Optimizing mmWave Spatial Reuse: Signal-to-Interference Aware Beamtraining

Guillermo Bielsa
IMDEA Networks Institute and
Universidad Carlos III de Madrid, Spain
guillermo.bielsa@imdea.org

Adrian Loch
IMDEA Networks Institute
Madrid, Spain
adrian.loch@imdea.org

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

***Abstract*—Both IEEE 802.15.3c and 802.11ad use the 60 GHz band for high datarate Wireless Personal/Local Area Network applications. These millimeter-wave communications use very directional antennas since the small wavelength allows to integrate many small antenna elements to form a beamforming antenna array, enabling very high spatial reuse as can be found in dense indoor and IoT settings. However, earlier work shows that current mmWave systems are not as directional as theory would suggest, with significant interference that may prevent spatial reuse. In this work, we propose a centralized system that allows the network to carry out the beamtraining process not only to maximize signal power, but also taking into account other stations in order to minimize interference. This system is designed to work with unmodified clients. We implement and validate our system on commercial off-the-shelf 60 GHz hardware, achieving an average throughput gain of 24.67% for TCP traffic, and up to a twofold throughput gain in specific cases.**

I. Introduction

Millimeter wave (mmWave) technology is available on the market and is being adopted by users. Given the high popularity of these devices, mmWave chipsets are getting cheaper, opening the path to new applications. The advantage of mmWave devices lies in their very high frequency and small wavelength, which allows them to integrate many small antenna elements creating a phased antenna array. This, in turn, allows for very directional beam forming to overcome the high path loss at these frequencies. Multiple standards, such as IEEE 802.15.3c or IEEE 802.11ad, take advantage of the mmWave characteristics as the high datarate and directionality make them a perfect fit for dense wireless deployments with very high datarates. Future use cases include dense deployments of Internet-of-Things (IoT) devices, such as surveillance cameras or streaming services.

In theory, mmWave devices with directional antenna beam patterns can achieve very high spatial reuse with many simultaneous links in the same area, providing almost interference-free communications. In contrast, the omni-directional transmissions and rich multi-path environment in legacy bands require a medium access control that avoids concurrent transmissions to prevent collisions. In practice, however, the difference between mmWave and lower frequency bands is not that pronounced. For simplicity and to lower manufacturing costs, current mmWave commercial off-the-shelf (COTS) use antenna arrays with beam shapes that are much wider and with many more side lobes than theory suggests [1], [2], which creates enough interference so as to prevent spatial reuse in most practical scenarios.

Beamtraining is the process where mmWave nodes select the most suitable beam pattern to be used from its codebook, which in the case of IEEE 802.11ad, is the one that achieves the highest Signal-to-Noise Ratio (SNR). In this paper we propose a novel approach for beamtraining that takes into account not only the SNR, but the possible interference that a beam pattern incurs. To this end, we design a centralized coordinated system that chooses the most efficient beam pattern for each Station (STA) that improves throughput for all the users. Our system requires only changes to the Access Points (APs) and works with unmodified clients.

There are a couple of theoretical and simulation-validated works that tackle the low spatial reuse on mmWave networks: [3] develops a scheduler to provide appropriate concurrent transmissions based on mmWave theoretical characteristics. [4] investigate the collision probability, showing that mmWave networks should improve the conventional collision avoidance procedure. [5] studies the overhead of aligning the Tx-beam pattern together with traffic scheduling on each time slot. [6] studies a hybrid beamforming optimization subject to Signal-to-Interference Ratio (SIR).

The main contributions of this paper are the following:

- We design four beamtraining algorithms to increase aggregated throughput by improving spatial reuse.
- We implement the different algorithms in COTS devices using alternatively the default omni-directional Rx-beam pattern or modified Rx-beam pattern chosen to be the same as the Tx-beam pattern. We modify the firmware of the devices to extract all the information needed to select and fix the utilized beam patterns.
- We first analyze the performance of the different algorithms in a proof of concept scenario and then evaluate them in a real-world environment where we test how the different algorithms behave for different positions and link combinations in an open-space office. We check both Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) as well as uplink, downlink and bidirectional traffic.

The paper is organized as follows: We explain the different beamtraining mechanisms in Section II. The experimental setup is presented in Section III. The evaluation of the different algorithms and scenarios is carried out in Section IV and

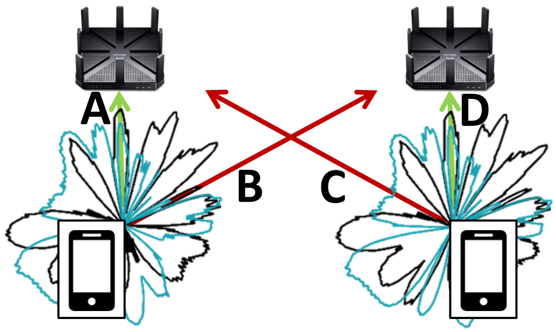

Fig. 1: Beamtraining example: A and D wants to be maximized while B and C minimized. Real beam patterns taken from [2].

discussed in Section V. Section VI concludes the paper.

II. Beamtraining Mechanisms

In this section we explain the different beamtraining algorithms for COTS mmWave devices. Instead of using the strongest beam pattern to reach their APs, the STAs' beam patterns will be chosen to balance the received power at the AP and the interference created to other links. Furthermore, we aim to maximize the aggregated throughput while taking into account fairness, to avoid possible starvation for some of the STAs. The algorithms work with unmodified STAs, requiring only a network of APs with centralized control. This way, any device entering the network can take advantage of this.

In order to establish directional communications following the IEEE 802.11ad standard, the nodes have to perform a beamtraining procedure which consists on the following: first, the AP tests all the possible Tx-beam patterns of its codebook with the initiator sector sweep, while the STAs listen with omni-directional Rx-beam pattern [7]. Then, the STAs reply similarly, sending replies over all their possible Tx-beam patterns as the responder sector sweep. These replies also include which one was the best AP beam pattern. Finally, the AP acknowledges which is the best beam pattern to be used by the STA. After this process, both ends of the link know which are the beam patterns leading to the highest SNR.

Our system needs to do an extra step in order to obtain all the information from the different STAs in the following way: All of the STAs follow the IEEE 802.11ad beamtraining process, sending frames through the different beam patterns. In parallel, *all* the APs overhear these frames and store the received power from every STA. Once all STAs finish their beamtraining, the APs will exchange among them the received power information, allowing them to know how much their STAs are interfering with other APs. Once the APs compute the most suitable beam patterns for their STAs, they send feedback to their STAs specifying which beam pattern to use.

Figure 1 shows a simple example consisting on two APs and two STAs. Each of the STAs can use two different beam patterns (black and light blue), which determine the power with which the signals are received at both APs.

The problem that we want to solve in Figure 1 is the following: the power arriving from STA_1 to AP_1 (named A) and the power arriving from STA_2 to AP_2 (named D) should be maximized, while the interference, which is the power arriving from STA_1 to AP_2 (named B) and the power arriving from STA_2 to AP_1 (named C) should be minimized. Following this example, STA_1 should use the light blue beam pattern. Even though the black beam pattern results in a slightly higher power towards AP_1, it also generates much more interference at AP_2. This interference is negligible if STA_1 uses the light blue beam pattern. On the other hand, STA_2 should use the black beam pattern, as this is the one that results in higher power at AP_2 and lower interference at AP_1.

To solve the problem of beam pattern selection, we investigate four different mechanisms and compare them with the baseline IEEE 802.11ad protocol:

- **IEEE 802.11ad (baseline):** The IEEE 802.11ad standard does not take interference into account and simply selects the beam pattern that provides the highest power to their pairing node. This is maximizing A and D.
- **Weighted SIR Fairness:** This mechanism takes into account all the combinations of beam patterns of all STAs and choose the ones maximizing the sum of SIRs, i.e., $(A - C) + (D - B)$. This is optimal for total channel capacity calculated according to the Shannon-Hartley theorem.
  This mechanism could lead into low SNR STAs meaning those links would disconnect. To obtain a fair allocation, we apply the α-fairness method as $(A-C)^{1-\alpha} + (D-B)^{1-\alpha}$, where α is chosen between 0 and 1. The higher the value, the more fair the allocation. In our implementation we choose an α of 0.99.
- **Argmaxmin:** From all the combinations of beam patterns, the 'Argmaxmin' mechanism selects the combination that maximizes the minimum SIR, i.e., the one where $\min(A-C, D-B)$ is largest. This provides fairness by maximizing the rate of the weakest link.
- **Power Threshold:** For each STA, this mechanism determines the set of beam patterns for which A is larger than a given threshold (corresponding to a Modulation and Coding Scheme (MCS)), and from this set selects the beam pattern that minimizes the interference B. This is repeated independently for all STAs. If there is no group of beam patterns that satisfy this threshold, it decreases to the preceding MCS.
  Since this selection could imply link starvation, we apply the α-fairness method as before.
- **Interference Threshold:** This mechanism checks for which beam patterns the interference B is below a threshold (typically selected to be the carrier sense threshold to enable concurrent communication), and from those selects the one that maximizes A. This is repeated for all STAs to find the beam pattern that maximizes D given that C must be below the threshold. If the threshold cannot be satisfied, default beamtraining is applied.

III. Experimental Setup

In this section, we first discuss the implementation details, from the custom firmware installation to the beam pattern selection. We then present the methodology as well as the different scenarios in which we test our mechanisms.

A. Router Implementation

We use the TP-Link Talon AD7200 router for our system, which follows the full IEEE 802.11ad standard. This router is equipped with the QCA9500 60 GHz chipset, which comes with a 32 element antenna array and a predefined codebook of 34 Tx-beam patterns. For the reception, it relies on a quasi omni-directional Rx-beam pattern.

With the default firmware running on the QCA9500 60 GHz chipset we cannot extract or set any information regarding the beamtraining . We thus use the Talon Tools framework [8] to:

1) Obtain beamtraining information: We can access the IEEE 802.11ad driver and record the received power for each of the beam patterns received by the router, together with the MAC address of the transmitting node.
2) Set the Tx-beam pattern to be used: We can fix the corresponding index of the utilized beam pattern.
3) Modify the omni-directional Rx-beam pattern to one of the predefined beam pattern from the codebook: We can load one beam pattern from the codebook and write it to the chipset to be used as Rx-beam pattern.

This way, the mmWave APs can share how much energy they are receiving from each of the different Tx-beam patterns from each of the different STAs, and choose the pattern which is best for each STA to use according to our algorithms.

B. Methodology

The use of our mechanisms in a real world system is simple. The only constraint is that the different APs are connected and have a minimum processing power, which is usually satisfied for example in managed Wireless Local Area Network (WLAN) deployments. Otherwise, the APs can share their data wirelessly.

For our experimental analysis, we have connected all our APs to a LAN and carried all the beamtraining mechanisms and control of the routers from a central server. In order to measure the performance of the different algorithms, we perform 30 second iperf3 measurements for TCP and UDP traffic. We also modify the traffic flow direction, having uplink, downlink and bidirectional communications. We will first run the measurements with the default beamtraining, then iterate for the different sets of beam patterns, fixing first the Tx-beam pattern with omni-directional Rx-beam pattern and then fixing both Tx and Rx-beam patterns to be equal.

C. Scenarios

All the algorithms explained in Section II are tested in two different scenarios. First, we validate the mechanisms in an empty auditorium and then check their behavior in the real open-space office environment seen in Figure 2. Given the Carrier Sense Multiple Access (CSMA) used by IEEE 802.11ad networks, we have simplified our system to have AP-STA pairs (rather than having multiple STAs per AP). An AP with several STAs would lead to similar results given the multiplexing among different STAs, but would significantly increase the complexity of our deployment.

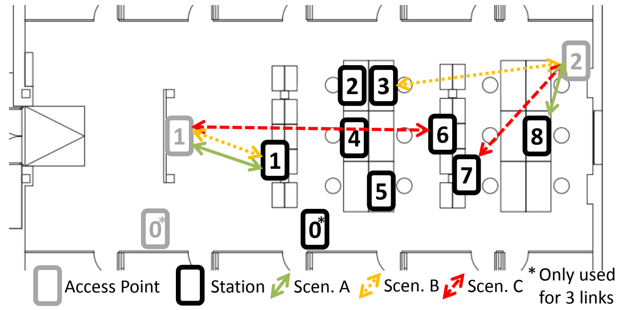

Fig. 2: Office environment setup: Grey boxes represent APs while black boxes STAs. Scenarios A, B and C represents high, medium and low gain scenarios. AP_0 and STA_0 marked with a star are only used for the three link setup.

1) Empty Auditorium: To validate our mechanisms and evaluate their performance we chose an empty auditorium of size 11 x 21 meters where we place two parallel links. For both links, the distance between AP and STA is equal to 3.2 m and the distance between the links increases from 1.6 to 9.6 meters with steps of 1.6 m for a total of 6 measured positions.

2) Office Environment: Our second environment is an open-space office environment that is in active use, with size 7.4 x 13.5 meters as shown in Figure 2. The APs are placed under the ceiling while the STAs are placed on top of regular desks in the working places of the office. In this environment we first do an extensive study of the mechanisms with two parallel links (using AP_1 and AP_2) and then study a few examples with three parallel links (where we add the AP_0–STA_0 link).

IV. Evaluation

In this section we evaluate the behavior of the four different mechanisms discussed in Section II. We first validate the algorithms in a simple setup and evaluate a more complex office scenario. For the latter we study both setups with two and three parallel links.

A. Empty Auditorium

The performance of the four different algorithms in the auditorium setup can be seen in Figure 3. As expected, in general the gains for the different algorithms increase together with the link distance, implying that with larger distances it is easier to achieve spatial reuse.

We also observe that UDP traffic obtains higher gains than TCP traffic; gains which are enlarged even more when choosing the Rx-beam pattern to be the same as the Tx-beam pattern. In contrast, for TCP traffic happens the opposite, and higher gains are achieved when only the Tx-beam pattern is fixed and the omni-directional pattern is used for reception.

Comparing the different algorithms in this validation measurement set, we observe the following characteristics for each of the different mechanisms:

- **Weighted SIR Fairness:** This mechanism achieves large gains for UDP traffic when the links are far apart.
- **Argmaxmin:** This mechanism provides some level of fairness for the different protocols. It has similar gains

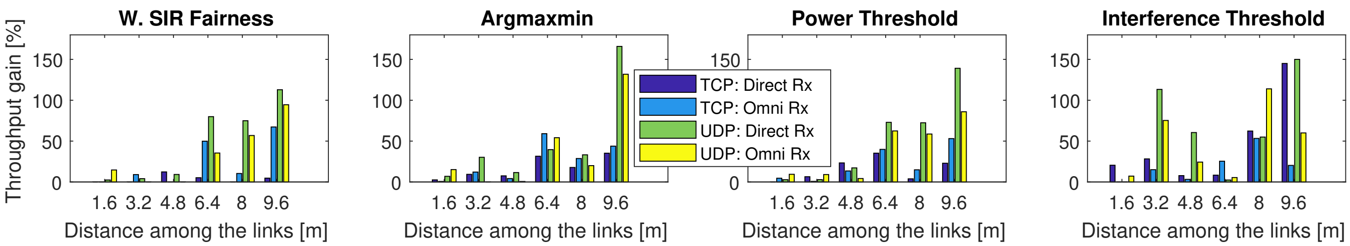

Fig. 3: Mechanisms validation result: Empty Auditorium. Throughput gains for TCP and UDP traffic, when fixing Tx- and Rx-beam pattern and when fixing only Tx-beam pattern with the default omnidirectional Rx-beam pattern.

TABLE I: Average gain for TCP traffic: Bidirectional (BD), Downlink (DL) and Uplink (UL) traffic. For directional Rx-beam pattern (Direc Rx), and for omni-directional Rx-beam pattern (Omni Rx).

	W. SIR Fair.	Argmaxmin	Power Th.	Int. Th.
BD Direc Rx	7.01%	7.65%	24.67%	22.71%
BD Omni Rx	13.86%	23.83%	21.56%	12.76%
DL Direc Rx	0.83%	5.19%	7.72%	9.94%
DL Omni Rx	5.00%	9.03%	1.81%	8.16%
UL Direc Rx	-2.53%	1.36%	2.67%	-1.45%
UL Omni Rx	-7.91%	4.81%	5.26%	-6.25%

TABLE II: Average gain for UDP traffic: Bidirectional (BD), Downlink (DL) and Uplink (UL) traffic. For directional Rx-beam pattern (Direc Rx), and for omni-directional Rx-beam pattern (Omni Rx).

	W. SIR Fair.	Argmaxmin	Power Th.	Interf. Th.
BD Direc Rx	14.89%	5.48%	18.41%	1.62%
BD Omni Rx	11.37%	8.35%	23.70%	2.60%
DL Direc Rx	15.36%	-0.07%	15.14%	4.64%
DL Omni Rx	15.47%	3.50%	16.79%	-0.89%
UL Direc Rx	-0.37%	-2.55%	6.68%	-3.58%
UL Omni Rx	-2.90%	-1.42%	7.61%	-13.25%

for TCP and UDP traffic, both for directional and omni-directional Rx-beam pattern. Gains increase up to 166% and 132% for UDP at 9.6 m.

- **Power Threshold:** This mechanism is the most promising and stable one of the four studied. For UDP traffic we achieve higher gains with directional Rx-beam patterns, and for TCP we achieve higher gains with the omni-directional Rx-beam pattern.
- **Interference Threshold:** This mechanism provides the largest gains among the algorithms when links are close, but it is not reliable for longer distances.

B. *Office Environment: 2 Links*

After validating our measurements in the simple auditorium scenario, we test our algorithms in the more realistic office environment shown in Figure 2.

Table I shows the average gain for each of the algorithms for TCP traffic. From this table we see that the most efficient algorithm is the 'Power Threshold' algorithm, which achieves the maximum average gain value of 24.67% for bidirectional communications with directional Rx-beam patterns. In fact, for all of the studied algorithms the highest gains are achieved in the bidirectional case. In bidirectional communications, the channel access times become critical as the medium usage increases due to the long frames from both nodes (as opposed to the regular data-ACK patterns of uni-directional traffic), obtaining higher gains with the use of our algorithms.

We can also see in Table I that the 'Argmaxmin' together with the 'Power threshold' are the only two methods which have positive gains for all the communication directions. In contrast, 'Weighted SIR Fairness' and 'Interference Threshold' have some sub-10% losses for uplink traffic.

Table II summarizes the measurements for UDP traffic. As for the case of TCP traffic, the best performing algorithm is again the 'Power Threshold', having a maximum average gain of 23.70% of aggregated throughput. It is the only algorithm that has positive gains for all traffic directions with UDP traffic, as the others have some performance issues with uplink traffic, and some 1% losses in downlink.

In order to see how the algorithms behave for the different location combinations, Figure 4 shows the Cumulative Distribution Function (CDF) for the different combinations in the bidirectional case. There are some link combinations where the gains are very high, but others where there are losses. The mechanism that finds the highest gains is the 'Argmaxmin' mechanism, but in turn, it is also the algorithm that has higher losses. It is thus a 'risky' algorithm whose performance very much depends on the specific link configuration.

As depicted in the previous tables I and II, we confirm that the best algorithm is 'Power Threshold', with gains for all the different link combinations and traffic types with the exception of two outliers: when using a directional Rx-beam pattern with UDP traffic and when using omni-directional Rx-beam pattern for TCP. These gains are due to the nature of the algorithm, which will always choose beam pattern combinations providing a sufficiently high MCS while minimizing interference.

From the CDFs seen in Figure 4, there are link combinations where the algorithms work efficiently, and others where they do not. Figure 2 shows how STAs are spread around the scenario. Here, we mark three links, one as good (A), one sub-average (B) and one as bad (C). In Figure 5 we can see the three examples where we can compare the aggregated throughput and the gain of each of the mechanisms with respect to the default beamtraining mechanism for TCP traffic. For the good combination, named A, all the algorithms have gains greater than 20%, with the maximum gain of 53% for the 'Power Threshold' algorithm for omni-directional Rx-beam pattern. With this network configuration, it is easy to obtain

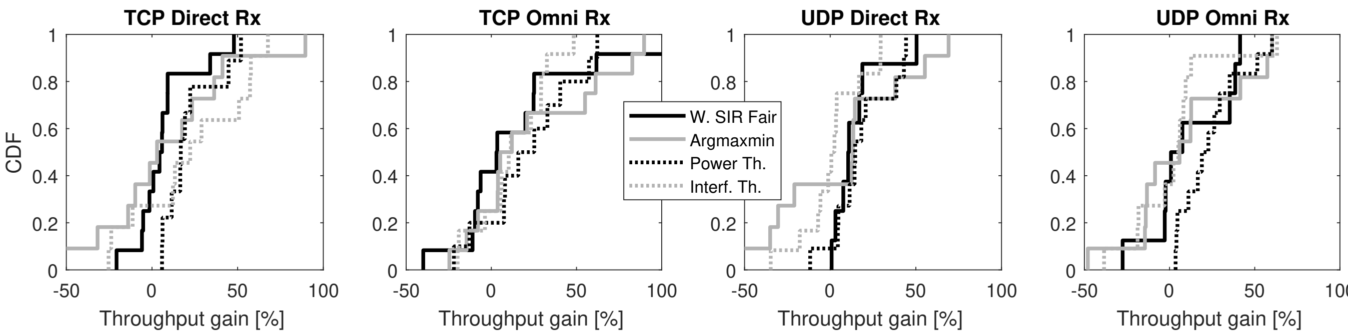

Fig. 4: CDF of the throughput gains for bidirectional TCP and UDP traffic, when using directional Rx-beam pattern and when using the default omni-directional Rx-beam pattern.

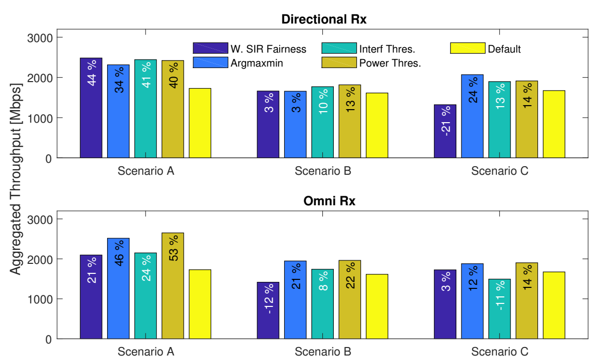

Fig. 5: Aggregated throughput for three measured positions with two parallel links, also showing the percentage gain.

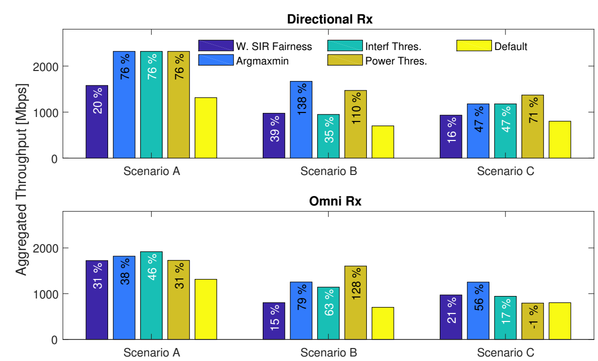

Fig. 6: Aggregated throughput for three measured positions with three parallel links, also showing the percentage gain.

gains as the STAs are far apart from each other, facilitating the selection of beam patterns that do not interfere.

Scenario B shows a below-average combination, achieving gains of 13% and 10% for the Power and Interference threshold techniques for directional Rx-beam pattern. For the omni-directional case we have gains of 22% for the 'Power Threshold' mechanism and 21% for 'Argmaxmin'.

Next, we can see a connection example with low gains, named C. For this network configuration it is difficult to obtain gains as the STAs are close to each other, and most of the beam patterns generate too much interfere to obtain spatial reuse. 'Power Threshold' has gains of 14% making it once again the most stable algorithm while 'Argmaxmin' achieves 24% gains for directional Rx-beam pattern and 12% for omni-directional. 'Interference Threshold' achieves gains of 13% and losses of 11% for directional and omni-directional Rx-beam pattern, while 'Weighted SIR Fairness' has losses.

With these examples we can see that when the STAs are very close to each other the gains that can be achieved are very low since no matter which beam pattern is chosen, interference is high. In contrast, links that are far apart can achieve very high gains, as it is easier to find patterns that minimize interference to neighboring nodes.

C. Office Environment: 3 Links

We now set up a third mmWave link in our scenario as seen in Figure 2. Figure 6 shows how the mechanisms behave for three parallel TCP links for three different link configurations for bidirectional traffic. For these cases, AP_0 always connects with STA_0, AP_1 always connects with STA_8, and AP_1 connects with STA_1, STA_2 or STA_4, naming those Scenarios A, B and C, respectively.

Again, similarly to the results shown in Section IV-B for 2 parallel links, we can see that the 'Power Threshold' algorithm outperforms the others for most of the cases. For scenario A, where AP_1 is connected to STA_1, we have gains of 76% for all the mechanisms for directional Rx-beam pattern except for 'Weighted SIR Fairness', which only has 20% gain. For omni-directional Rx-beam pattern, the gains decrease for the three best mechanisms, having now gains of 38%, 46% and 31% respectively, whereas the gain obtained by the 'Weighted SIR Fairness' method increases to 31%.

For scenario B with AP_1 connected to STA_2, the 'Argmaxmin' and 'Power Threshold' algorithms again outperform the others having 138% and 110% gains for directional Rx-beam pattern. For omni-directional Rx-beam pattern, the 'Power Threshold' algorithm gives the highest gain of 128%.

In the last scenario C, where AP_1 is connected to STA_4, we can see that the highest gain is again for the 'Power Threshold' mechanism for directional Rx-beam pattern, resulting into a gain of 71%. For omni-directional Rx-beam pattern, we have a maximum gain of 56% with the 'Argmaxmin' mechanism.

As for the case of two links, the algorithm that consistently delivers good performance is the 'Power Threshold' algorithm, whereas other algorithms might have higher gains in some specific configurations but underperform in others.

Our three link measurements shows that the aggregated throughput decreases when more than two links work in parallel in the same collision domain. Even though our algorithms provide large gains, having so many nodes increase the delays on the medium access. This result in a three link scenario with lower aggregated throughput than a two link scenario.

V. Discussion

During the development and validation of these mechanisms we found different insights regarding the possibilities of spatial sharing for 60 GHz COTS devices.

Weighted SIR Fairness leads to the highest Signal-to-Interference ratio. This mechanism results into the optimum channel capacity for a wireless communication system given by the Shannon-Hartley capacity theorem. But reality is different for our system, as the IEEE 802.11ad standard does carrier sensing, which implies that even if we could achieve very high Signal-to-Interference ratios which could result in very high MCS, if a node receives interference above the carrier sense threshold it needs to wait until it can access the medium. This beamtraining mechanism can lead to rate losses, as the link SNR might get reduced with the change of beam pattern but still receive too much interference to achieve the desired spatial reuse. This occurs for half of the cases, which have sub-25% bitrate losses. However, gains are achieved the rest of the time when using this mechanism.

Argmaxmin chooses the beam pattern combination in a way that the SIR of worst link in the whole scenario is maximized. This selection can lead us to an average poor scenario with poor connections among the links. Due to this, the different nodes might use worse MCS values and may still carrier sense with other nodes. In contrast, the default IEEE 802.11ad algorithm will carrier sense, but it will use stronger link budgets resulting in higher MCS and bitrates. Nevertheless, the gains of this mechanism are usually positive.

Power Threshold is the most reliable mechanisms from the ones developed. Here, we select a given group of beam patterns that provide a sufficiently high SNR and thus MCS, and from those select the one that produces the least interference. This way, the MCS is often equal to or not much below the one of the default beamtraining mechanism, but with the advantage that the devices will carrier sense less often.

Interference Threshold can fail due to low directional beam patterns. While this mechanism should provide very efficient communications, the low-directionality of the current beam pattern codebook makes very difficult to find a beam pattern capable of providing good communication conditions while having zero-interference with neighboring nodes. Also, as we developed our beamtraining mechanisms to work for unmodified clients, it is impossible to achieve 100% spatial reutilization as only one link direction is optimized.

Other important insights that we discovered during this study are the following: For our 2-link measurement scenarios, we have an upper-bound of 3000 Mbps of aggregated throughput, which, given that the Talon AD7200 routers can achieve individual rates of up-to 2000 Mbps, means that we never achieve full spatial reuse. As a consequence, there is still a large performance gap to improve spatial sharing for mmWave devices, even though we have shown gains larger than 100%. Related with this previous observation, we have seen that the aggregated throughput in our three parallel link scenario is lower than the one in our two parallel link scenario. Despite some spatial reuse, the increased overhead due to the carrier sensing of a larger number of nodes and the medium access delays, the overall medium efficiency decreases. Nevertheless, our algorithms obtain very large gains for this 3-link scenario.

VI. Conclusions

mmWave devices choose the strongest communication beam pattern from a predefined codebook without taking into consideration other parallel links that are in the same collision domain. This, together with the non-ideal shapes of the beam patterns, prevents efficient spatial sharing among the nodes, reducing aggregated and individual throughput. We thus implement four different beamtraining schemes which select the most efficient beam patterns, not only in terms of received power, but also in terms of interference. This system is implemented on the AP side and is transparent for the STAs, making possible for any unmodified STA joining the wireless network to benefit from the system. It provides gains of up to 100% in some cases, and average gains of 25% for the aggregated throughput of two links with TCP and bidirectional transmissions. We implement and test our four beamtraining mechanisms in mmWave IEEE 802.11ad COTS devices. We first validate the implementation and mechanisms in a simple scenario and then test them in a real world open-area office environment. In this second scenario we test our algorithms for 2 and 3 parallel links.

Acks

This work is partially supported by the European Research Council grant ERC CoG 617721 and the Region of Madrid through the TAPIR-CM program (S2018/TCS-4496).

References

[1] 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*, 2018.

[2] D. Steinmetzer, D. Wegemer, M. Schulz, J. Widmer, and M. Hollick, "Compressive millimeter-wave sector selection in o-the-shelf IEEE 802.11 ad devices," in *Proc. ACM CoNEXT*, 2017.

[3] L. X. Cai, L. Cai, X. Shen, and J. W. Mark, "Rex: A randomized exclusive region based scheduling scheme for mmwave wpans with directional antenna," *IEEE Transactions on Wireless Communications*, vol. 9, no. 1, 2010.

[4] H. Shokri-Ghadikolaei and C. Fischione, "The transitional behavior of interference in millimeter wave networks and its impact on medium access control," *IEEE Transactions on Communications*, vol. 64, no. 2, pp. 723–740, Feb 2016.

[5] H. Shokri-Ghadikolaei, L. Gkatzikis, and C. Fischione, "Beam-searching and transmission scheduling in millimeter wave communications," in *2015 IEEE International Conference on Communications (ICC)*, June 2015, pp. 1292–1297.

[6] S.-H. Wu, K.-Y. Lin, and L.-K. Chiu, "Hybrid beamforming using convex optimization for sdma in millimeter wave radio," in *Personal, Indoor and Mobile Radio Communications, 2009 IEEE 20th International Symposium on*. IEEE, 2009, pp. 823–827.

[7] 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.

[8] D. Steinmetzer, D. Wegemer, and M. Hollick. Talon tools: The framework for practical ieee 802.11ad research. [Online]. Available: https://seemoo.de/talon-tools/