A Comprehensive Analysis and Performance Enhancements for the IEEE 802.11ay Group Beamforming Protocol

Nina Grosheva[⋆] [†], Hany Assasa[⋆], Tanguy Ropitault[‡], Pablo Jiménez Mateo[⋆] [†], Joerg Widmer[⋆], and Nada Golmie[§]

[⋆]IMDEA Networks Institute, Madrid, Spain
[†]Universidad Carlos III de Madrid, Madrid, Spain
[‡] Associate, National Institute of Standards and Technology Prometheus Computing LLC, Gaithersburg, MD, USA
[§]National Institute of Standards and Technology, Gaithersburg, MD, USA

***Abstract*—Millimeter-wave technology provides the necessary improvements in capacity and performance for the next generation of wireless networks. The new IEEE 802.11ay amendment extends IEEE 802.11ad to offer 100 Gbit/s connectivity in the unlicensed 60 GHz band through technical advancements such as Multiple-Input and Multiple-Output (MIMO), channel bonding and aggregation. Additionally, it offers improvements to the Beamforming Training (BFT) process in order to increase its efficiency and accuracy. One new technique defined by IEEE 802.11ay is Group Beamforming, which allows to simultaneously train all stations, and significantly reduces training overhead, especially in very dense networks. In this paper, we provide an implementation of IEEE 802.11ay in ns-3 and perform, to the best of our knowledge, the first detailed system-level evaluation of the performance of the novel IEEE 802.11ay protocol. We specifically study the performance of Group Beamforming and compare it against the legacy 802.11ad BFT. We explore how different BFT approaches scale in large networks, identify the possible problems and evaluate at how the BFT process influences the performance of the network overall. Our analysis shows that Group Beamforming can outperform the legacy approach, resulting in lower overhead and improved network performance. However, we also found that the Access Point (AP) training is quite vulnerable to interference in dense networks, introducing severe limitations to the performance, especially in large rooms where precise BFT is crucial to maintain the communication link. Therefore, we propose several improvements to Group Beamforming that improve performance and provide robust beamforming even in very dense scenarios.**

***Index Terms*—Millimeter Wave, IEEE 802.11ay, Group Beamforming, 60 GHz, WiGig, ns-3, Simulations**

I. Introduction

The exponential growth of mobile devices and mobile traffic, coupled with the ever-increasing need for higher data rates, demands an increase in the capacity and performance of wireless technologies. Existing Wireless Local Area Networks (WLANs) operating in the sub-6 GHz band are becoming overcongested and are unable to cope with emerging applications such as Augmented Reality (AR)/Virtual Reality (VR) and wireless backhauling [1]. Millimeter-Wave (mmWave) technology is foreseen as a key enabler for 5G and beyond wireless communication due to the unprecedented amount of spectrum available. The IEEE 802.11ad amendment [2], introduced in 2012, was the first WLAN standard to support multi-Gbit/s throughput in the unlicensed mmWave band. However, it was never widely deployed as its performance gains were minor compared to the sub-6 GHz IEEE 802.11ac [1]. One of the drawbacks of IEEE 802.11ad is the large overhead and suboptimality of the Beamforming Training (BFT) process [3]. BFT is one of the most crucial processes at mmWave frequency as it determines the beamforming steering at the transmitter and receiver to align the antenna beams and thus overcome the large propagation loss experienced in the mmWave band.

IEEE 802.11ay has only just been approved at the beginning of 2021. It incorporates various technical advancements such as Single-User MIMO (SU-MIMO), downlink Multi-User MIMO (MU-MIMO) and channel bonding/aggregation [4] and is envisioned to enable communication rates close to 100 Gbit/s. Moreover, the BFT process has been redesigned to be more efficient, flexible, and adaptive to cope with different deployments and heterogeneous end-devices. One of the most novel BFT procedures introduced in IEEE 802.11ay is Group Beamforming. This technique is a significant redesign and extension of the IEEE 802.11ad BFT paradigm. Instead of relying on a per Access Point (AP)-Station (STA) pair training, Group Beamforming allows to simultaneously train all STAs associated with a given AP, which drastically reduces the BFT overhead and improves the network performance. Thus, Group Beamforming not only improves the BFT scalability, but also provides an efficient and simple solution to the complex problem of BFT in mobile and dense networks.

Very few papers have investigated IEEE 802.11ay BFT performance. In [3], the new BFT techniques introduced in IEEE 802.11ay are presented, with a special focus on the Beam Refinement Protocol (BRP) Transmit Sector Sweep (BRP TXSS) and asymmetric BFT performance. Asymmetric beamforming is also evaluated in [5] and [6]. Finally, MIMO BFT overhead is studied in [7] and [8]. However, to the best of our knowledge, no investigation of IEEE 802.11ay Group Beamforming has been performed.

The major contributions of this paper are as follows:

- We implement IEEE 802.11ay Group Beamforming in ns-3, based on the existing IEEE 802.11ay model [9]. The implementation extends the ns-3 IEEE 802.11ay model and will be made available to the research community.
- We perform a comprehensive analysis of the system-level performance of Group Beamforming against legacy IEEE 802.11ad BFT using typical IEEE 802.11ay use cases.
- Based on our insights into Group Beamforming performance, we propose and evaluate different improvements. Our analysis shows that our modified version of Group Beamforming outperforms both the legacy 802.11ad BFT and the standard Group Beamforming and provides the best network performance in all scenarios we tested.

II. A Deep Dive into Beamforming Training

The mmWave band exhibits higher propagation loss than the sub-6 GHz band, requiring the use of beamforming to focus the energy in a specific direction and thus increase the transmission and reception gain. This introduces the need for BFT procedures through which devices determine the optimal antenna steering directions in order to establish a directional link. The efficiency and accuracy of the BFT procedure is a key factor in the performance of mmWave networks as incorrect steering can cause link breakage. Moreover, BFT in large and dense networks is an even more complex challenge due to the dynamic wireless channel, the high-interference conditions that can disrupt the BFT and the high number of BFT procedures that need to be performed. IEEE 802.11ad and IEEE 802.11ay define several different BFT mechanisms, as explained below.

A. IEEE 802.11ad Beamforming Training

IEEE 802.11ad uses a two-stage process for the BFT. The first Sector Level Sweep (SLS) phase determines the initial beam patterns to enable communication, while the subsequent BRP phase allows to further refine the beam alignment to obtain maximum gain. The SLS performs the BFT with a sequence of packets called Sector Sweep (SSW) frames, each one transmitted or received with a different beampattern. Each SSW packet is used to train a single directional beampattern, and by measuring the quality of all received SSW frames, a STA can determine the optimal beampattern. The BFT requires two distinct SLSs (an initiator and responder SLS) performed sequentially to determine the optimal transmit (TX) configuration for a pair of nodes that want to communicate. The SLS procedure is simple and reliable, but it suffers from high overhead and limited scalability. Separate training is required in both communication directions and each beampattern is trained in a separate training packet. Therefore, the training duration grows with the number of beampatterns trained and the number of trainings needed grows with the network size.

The optional BRP phase introduces special elements called Training (TRN) subfields. TRN subfields are composed of multiple orthogonal Golay sequences with good correlation properties and each subfield can be used to train a single directional beampattern. Multiple subfields are grouped in a TRN field which is added at the end of BRP training packets. The TRN fields of a packet can be used either for TX or receive (RX) training. For TX training, the transmitter switches beampatterns when transmitting each Transmit Training (TRN-T) subfield, while for RX training all Receive Training (TRN-R) subfields are transmitted with the same beampattern, but are received by switching between different beampatterns. Thus, multiple beampatterns can be trained within a single packet, reducing training overhead and improving performance.

B. IEEE 802.11ay Beamforming Training

The IEEE 802.11ad BFT procedures provide robust BFT, however, they suffer from large overhead and do not scale well to dense and high-mobility networks. To overcome these challenges, IEEE 802.11ay introduces many changes to the BFT process of its predecessor, with four approaches used:

- Reducing SLS overhead and duration: IEEE 802.11ay introduces short SSW frames, which eliminate the Medium Access Control (MAC) overhead of SSW packets. It also allows the use of a partial SLS, which reduces the number of SSW packets needed to perform an SLS.
- Improving TRN field functionality: IEEE 802.11ay significantly increases the number of TRN subfields that can be used in one field, and thus the number of beampatterns that can be trained in a single packet. Additionally, it introduces Receive/Transmit training (TRN-R/T) subfields that enable both TX and RX training in the same packet. Finally, it relaxes the switching time requirements for the antenna array.
- Reducing the BFT duration: IEEE 802.11ay relies much more on TRN subfields rather than SSW packets, as the overhead is significantly lower and scales better. This includes the introduction of BRP TXSS where once a link is established, the BRP protocol is used for BFT. It is also reflected in the design of the new SU-MIMO and MU-MIMO beamforming protocols which mostly train with BRP rather than SSW packets.
- Simultaneous multi-STAs BFT: IEEE 802.11ay defines two new procedures, Group Beamforming and asymmetric beamforming which exploit antenna reciprocity and use TRN-R subfields to allow multiple stations to train *simultaneously*.

C. Group Beamforming

Group Beamforming is a promising novel technique that aims at increasing BFT efficiency. It can provide full BFT for a whole Basic Service Set (BSS) in a single step with very low training overhead. A BSS is composed of an AP and all its associated STAs.

Group Beamforming uses different approaches for training in the AP-STA and the STA-AP direction. Training in the AP-STA direction is similar to the IEEE 802.11ad SLS phase, and reuses periodic beacon transmissions by the AP to train the AP TX sectors. These beacons are transmitted at the start of each Beacon Interval (BI) and serve for initial BSS discovery, distribution of BSS parameters and BI scheduling. To ensure that STAs all around the AP can detect the network, the beacons are transmitted sequentially in a directional manner and therefore can also function as TX SSW frames for the AP.

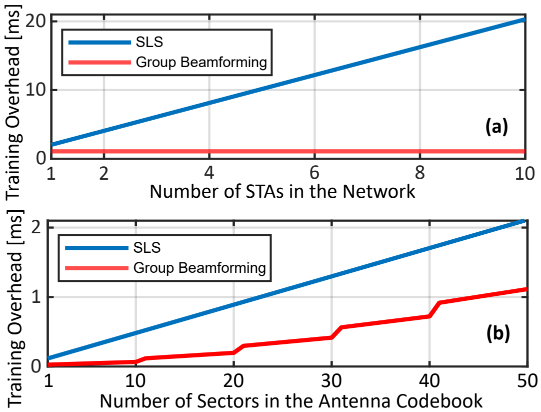

Fig. 1: Overhead of Group Beamforming and SLS with different (a) number of STAs, (b) antenna codebook size.

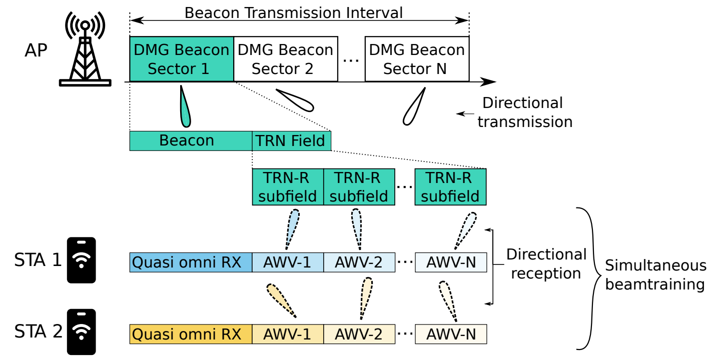

Fig. 2: An example of Group Beamforming with 2 STAs.

In this way training in the AP-STA direction is done with no extra overhead, as the beacons are already a necessary part of network operation.

The novelty of Group Beamforming is in the STA-AP training which proposes adding TRN-R subfields to each beacon. All STAs in the BSS that receive the beacons can then use the TRN-R subfields to perform simultaneous RX training. Assuming antenna reciprocity at the STAs antennas, i.e., the TX and RX beampatterns are equivalent, the best STA RX beampattern will also be the best TX beampattern, which thus provides the full antenna configuration for communication with the AP. Note that this adds redundancy to the training, as STAs repeat the training with each received beacon. Since all STAs train with the same TRN-R subfields, the training overhead of Group Beamforming does not depend on the number of STAs in the network. Additionally, the TRN fields have a very low overhead - each subfield has a duration of only 437 ns.

Therefore Group Beamforming has a significantly lower overhead than other BFT approaches defined in IEEE 802.11ad and IEEE 802.11ay. To illustrate this, in Figure 1 we show a comparison of the training overhead for Group Beamforming and legacy SLS with respect to the number of STAs in the network and the antenna codebook size. In Figure 1 (a) we assume that the initiator and responder antenna codebooks both have 48 sectors, and in Figure 1 (b) we calculate the overhead with one STA in the network. We can clearly see that while the SLS overhead grows with the addition of STAs in the network, the Group Beamforming overhead remains constant regardless of the size of the BSS, as it involves only the addition of TRN-R subfields to the beacons. The jumps in the Group Beamforming overhead in Figure 1 (b) come from the TRN-R field structure. Namely, the TRN-R field is composed of multiple Units, each with 10 subfields for training. Therefore, every 10 additional sectors in the STA codebook require one extra TRN Unit for training. The additional growth between the jumps is because the TRN-R subfields are added to each beacon sent by the AP. Therefore, the Group Beamforming overhead also grows with the size of the AP codebook, although the growth is still much lower than the corresponding SLS overhead.

Figure 2 shows an example of Group Beamforming training. The AP iterates over its TX beampatterns and appends TRN-R subfields to every beacon. The STAs receive the beacon with a quasi-omni RX beampattern (i.e, the widest beampattern), and then iterate over their RX beampatterns when receiving the TRN-R subfields, measuring the reception quality. Once every beacon and TRN-R subfield are received, the STAs are able to determine: 1) the best TX beampattern for the AP, 2) their best RX beampattern, and 3) assuming antenna reciprocity, their best TX beampattern and the best RX beampattern for the AP.

III. Group Beamforming performance enhancements

In addition to the standard compliant version of Group Beamforming, we also design and implement different performance enhancements. As explained in Section II-C, Group Beamforming has two different approaches for the AP and STA training. Our analysis in Section V shows that the STA training using TRN-R subfields is much more resilient to interference than the AP training using beacons, due to the ability to redundantly repeat the training with each received beacon. Therefore, we designed a version of Group Beamforming where the AP training also uses TRN-T subfields appended to beacons. In this case, the AP and STA training is done in two consecutive BIs. In the first one, the AP appends TRN-R subfields to each transmitted beacon and all STAs perform RX training and update their optimal antenna configurations. In the second BI, the AP appends TRN-T subfields to the beacons and switches the TX beampattern when transmitting each subfield, while the STAs receive the subfields in quasi-omni mode and record the received Signal-to-interference-and-noise ratio (SINR).

We also explore the effect of directional transmission and reception on the RX training of the STAs when using TRN-R subfields. Normally, the BFT in IEEE 802.11ad and IEEE 802.11ay networks uses directional transmission and quasi-omni reception. To study how this affects the accuracy of the BFT decisions, we modified the transmissions of beacons with TRN-R subfields so that the AP switches from directional to quasi-omni transmission once the TRN field begins.

IV. Implementation

We implemented IEEE 802.11ay in the system-level simulator ns-3, allowing us to evaluate high-density IEEE 802.11ay networks, taking into account different MAC layer aspects such as signalling overhead, channel access, packet collision, etc. We have expanded our previous implementation described in [9] to fully support the IEEE 802.11ay standard, including among other things Group Beamforming. Our implementation includes support for the transmission of the new Physical Layer Convergence Protocol (PLCP) frame format, the new Modulation and Coding Schemes (MCSs), all necessary new MAC headers and WiFi information elements, as well as the state machines for transmission and reception of the new TRN fields. Our full IEEE 802.11ay implementation, including Group Beamforming will be made publicly available on GitHub.

For our Group Beamforming implementation, we first enabled APs to add TRN-R subfields to Directional Multi-Gigabit (DMG) Beacons. The user can configure the periodicity of Group Beamforming by specifying the periodic interval, in number of BI, at which TRN-R subfields are appended to the beacons. Then we enabled STAs to sweep though their RX beampatterns while receiving the TRN-R subfields, measuring the received SINR. The STAs use these measurements to choose the best antenna configurations for communication with the AP. If they detect a change in the optimal TX sector of the AP, they send an unsolicited `Information Response` frame to the AP informing it of the change.

V. Simulations

We conducted a set of simulations to compare the performance of Group Beamforming and legacy SLS. For our simulations we use the Quasi-Deterministic (Q-D) channel realization software [10], which uses a combination of ray tracing and statistical modelling to generate the channel model with all Line-Of-Sight (LOS) components and first order reflections. We perform simulations in an indoor scenario with rectangular rooms of two sizes: a smaller room with dimensions $7.4\ m \times 13.5\ m \times 3\ m$ and a larger room with dimensions of $29.6\ m \times 54\ m \times 3\ m$. We simulate varying network densities with 1 to 8 APs and 1 to 64 STAs. The STAs are randomly placed within the room in a fixed position for the duration of the simulation. The association to the AP is according to the closest AP. However, we use load balancing to ensure an equal number of STAs per AP, meaning that in some cases STAs might not be associated with the closest AP. Unless otherwise specified, the number of STAs associated with each AP is 8. All APs are mounted on the ceiling at a height of 3 m while STAs are placed at a height of 1.2 m. Figure 3 shows an example network topology with 4 APs and 32 STAs. Due to the random node distribution, in some cases STAs can be close to each other and suffer from high interference. Additionally, STAs located far from the AP may experience low Signal-to-Noise Ratio (SNR). Finally, as mentioned, due to the greedy load balancing, some STAs may not be associated with the closest AP. These devices can suffer both from the large distance to the AP and the interference from the other close-by BSS.

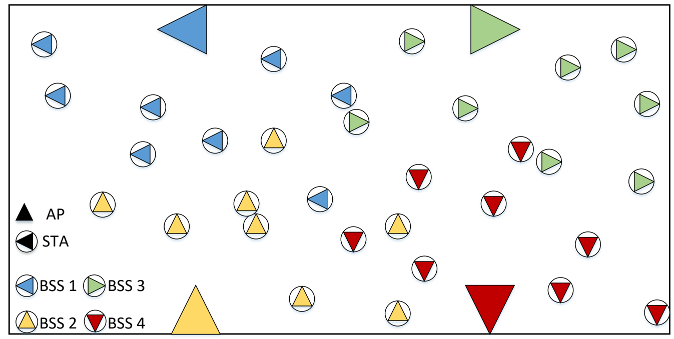

Fig. 3: Example network topology for a 4 BSS network.

All devices use 2x8 element Phased Antenna Arrays (PAAs) in the form of a Uniform Rectangular Array (URA) generated with the IEEE 802.11ad Codebook Generator [10]. The devices perform Group Beamforming or SLS each 10 BIs. In the SLS simulations, the training is initiated by the STA at the start of the Data Transmission Interval (DTI). The results are averaged over 40 simulation instances with a duration of $50\ s$, using 20 scenarios with different STA locations, each run two times with different random seeds.

For the data communication, we use a rate adaptation algorithm that chooses the MCS based on the last received SINR. We evaluate network performance with downlink (DL) and uplink (UL) User Datagram Protocol (UDP) traffic. We vary the station data rate between 100 Mbit/s and 1 Gbit/s, allowing us to study the performance with different levels of channel saturation. Unless otherwise specified, the simulation data rate is set to 100 Mbit/s, as we found that this allowed us to study and compare networks of different densities without completely oversaturating the channel. Finally, although Group Beamforming enables communication with directional TX and RX beam patterns, we use quasi-omni reception when receiving data packets for a fair comparison with SLS.

We evaluate the BFT performance by assessing the accuracy and quality of the chosen beams. We define accuracy as the percentage of time that the optimal sector is chosen. We consider as optimal the sector that has the highest SNR in the absence of any interference, measured with directional transmission and quasi-omni reception. In addition, we calculate the SNR loss due to BFT errors as the difference between the maximum possible SNR of the optimal sector and the SNR of the actual sector chosen by the BFT. For this purpose we look at pure SNR without interference. In this way we compare the BFT quality based on the distribution of the SNR loss values and the distribution of the average SNR loss of STAs. We analyze the network performance in terms of the per STA and aggregate network throughput.

A. *Performance in a Single BSS Network*

1) BFT performance: We begin the analysis by looking at simple scenarios with 1 AP in the network. The analysis of networks with a single AP is separated as this is the most favorable case for Group Beamforming compared to SLS.

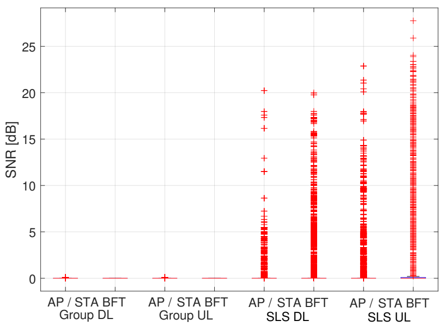

Fig. 4: SNR loss from BFT choices, 1 AP, 8 STAs, large room.

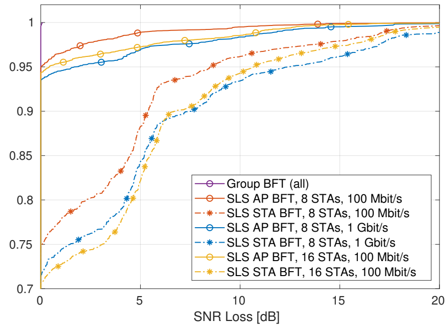

Fig. 5: CDF of SNR loss from BFT errors, 1 AP, UL traffic, large room.

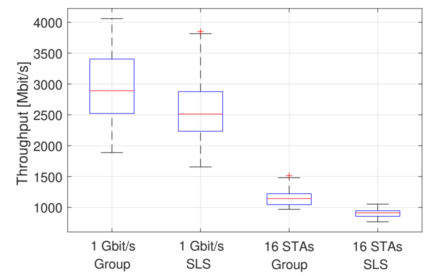

Fig. 6: Aggregate network throughput depending on the BFT method, 1 AP, UL traffic, large room.

This can be clearly seen in Figure 4 which shows the SNR loss due to BFT decisions for the large room with both UL and DL traffic of 100 Mbit/s and 8 STAs in the network. We show separately the SNR loss of the APs and the STAs since Group Beamforming uses different training approaches for training in the different directions. It is evident Group Beamforming is extremely accurate in choosing the optimal sector, resulting in SNR losses of zero. This is because Group Beamforming moves the BFT from the DTI, where typically STAs compete for access to the channel in a contention-based manner, to the Beacon Transmission Interval (BTI), where only the AP is allowed to transmit. Thus, the BFT is free from any interference caused by STAs within the BSS, regardless of the size of the BSS and the amount of traffic. This eliminates interference in the single BSS case and can also improve BFT accuracy in multi-BSSs networks as STAs within our BSS are typically the closest and strongest interferers.

The few outliers that show an SNR loss slightly higher than zero come from an interesting effect where in very rare cases Group Beamforming never chooses the optimal sector. This is because Group Beamforming is performed with both the transmitter and receiver in directional mode, as described in Section III. In very rare cases, this can result in different BFT decisions than the ones made with the transmitter in directional and the receiver in quasi-omni mode, as the RX beampattern affects the amplification or attenuation of different multi-path components. In low-interference scenarios the resulting SNR loss is low (less than 1dB in Figure 4) and there is no performance degradation. However, the effect can be more detrimental in dense, high-interference scenarios.

Analysing legacy SLS performance shows that interference can lead to BFT errors even in small networks. The exact SLS performance depends on the saturation of the wireless channel and we found that it degrades as the interference in the network increases. The results presented in Figure 4 show that the average SLS performance matches that of Group Beamforming as most BFT decisions are correct. However, we can also observe outliers with SNR loss of up 28 dB caused by SLS errors, leading to link instability and even loss of service.

Additionally, in Figure 4 we can also see that UL traffic causes more SLS errors than DL traffic and that the STA training is more vulnerable than the AP training. Both effects are due to the higher level of contention that STAs experience. When analyzing the effect of the room size, we found that, as expected, Group Beamforming remained perfect in both rooms. For SLS, we noticed a performance degradation in the large room. In the small room, the measured SLS accuracy was above 95%, while in the large room it went down to 75%. Lastly, we found that the highest SNR losses came from collisions with other SLS packets, rather than data packets. In these cases, the interference affects most or all SLS packets and constantly varies as the interfering STA sweeps through TX sectors, making it more difficult to find the optimal sector.

2) Network performance: We analyze network performance in terms of achieved throughput to answer the following questions: (i) how do errors in BFT impact the throughput and (ii) how does the choice between Group Beamforming and SLS affect the overall performance of mmWave networks?

When we analysed the network performance we found that the wireless channel is under-saturated with a data rate of 100 Mbit/s so most SLS errors are not reflected in the data throughput. In the small room, both Group Beamforming and SLS lead to optimal network performance of 797 Mbit/s and average STA throughput of 99.65 Mbit/s, with both UL and DL traffic. However, in the large room UL results we found a slight performance degradation when using SLS. In some SLS scenarios the aggregate network throughput drops to 765 Mbit/s, while the average per STA throughput can go down to 88 Mpbs. This degradation is directly caused by errors from interference in the SLS BFT, since the throughput performance in the exact same scenarios with Group Beamforming remains optimal.

3) Performance under higher channel saturation: Next, we explore how increasing the channel saturation affects the BFT and the overall network performance. For this purpose we analyse two additional single-BSS scenarios, one with a data rate of 1 Gbit/s (instead of 100 Mbit/s) and another with 16 STAs (instead of 8). Figure 5 compares the Cumulative Distribution Function (CDF) of the BFT SNR loss. We notice that the SLS accuracy drops in both new scenarios. It is also interesting to see that although the data rate increase adds much more traffic than the STA increase, the SLS accuracy is pretty similar in both cases. This is because increasing the number of STAs in the network increases both the training overhead and the collisions from carrier-sensing failure. We can also see

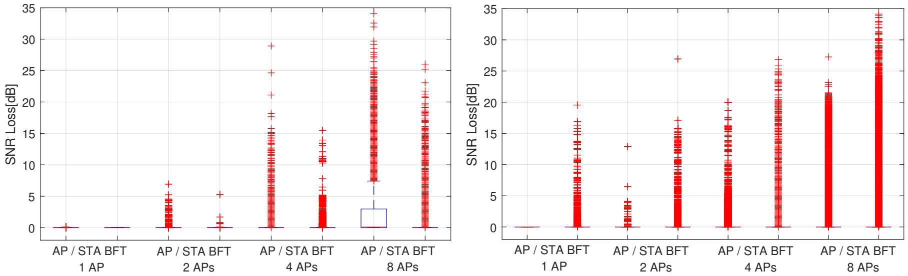

(a) Group Beamforming

(b) SLS

Fig. 7: SNR loss depending on network density, DL traffic, small room.

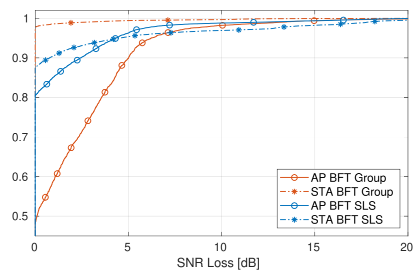

Fig. 8: SNR loss CDF in a network with 8 APs, DL traffic, small room.

that Group Beamforming still provides perfect BFT in both cases, while also keeping the overhead constant. The benefits of this can be seen in Figure 6 which shows a comparison of the aggregate network throughput with UL traffic in the large room. In both cases the wireless channel is over-saturated and the desired data rate can not be achieved. However, in both cases there is a gain when using Group Beamforming. This gain is on average 370 Mbit/s in the 1 Gbit/s case and 230 Mbit/s in the 16 STAs scenario and comes both from the perfect BFT and the lower overhead.

B. Performance in Networks with Multiple BSSs

1) BFT performance: Networks with multiple BSSs introduce challenges to the BFT, particularly for Group Beamforming. Unlike the single BSS case, Group Beamforming is not protected from interference by STAs in other BSSs. Therefore we vary the number of BSS and the network density and study the effects on BFT performance. Figures 7 (a) and (b) show the SNR loss distribution for Group Beamforming and SLS respectively. The presented results are for simulations of the small room with 100 Mbit/s DL traffic and 8 STAs per AP. The trends observed also occurred in all other scenarios tested.

The first conclusion that can be drawn is that in high interference scenarios, the BFT procedure can be compromised, leading to the high SNR losses seen in Figure 7. Depending on the resulting SNR, this instability can lead to drops in the achievable data rate and even link outage. Secondly, in Figure 7 we observe an asymmetry in the AP and STA BFT performance, both for Group Beamforming and SLS. In order to better illustrate this asymmetry, we show in Figure 8 the CDF of the SNR loss for the 8 BSSs network. In the case of Group Beamforming, it is evident that the APs training is much less reliable and has the worst performance overall. For SLS, there is a less pronounced asymmetry where the STA training results in higher SNR losses. In many scenarios tested, we also found that the SLS STA training had a lower accuracy.

This SLS asymmetry can be explained by the STA-initiated SLS procedure and higher channel contention that STAs experience. Additionally, most APs are located at the room's edges which exposes them to less interference. We found that this asymmetry in BFT performance is very slight and does not seem to be reflected in the network throughput performance.

In the case of Group Beamforming, however, the training difference is very large, requiring further investigation. In Figure 7 it is evident that the Group Beamforming STA training has the least amount and lowest SNR losses, compared both with the Group Beamforming AP training and the SLS training. In fact, if we look at the accuracy of the STA training in Figure 8, we can see that it is above 97%, which shows remarkable robustness and stability. This is particularly impressive if we consider that Figure 8 shows a very dense, high-interference scenario, with 8 APs and 64 STAs distributed in a relatively small room. Our simulations showed the same behaviour across all other simulation scenarios. Although we found that UL traffic was slightly more disruptive, the Group Beamforming STA accuracy was always above 93%, in both the large and small room scenarios. For comparison, the worst SLS accuracy we measured was 70% and 68% for the STA and AP training, respectively. The Group Beamforming AP training, on the other hand, suffered much more in high-interference scenarios. In Figure 8 we see that the accuracy is 48%, far below the SLS accuracy and that this is also reflected in Figure 7, where the Group Beamforming AP training is the only one with an upper quartile above 0 dB. In addition to this, when we looked at the average per STA SNR loss, the Group Beamforming AP training had by far the worst performance. The mean SNR loss over the course of the simulation reached up to 13 dB, while it did not go over 5 dB for the SLS training or the Group Beamforming STA training.

This stark difference in the Group Beamforming performance comes from the fundamentally different approaches to the AP and STA BFT. The TRN-R subfields used for STA training are appended to each transmitted beacon, so STAs which receive multiple beacons perform the RX training multiple times, leading to increased reliability. Additionally, training with TRN subfields allows for all RX beampatterns to be tested in a single packet. This considerably reduces the probability that there will be significant interference variations during the training and enables obtaining coherent measurements, which can also increase the performance. This results in highly resilient BFT, which is not only robust to different types of interference, but also consistently outperforms legacy SLS. In contrast to this, the AP BFT, done with directionally transmitted beacons, seems to be very sensitive to interference and more vulnerable than SLS. This is a consequence of the simultaneous training of all STAs in the BSS introduced by Group Beamforming. This greatly reduces the training overhead, but also means that

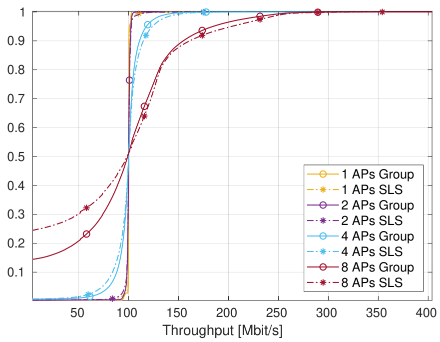

Fig. 9: Throughput CDF depending on BFT method, DL traffic, small room.

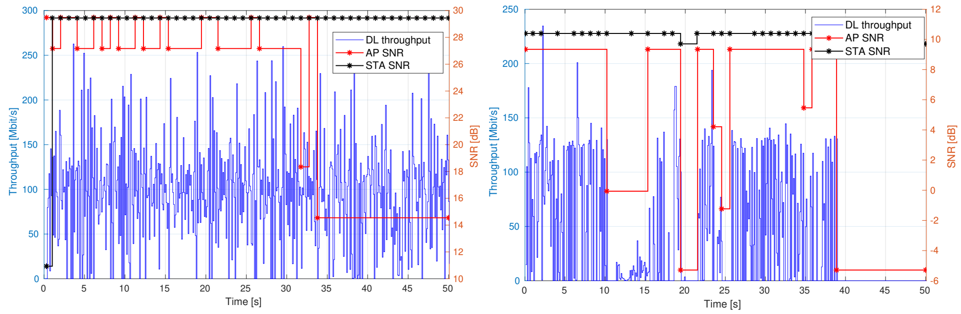

(a) Small room scenario

(b) Large room scenario

Fig. 10: Throughput and SNR over the course of the simulation, 8 APs, DL traffic, Group Beamforming.

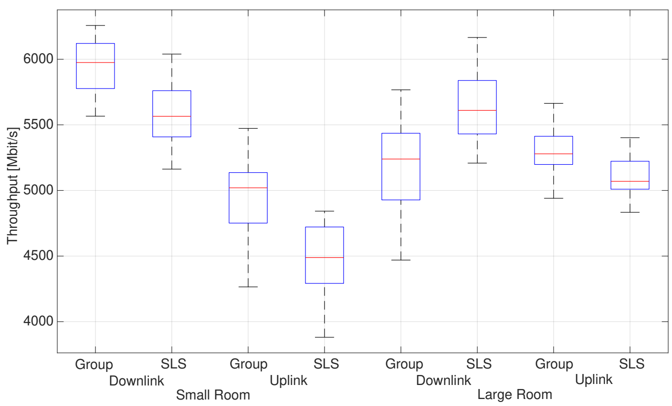

Fig. 11: Aggregate network throughput depending on the BFT method, traffic type and room size, 8 APs, 64 STAs.

BFT decisions for the whole BSS are done based on the same beacon frames. Interference during the beacon transmissions will therefore affect the AP's training with all STAs in the BSS. Errors in the SLS procedure, on the other hand, are detrimental solely for the AP-STA pair performing the procedure.

Next, we analysed the effect of the room size on BFT performance. We noticed a performance degradation in the larger room, which was stronger in lower density networks and decreased with the number of devices in the network. As the average distance between devices is higher in the large room, there are more collisions due to carrier-sensing failure leading to BFT errors. The effect of these errors is more evident in lower density networks where the wireless channel is not very saturated. Finally, when comparing the BFT performance with DL and UL traffic we found that UL traffic tends to be more harmful. This is due to two reasons. Firstly, in DL transmissions, there are fewer devices transmitting and therefore the collision probability is smaller. Secondly, most of the APs are located at the room's edges and cause lower interference than the STAs which are distributed inside the room.

2) *Network performance:* Figure 9 shows the CDF of the DL throughput measured in 0.1 s intervals, varying the network density and BFT method. The throughput performance here corresponds to the BFT performance shown in Figure 7. We can clearly see how increasing the wireless channel saturation affects STA performance. In the 1 and 2 AP networks, the performance is extremely close to optimal and there is no discernible difference between Group Beamforming and SLS. However, in the 4 AP network we begin to see a slight performance gain in using Group Beamforming, and this gain grows in the very dense 8 AP network. In the 4 and 8 AP networks, SLS causes more throughput values in the 0 to 100 Mbit/s interval, reflecting the time spent on SLS instead of transmitting data. In very dense networks, the extra time for data transmissions given by Group Beamforming can lead to significant increases in network performance. In Figure 11, we show the aggregate network throughput for the densest 8 AP, 64 STA scenario, comparing the performance with Group Beamforming and SLS. We can see that in the small room, DL traffic scenarios, Group Beamforming outperforms SLS by 400 Mbit/s on average.

This gain does seem somewhat surprising if we consider that Group Beamforming had significantly worse BFT performance. However, due to the closeness of devices in the small room, BFT errors rarely cause link outage and mostly result in a reduced MCS. Therefore, we have found that for the small room, the overhead reduction of Group Beamforming benefits network performance, despite the reduction of BFT accuracy. The overhead gains increase with the network density as SLS overhead grows, counterbalancing the BFT accuracy reduction.

As we can see from Figure 11, this effect is more pronounced with UL traffic, leading to an average increase of 500 Mbit/s in aggregate network throughput. Contributing to the better UL performance is the fact that errors in Group Beamforming are overwhelmingly done for the AP sector, while the STA sector is chosen extremely accurately. When considering UL transmissions done in a contention-based manner, the STAs are in directional TX mode, while the AP is in quasi-Omni RX mode. Therefore, the wrongly chosen AP sectors have little influence on network performance, as they are only used to send back Block Acknowledgments, sent with a low MCS that is decodable even at low SNR.

We also found that in the small room the throughput performance was predominantly affected by the level of wireless medium saturation and failures in carrier-sensing caused by directional transmissions. For this reason, in these

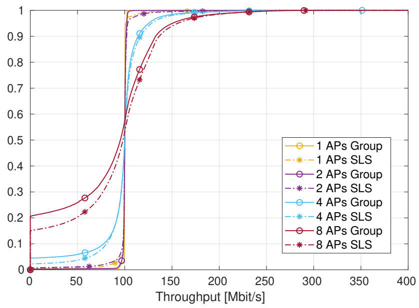

Fig. 12: Throughput CDF depending on network density, DL traffic, large room.

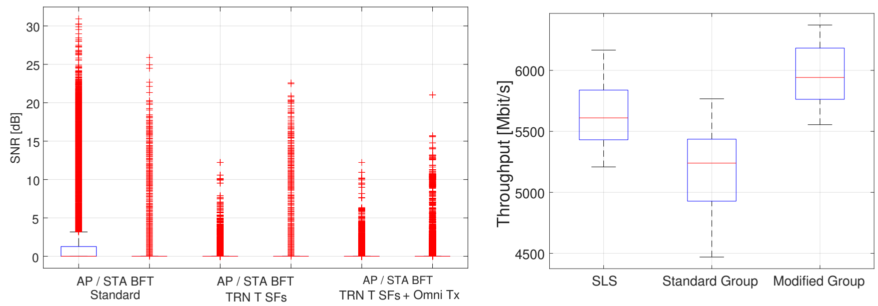

(a) SNR loss due to errors in BFT

(b) Aggregate network throughput

Fig. 13: Performance with modified versions of Group Beamforming, 8 APs, DL traffic, large room.

types of scenarios, Group Beamforming is a very good option due to the reduction in overhead that it provides.

In the large room, however, devices are much further apart. This results in lower optimal SNR and fewer sectors that still provide satisfactory SNR for communication. SNR losses caused by errors in BFT often mean link breakage until the device can recover from the error. In addition to this, there is more potential for spatial re-use, so the throughput is not as affected by collisions. Therefore, we found that the throughput performance in these scenarios had a significant connection to the BFT performance. To highlight the difference between operation in the small and large room, we show in Figure 10 a throughput comparison of two STAs that both experience severe SNR losses due to BFT errors. The network density (8 APs, 64 STAs), data rate (100 Mbit/s DL traffic) and BFT method (Group Beamforming) are the same for both scenarios and the only difference is the size of the simulated room. In both scenarios, BFT errors in the DL direction lead to drops of 10 dB to 15 dB in the achievable SNR for a significant portion of time. In the small room, we found that the errors were caused by the fact that the STA was located in the center of the room and experienced high interference from surrounding devices. In addition to this, due to collisions, some of the BFT feedback packets were not correctly received, and as a consequence the AP was not able to update its antenna configuration in time. In the large room, due to load balancing, the STA was not associated to the closest AP and the large distance caused problems in the BFT. When observing the effect of the drop in SNR on the throughput, we can see that in the small room, the data transmissions are not affected, as the SNR is still well above the decodability limit. In the large room, however, we see that multiple BFT errors cause communication link breakage, as the SNR is too low for successful data transmissions.

We can see the effects of this in Figure 12, which shows the DL performance in the large room. In contrast to Figure 9, where we observed that Group Beamforming provided better performance than SLS, we can see that in the 4 and 8 AP scenarios the network performance with Group Beamforming suffers due to the poor BFT performance. These were the only scenarios where we found that SLS outperformed Group Beamforming. In Figure 11 we can see that there is a difference of approximately 400 Mbit/s in the 8 AP network. The UL results show the opposite trend, as the high accuracy of the STA training combined with the higher level of network contention result in better performance than SLS. The performance gains, however, are lower than in the small room and only go up to 200 Mbit/s. This is a combination of the more severe effects of the Group Beamforming BFT errors and the increased spatial re-use which reduces the benefit of the overhead reduction.

Similar as for the BFT performance, we found that in general UL transmissions had worse throughput performance. The exception to this was the performance in the Group Beamforming, large room scenarios, with 4 and 8 APs, as can be seen in Figure 11. Due to the reasons discussed above, in these cases the DL performance suffered significantly and was equal to the UL performance.

Finally, a direct comparison of the results in the small and large room for the same scenarios showed that the small room had better throughput performance due to the link breakages caused by BFT errors and the lower achieved SNRs which resulted in lower MCSs used. The exception to this were the UL scenarios with 8 APs in the network, as seen in Figure 11. In these cases the wireless medium was over-saturated and therefore the higher spatial re-use in the large room caused gains in performance.

C. *Performance with modified versions of Group Beamforming*

Our performance analysis of Group Beamforming detected several causes of BFT errors, as described in the above sections. Therefore, we propose modifications that deal with these issues. The main performance limitation was the vulnerability of the AP training in high-interference scenarios. The contrast with the extremely robust STA training motivated us to implement a similar approach for the AP training. As explained in Section III, we modified the AP training to use TRN-T subfields rather than rely on the beacon packets. This doubled the Group Beamforming overhead but significantly improved the training accuracy. In addition to this, we also attempted to further improve the STA training, by eliminating having both the transmitter and receiver in directional mode. Since this sometimes led to errors, we modified the APs to transmit the TRN-R subfields with a quasi-omnidirectional beampattern.

Figure 13 shows a comparison of the performance of Group Beamforming with the different proposed versions both in terms of the SNR loss and aggregate network throughput. The results show the performance in a network with 8 APs and DL traffic in the large room. As described in the above section, this was one of the scenarios where Group Beamforming had poor performance and was outperformed by SLS.

We can clearly see that our modifications significantly improve the BFT performance, particularly in terms of the AP training accuracy. The accuracy increased from 60 % to 95 % and the highest SNR loss measured dropped from 30 dB to 12 dB. In addition, we noticed that the overhead increase did not cause any throughput degradation, as it was offset by the better BFT. In fact, the BFT improvements resulted in significant throughput gains, as seen in Figure 13 which shows an average increase of 700 Mbit/s in the aggregate network throughput. The STA BFT accuracy increase was much smaller, as the STA training was already quite close to optimal, resulting in only minor throughput gains. However, this modification also decreased the SNR losses, leading to fewer link breakages and an improvement in the quality of service.

We found that with our performance enhancements Group Beamforming was able to outperform SLS in all tested scenarios, both in terms of BFT accuracy and achieved throughput.

VI. Discussion

There are several lessons we learned from our analysis, which is the first known to evaluate the performance of Group Beamforming, both regarding BFT procedures and the performance of dense mmWave WLAN networks.

First, we found that although all mmWave networks suffer from high path loss and require directional communication, the room size determines how much the BFT accuracy affects network performance. In larger rooms, where the distance to the AP is big and there are no strong reflections, only a few sectors offer a sufficiently high SNR for communication and extremely accurate BFT is crucial for good performance. Errors in the BFT lead to link outage and it can take a long time until the STA recovers. In contrast, in smaller rooms, issues with carrier-sensing failure cause much more performance problems than BFT failure. As long as the SNR remained high enough to maintain a communication link, we found that network performance was not affected by BFT errors as even the lower MCSs can offer high data rates. However, depending on the nature of the application these errors may affect the quality of service and the user experience.

Another observation was that carrier-sensing failure can be a significant problem in dense mmWave networks that limits the network performance. In many cases, although the wireless channel can handle more traffic, it is limited by collisions caused by the failures, which are much more common in mmWave networks due to the directional communication. This is especially a problem in UL communication, where there is a lot of channel contention. More research is necessary to find solutions for this problem and improve the channel access procedure for mmWave WLAN networks.

When considering BFT performance, we found that even in very dense networks, the BFT was mostly successful with both approaches tested. BFT errors were often transient and devices were able to recover in the next training. However, these errors could lead to significant drops in achievable SNR.

The legacy SLS mechanism was shown to be quite robust to different levels of interference. However, the high training overhead affected the overall network performance and therefore this approach is not suitable for high density networks.

In terms of Group Beamforming, our analysis showed that it is a very promising approach. First, shifting the BFT to the DTI ensures complete protection from interference from STAs in the same BSS, providing perfect BFT with lower overhead. Therefore, single-BSS networks benefit significantly from its use. Moreover, the use of TRN-R subfields in Group Beamforming was proven to be extremely efficient and reliable. The drastically lower overhead offers the possibility to add redundancy and repeat the training multiple times at a fraction of the overhead needed with full packets. This resulted in extremely high BFT accuracy which outperformed the other approaches tested. We did find, however, that the simultaneous training of all STAs in the BSS has some drawbacks. As all STAs make their BFT decisions based on the same transmissions, training corrupted by interference can cause errors for all STAs in the network, leading to degradation in the performance of the whole BSS. Switching the AP training to the TRN approach, however, fixed this problem and significantly improved the performance. It is important to highlight that the standard-compliant version of Group Beamforming might have problems in high-density scenarios, particularly for DL traffic. Moreover, as typically the gains of the antennas in APs are much larger than the gains of the antennas in user devices, the AP BFT errors can cause more harm to the overall network performance. Lastly, the training overhead reduction can improve network performance as compared to legacy SLS, even when the BFT suffers more.

VII. Related Work

Most existing work on IEEE 802.11ay is theoretical and either provides a protocol overview or proposes solutions to some open design issues. In [1] the authors give an excellent overview of IEEE 802.11ay, while in [11] the authors focus on detailed technical explanations of aspects like spatial sharing, channel bonding and interference mitigation.

Some work focuses specifically on the BFT aspects of IEEE 802.11ay, such as [12], where the authors study scheduling, beamforming, and link maintenance protocols. Lastly, the authors of [3] explain the legacy IEEE 802.11ad BFT and show theoretical results on asymmetric beamforming. The authors in [13] propose a dynamic scheduling mechanism for asymmetric BFT based on SINR with random backoff. The authors in [14] focus on dense scenarios and propose an enhanced BFT which is backward compatible with IEEE 802.11ad. Their implementation divides the Association Beamforming Training (A-BFT) slots to further improve the association. Another backward compatible BFT solution is presented in [15], where the authors extend the IEEE 802.11ad channel model to support Hybrid MIMO configurations.

Practical analysis of IEEE 802.11ay is out of the scope of these papers, which limits the potential to find new open research issues. To the best of our knowledge [16] is the only paper with system-level validation regarding IEEE 802.11ay. This paper focuses on channel bonding in IEEE 802.11ay. The authors briefly mention the implementation of an ns-3 module to validate their theoretical work, but the implementation details are unclear and the source code has not been released. The authors leave out some important considerations about the implementation, such as the BFT mechanism used.

Multiple measurement campaigns have evaluated the performance of IEEE 802.11ad BFT in Commercial Off-the-Shelf (COTS) devices. The authors in [17] were the first to provide an in-depth analysis on the performance and operations of WiGig compliant COTS devices, but lack an analysis of the beamforming mechanism. IEEE 802.11ad BFT can be improved using compressed sensing [18]. The implementation on COTS routers provides a significant reduction in training overhead compared to plain IEEE 802.11ad BFT. The authors of [19] focused on the interplay between BFT and the rate adaptation algorithm as well as the impact of BFT errors due to mobility or blockage on the mmWave link stability.

To the best of our knowledge, our work is the first to focus on IEEE 802.11ay Group Beamforming and evaluate its performance in various scenarios and network configurations.

VIII. Conclusions

In this paper we presented a study on the performance of Group Beamforming, a new BFT technique introduced in the novel IEEE 802.11ay amendment. For this purpose we extended the existing ns-3 mmWave module to support Group Beamforming and evaluated its performance in a variety of simulation scenarios. Our evaluation showed that the new concept of Group Beamforming which includes simultaneous training of all STAs in the BSS can boost network performance by lowering the training overhead of legacy BFT procedures like SLS. We found, however, that the AP training faced challenges in high-density networks. Therefore, we implemented and tested two modifications of Group Beamforming that significantly improved its performance. With the suggested modifications, the new Group Beamforming approach was able to outperform the legacy SLS procedure in all tested scenarios, both in terms of BFT and throughput performance. Therefore, we consider it an extremely promising solution for BFT in dense mmWave networks.

IX. Acknowledgments

This work has been funded by the European Union's Horizon 2020 research and innovation program 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] Y. Ghasempour, C. R. C. M. da Silva, C. Cordeiro, and E. W. Knightly, "IEEE 802.11ay: Next-Generation 60 GHz Communication for 100 Gb/s Wi-Fi," *IEEE Communications Magazine*, vol. 55, no. 12, pp. 186–192, Dec. 2017.

[2] IEEE, "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.

[3] C. R. C. M. Da Silva, J. Kosloff, C. Chen, A. Lomayev, and C. Cordeiro, "Beamforming Training for IEEE 802.11 ay Millimeter Wave Systems," in *2018 Information Theory and Applications Workshop (ITA)*. San Diego, CA: IEEE, Feb. 2018, pp. 1–9.

[4] IEEE, "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.

[5] A. Maltsev, A. Pudeyev, I. Bolotin, and O. Bolkhovskaya, "Asymmetric Links Beamforming for mmWave Overlay in LTE-Based Heterogeneous Network," in *2017 IEEE 86th Vehicular Technology Conference (VTC-Fall)*, 2017, pp. 1–5.

[6] Y. Kim, S. Lee, and T. Ropitault, "STS Adaptation for Beamforming Training of Asymmetric Links in IEEE 802.11ay-based Dense Networks," in *IEEE Vehicular Technology Conference (VTC2020-Spring)*, 2020, pp. 1–6.

[7] M. Kim, T. Ropitault, S. Lee, and N. Golmie, "Efficient MU-MIMO Beamforming Protocol for IEEE 802.11ay WLANs," *IEEE Communications Letters*, vol. 23, no. 1, pp. 144–147, 2019.

[8] M. S. Kim, T. Ropitault, S. Lee, N. Golmie, H. Assasa, and J. Widmer, "A Link Quality Estimation-Based Beamforming Training Protocol for IEEE 802.11ay MU-MIMO Communications," *IEEE Transactions on Communications*, vol. 69, no. 1, pp. 634–648, 2021.

[9] H. Assasa, N. Grosheva, T. Ropitault, S. Blandino, N. Golmie, and J. Widmer, "Implementation and Evaluation of a WLAN IEEE 802.11ay Model in Network Simulator ns-3," in *in the Workshop on ns-3*, 2021.

[10] H. Assasa, J. Widmer, T. Ropitault, and N. Golmie, "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, 2019.

[11] P. Zhou, K. Cheng, X. Han, X. Fang, Y. Fang, R. He, Y. Long, and Y. Liu, "IEEE 802.11ay based mmWave WLANs: Design Challenges and Solutions," *IEEE Communications Surveys & Tutorials*, vol. 20, no. 3, pp. 1654–1681, 2018, arXiv: 1803.07808.

[12] C. Chen, O. Kedem, C. R. C. M. da Silva, and C. Cordeiro, "Millemeter-Wave Fixed Wireless Access Using IEEE 802.11ay," *arXiv:1907.00082 [cs]*, Jun. 2019, arXiv: 1907.00082.

[13] Y. Kim, S. Lee, and T. Ropitault, "Adaptive Scheduling for Asymmetric Beamforming Training in IEEE 802.11ay-based Environments," *IEEE Wireless Communications and Networking Conference (WCNC)*, 2019.

[14] P. Zhou, X. Fang, Y. Fang, Y. Long, R. He, and X. Han, "Enhanced Random Access and Beam Training for Millimeter Wave Wireless Local Networks With High User Density," *IEEE Transactions on Wireless Communications*, vol. 16, no. 12, p. 7760–7773, 2017.

[15] F. Fellhauer, N. Loghin, D. Ciochina, T. Handte, and S. Brink, "Low Complexity Beamforming Training Method for mmWave Communications," *2017 IEEE 18th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, 2017.

[16] L. Lei, B. Li, M. Yang, and Z. Yan, "System Analysis and Performance Evaluation for the Next Generation mmWave WLAN: IEEE 802.11ay," in *IEEE International Conference on Signal Processing, Communications and Computing (ICSPCC)*. Qingdao: IEEE, Sep. 2018, pp. 1–6.

[17] G. Bielsa, A. Loch, I. Tejado, T. Nitsche, and J. Widmer, "60 GHz Networking: Mobility, Beamforming, and Frame Level Operation from Theory to Practice," *IEEE Transactions on Mobile Computing*, vol. 18, no. 10, pp. 2217–2230, Oct. 2019.

[18] D. Steinmetzer, D. Wegemer, M. Schulz, J. Widmer, and M. Hollick, "Compressive Millimeter-Wave Sector Selection in Off-the-Shelf IEEE 802.11ad Devices," in *Proceedings of the 13th International Conference on Emerging Networking EXperiments and Technologies*, ser. CoNEXT '17, New York, NY, USA, 2017.

[19] S. K. Saha, H. Assasa, A. Loch, N. M. Prakash, R. Shyamsunder, S. Aggarwal, D. Steinmetzer, D. Koutsonikolas, J. Widmer, and M. Hollick, "Fast and Infuriating: Performance and Pitfalls of 60 GHz WLANs Based on Consumer-Grade Hardware," in *2018 15th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON)*, Hong Kong, China, 2018.