comes from our observation that the measurements at one RX PAA are independent of the configuration of the other RX PAAs. This means that 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. Therefore, for the receive training in the MIMO phase, we implement a simultaneous sweeping with all PAAs across all sectors. This allows us to greatly improve the scalability of the MIMO phase 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.

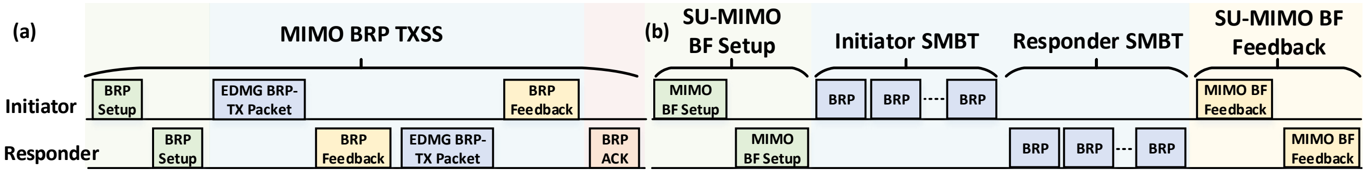

Figure 5: SU-MIMO Beamforming Training Phases: a) SISO Phase; b) MIMO Non-reciprocal Phase.

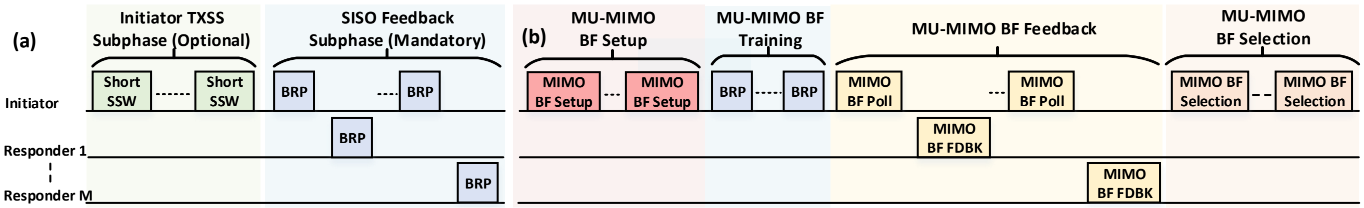

Figure 6: MU-MIMO Beamforming Training Phases: a) SISO Phase; b) MIMO Non-reciprocal Phase.

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 the accurate estimation of the inter-stream interference is crucial to the MIMO phase, if this option is activated, all possible combinations of the transmit AWVs are tested. The number of possible 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.

After the MIMO phase is completed, it is necessary to rank the performance of the different combinations tested and determine what is the optimal MIMO configuration. To this end, we choose the combinations that maximize the minimum per stream SINR as it will maximize the possibility 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 significant insight into the MIMO performance and evaluate the MIMO BFT algorithms.

We implement 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, meaning that we can consider that the transmit antenna configurations will be 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 STAs. It includes training of the transmit and corresponding receive antenna configurations for both STAs involved, which means that after the conclusion of the BFT SU-MIMO communication can be established in both directions.

In the SISO phase, only transmit training is performed using BRP packets with transmit training (TRN-T) SFs 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 SFs, 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 SF. 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.

Figure 5 shows our SU-MIMO BFT algorithm implementation.

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 described in Section 3.5.1, 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 SFs. 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 and it 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 MIMO training is performed using TRN-R/T SFs, same as for SU-MIMO. However, it requires an additional Selection Subphase where the Initiator informs the MU group of the Responders and optimal MIMO configurations that have been selected for MU-MIMO communication. This allows the Responders to activate 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 SU-MIMO channel access before data transmission can take place. We implement a RTS/DMG 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 that were previously trained.

Additionally, for the data transmission, we extend the `DmgWifiMac`, `MacLow`, `DmgWifiPhy` and `InterferenceHelper` classes to support the 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 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 utilize 2x8 elements Uniform Rectangular Array (URA) PAA which yields a narrow beam in the azimuth plane, and a wide beam in the elevation plane.

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 separated apart by one meter and have a Line-of-sight (LOS) link. We configure the two devices to use a broadside 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 MAC layer protocol. This MAC implementation allocates the whole Beacon Interval (BI) for data transmission.

Figure 7 depicts our simulation results for both 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 can observe that the maximum achievable throughput with four bonded channels is around 29.6 Gbps for EDMG SC and 31.25 Gbps for EDMG OFDM. We notice a degradation in the throughput for EDMG-MCS-17. it is worth mentioning that this might cause issues with rate adaptation algorithm (RAA) algorithms as they would

Table 1: Simulations Parameters

Parameter Name	Parameter Value
Application Type	OnOffApplication
Payload Size	1472 Bytes
Transport Protocol	UDP
MAC Queue Size	4000 Packets
Aggregation Type	A-MSDU and A-MPDU
A-MSDU Max. Size	7935 Bytes
A-MPDU Max. Size	4194303 Bytes
PPDU Max. Duration	2 ms
Block ACK Size	1024 Frames
MAC Protocol	CSMA/CA
Codebook Type	Parametric
Number of Transmit Sectors	27 Sectors
Sectors Azimuth Steering Angles	-80°:20°:80°
Sectors Elevation Steering Angles	-45°, 0°, 45°
A-BFT Sectors	13 Sectors
Guard Interval Size	Normal
Transmit Power	10 dBm
Rx Noise Figure	10 dB
Operating Frequency	60.48 GHz (CH2)

expect a monotonic increase in throughput when increasing the MCS.

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

4.2 SU-MIMO Beamforming Training Validation

To validate our SU-MIMO implementation, our scenario is made of an AP and a STA, each one equipped with two PAAs separated by 3cm along the x-axis, deployed in a 5x10x3m room as depicted in Figure9. Each PAA is connected to a separate transmit chain which allows 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 (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 (Figure 8 (c)) resulting in a set of SINR measurements. For this scenario, we used 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 minimum per stream SINR for each candidate tested. On the x-axis, we show the different TX candidates according to

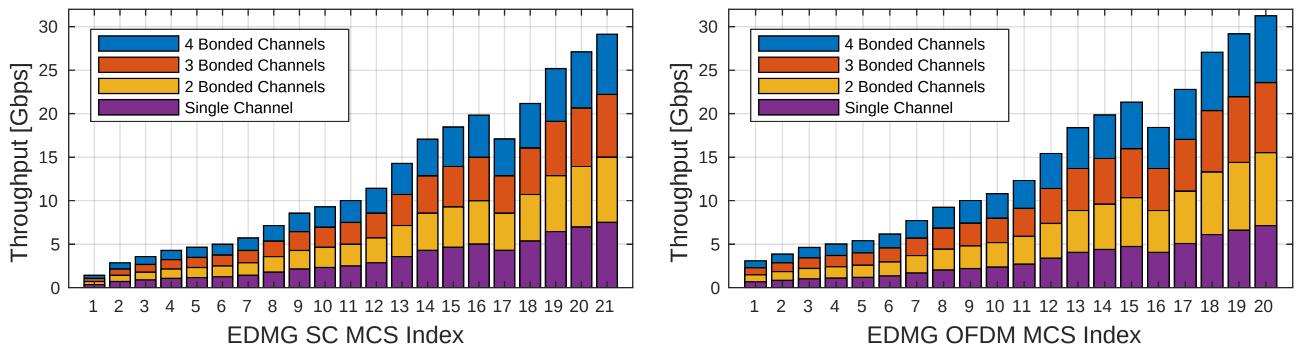

Figure 7: EDMG MCSs Throughput for Different Channel Sizes.

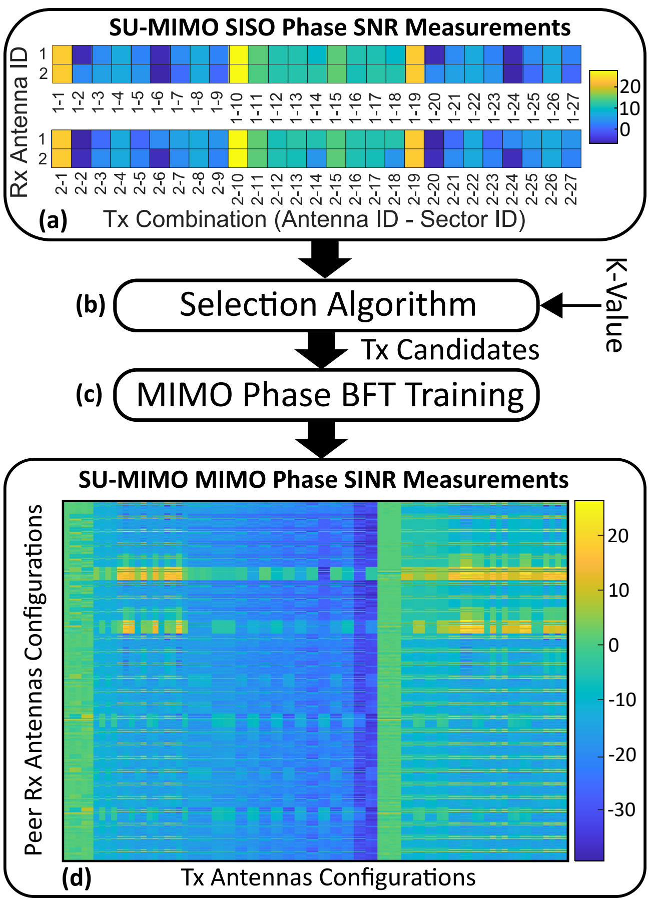

Figure 8: MIMO Beamforming Training Results.

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 (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. This shows the significance of the MIMO phase, as the optimal SISO configurations can sometimes result in poor MIMO performance. Additionally, we can observe a high diversity 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, in the top right half of the map, we can also see a high SINR area where the best SU-MIMO configuration is located.

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, (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 training to setup transmit and receive antennas. The large SINR experienced by the two streams enables the use of EDMG-SC MCS-21 (8 Gbps). We observe an aggregate throughput of around 14 Gbps, validating the multi-stream transmission implementation.

4.3 MU-MIMO Beamforming Training Validation

In this scenario, we deploy an EDMG AP and two STAs in the same room as depicted in Figure 10. The AP is equipped with two RF chains where each chain is connected to a separate PAA, while the two STAs are each equipped with a single PAA. As a result, the AP

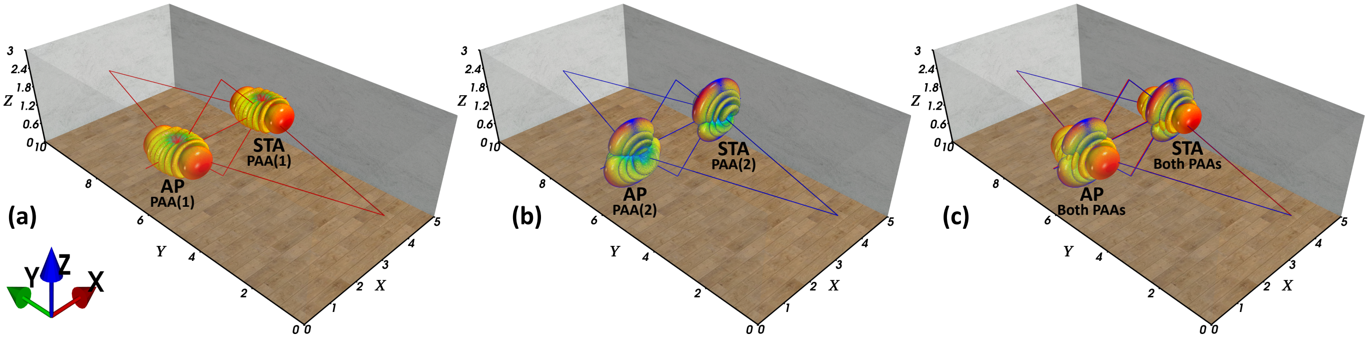

Figure 9: SU-MIMO Beamforming Training Qualitative Results: (a) PAAs Beam Patterns Corresponding to Stream 1, b)PAAs Beam Patterns Corresponding to Stream 2, (c) Combined PAAs Beam Patterns for Stream 1 and 2.

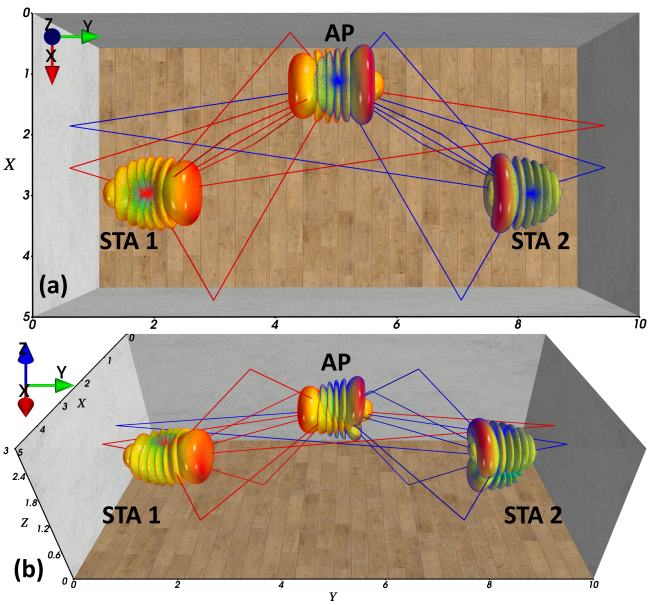

Figure 10: MU-MIMO Beamforming Training Qualitative Results (a) Top View; (b) Side View.

can transmit two spatial streams thus allowing data communication with two users at the same time.

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 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 and the National Institute of Standards and Technology (USA) under the Measurement Science and Engineering (MSE) Research Program, grant n° 60NANB19D086.

REFERENCES

[1] 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).

[2] H. Assasa and N. 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

[3] H. Assasa and J. 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.

[4] H. Assasa and J. 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.

[5] H. Assasa, J. Widmer, T. Ropitault, and N. 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.

[6] A. Bodi, S. Blandino, J. Zhang N. Varshney, T. Ropitault, P. Testolina M. Lecci, J. Wang, C. Lai, and C. Gentile. 2021. *The NIST Q-D Channel Realization Software*. https://github.com/wigig-tools/qd-realization

[7] F. Fellhauer, N. Loghin, D. Ciochina, T. Handte, and S. 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

[8] Y. Ghasempour, C. R. C. M. da Silva, C. Cordeiro, and E. W. Knightly. 2017. IEEE 802.11ay: Next-Generation 60 GHz Communication for 100 Gb/s Wi-Fi. *IEEE Communications Magazine* 55, 12 (2017), 186–192. https://doi.org/10.1109/MCOM.2017.1700393

[9] IEEE. 2012. 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* (2012).

[10] N. Varshney, J. Zhang, J. Wang, A. Bodi, and N. 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.