An Experimental Study of Rate and Beam Adaptation in 60 GHz WLANs

Shivang Aggarwal, Urjit Satish Sardesai, Viral Sinha, Dimitrios Koutsonikolas
University at Buffalo, The State University of New York, NY, USA
{shivanga,urjitsat,viralsin,dimitrio}@buffalo.edu

ABSTRACT

In this paper, we conduct an extensive experimental study of the two primary link adaptation mechanisms in 60 GHz WLANs, namely rate adaptation and beam adaptation, using a large data set collected from a 60 GHz software-defined radio testbed. First, we compare the effectiveness of the two mechanisms in a variety of indoor environments and scenarios, including linear and angular displacement, mobility, blockage, and interference. Next, we study the effectiveness of two rate adaptation approaches – SNR-based rate adaptation, which has been proposed by recent works, and a learning-based approach using PHY layer information. Our results show that the former performs poorly in practical scenarios, while the latter is promising, especially when combined with online training. Finally, we explore the effectiveness of maintaining backup beams to speedup link recovery and reduce the beam training overhead. We show that this heuristic fails in scenarios involving angular displacement on the receiver side but is quite effective in most other scenarios.

CCS CONCEPTS

• **Networks** → **Network protocols**; **Network performance analysis**; **Wireless local area networks**.

ACM Reference Format:
Shivang Aggarwal, Urjit Satish Sardesai, Viral Sinha, Dimitrios Koutsonikolas. 2020. An Experimental Study of Rate and Beam Adaptation in 60 GHz WLANs. In *23rd International ACM Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM '20), November 16–20, 2020, Alicante, Spain.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3416010.3423219

1 INTRODUCTION

Rate adaptation (RA) – dynamically adjusting the modulation and coding scheme (MCS) based on the channel quality – has been traditionally considered as the main link adaptation mechanism in 802.11-based WLANs. Since the introduction of the first RA algorithm for 802.11b [16], all 802.11 standards mandate support for multiple rates at the PHY layer but they do not specify a RA algorithm. Consequently, a very large number of RA algorithms have been proposed for legacy WiFi, e.g., [1, 5–7, 19, 21, 32, 33].

MSWiM '20, November 16–20, 2020, Alicante, Spain

ACM ISBN 978-1-4503-8117-8/20/11…$15.00
https://doi.org/10.1145/3416010.3423219

Millimeter-wave (mmWave) wireless, e.g., in the 60 GHz unlicensed frequency band, is fast emerging as the prime candidate technology for providing wireless multi-Gbps data rates. To cope with the high propagation loss in the mmWave frequency bands, mmWave transceivers establish highly directional links. However, high directionality introduces new challenges – vulnerability to blockage and beam misalignment due to mobility. While RA continues to play a critical role in the performance of 60 GHz WLANs, it cannot always address these new challenges alone. Hence, mmWave radios are typically equipped with electronically steerable phased antenna arrays, and employ a second link adaptation mechanism, beam adaptation (BA), also referred to as *beamforming* or *beam searching*, to maintain transmitter-receiver (TX-RX) beam alignment. The 802.11ad/ay 60 GHz WLAN standards provide a well-defined procedure for BA and several other BA algorithms have also been proposed over the past few years, e.g., [8, 11, 26, 29, 35].

In spite of their importance, several aspects of the two link adaptation mechanisms are not well understood in the case of 60 GHz WLANs. First, the standards do not specify when each of the two adaptation mechanisms, RA and BA, should be used or in what order, and wireless chipset vendors resort to simple heuristics [2, 22] to select the right mechanism. It remains unclear whether both mechanisms are equally effective in different scenarios, when one of them suffices to repair a link, and when both should be triggered. Second, in contrast to legacy WiFi, where a very large number of RA algorithms have been proposed over the past 20 years, RA in 60 GHz WLANs has not received due attention. Third, in spite of the very large number of BA algorithms proposed over the past few years, the majority of the proposed algorithms have been evaluated with horn antennas that feature perfect cone-shaped beam patterns and in limited settings. It is unclear how beam adaptation is affected by imperfect beam patterns, featured by commercial phased antenna arrays, and by different indoor environments.

This paper fills this gap by conducting the first large-scale experimental study of the two link adaptation mechanisms using a large data set collected from a 60 GHz software-defined radio (SDR) testbed with phased arrays. The dataset is publicly available at http://bit.ly/60ghz-link-adaptation. We make three contributions:

- We first explore the performance of RA and BA in a variety of indoor WLAN environments and scenarios, including linear and angular displacement, mobility, blockage, and interference, and compare the effectiveness of the two mechanisms. Our results show that none of the two mechanisms performs consistently well alone and often a combination of the two is required for optimal performance.
- Next, we study the effectiveness of two RA approaches – SNR-based RA, which has been proposed by recent works, and a learning-based approach using PHY layer information. We show that the former performs poorly, as it is difficult to obtain a direct

SNR-MCS mapping in practice. In contrast, the latter appears promising, but it requires online training.
- Finally, we turn our attention to BA and study the effectiveness of a simple heuristic: maintaining backup beams to speed up link recovery and reduce the beam training overhead. We show that this heuristic fails when RX angular displacement is involved, but works well in all other scenarios; maintaining the top 2-3 beam pairs as backup pairs is often enough to yield satisfactory performance with minimal beam training overhead.

2 BACKGROUND AND RELATED WORK

Rate Adaptation. RA is the primary link adaptation mechanism in legacy WiFi. The goal of RA is to select a PHY data rate, expressed as a combination of a modulation and a coding scheme (MCS), that matches the observed channel quality at the receiver. Since the RX channel quality is not known on the TX side without explicit feedback, the majority of RA algorithms for legacy WiFi (e.g., [1, 5, 7, 19, 21, 32, 33]) estimate the channel quality using link layer statistics of transmitted frames and employ simple heuristics based on these statistics. The use of SNR and other PHY layer metrics has also been proposed in the literature (e.g., [6, 9, 15, 25, 30]) but has not been used by 802.11 chipset vendors. Some works (e.g., [14, 18, 31]) have also proposed the use of ML to guide rate selection and they have shown that ML-based RA schemes outperform SNR-based or frame-based schemes.

In contrast to legacy WiFi, RA has not been extensively studied in the context of 60 GHz WLANs. Commercial off-the-shelf (COTS) devices employ heuristics similar to those used by legacy WiFi devices, e.g., they lower the MCS until they find a working MCS [2]. A few recent works argued that 60 GHz links are much more stable than legacy WiFi links due to the high directionality, and suggested the use of simple SNR-based RA algorithms via a direct SNR-MCS mapping [8, 10, 28, 36, 37]. However, a more recent work showed experimentally that MCS is only weakly correlated with SNR in 60 GHz WLANs [23]. In this paper, we conduct a detailed study of the effectiveness of SNR-based RA using a much larger dataset than the one in [23], including more locations, as well as blockage and interference scenarios. The work in [17] uses ML for channel classification (LoS vs. NLoS) and leverages this information to guide RA in 802.11ad links. In this paper, we consider for first time an ML-based approach that uses various PHY layer features to directly predict the best MCS (similar to the approach in [18] for legacy WiFi) rather than indirectly inferring it from channel classification and compare it against SNR-based approaches.

Beam Adaptation. The goal of BA is to find the TX-RX beam (sector) pair that maximizes the SNR. A naive approach is to perform a full sector level sweep (SLS), i.e., to test all possible pairs, but the overhead of this approach ($O(N^2)$, where N is the number of available beams) can be prohibitive, especially in the case of a large number of beams. The 802.11ad standard takes a different approach having each side train their TX and RX beams separately [12], which reduces the complexity from $O(N^2)$ down to $O(N)$. Additionally, a large number of recent works have proposed different approaches to further reduce the overhead, e.g., [8, 11, 26, 29, 35]. The majority of these algorithms have either been evaluated via simulations or using horn antennas with perfect cone-shaped beam patterns. In this paper, we perform an extensive study of the efficiency of the baseline BA algorithm in a variety of indoor WLAN scenarios, using phased arrays with imperfect beam patterns. We also evaluate the effectiveness of a simple heuristic – maintaining backup beams, which can be part of different BA algorithms.

Surprisingly, the standard does not specify when each of the two adaptation mechanisms should be triggered, and the problem has been largely overlooked by the research community. COTS 802.11ad devices trigger RA in the case of a missing Block ACK, successively trying lower MCSs, and only resort to BA if a working MCS cannot be found [22]. The authors of [2] point out that this approach may often be suboptimal and propose to first perform BA in the event of link degradation and then RA. The work in [8] proposes a non-standard-compliant approach based on beam sounding, i.e., the use of a short control frame exchange before each data transmission to test the link and determine the right action. Our work is the first to compare the effectiveness of the two link adaptation mechanisms in a variety of indoor environments and scenarios.

Performance Studies A number of experimental studies have explored different aspects of 60 GHz WLANs, such as beamforming [20, 23], link layer performance and the impact of imperfect beam patterns generated by phased arrays [22, 23, 28], frame aggregation [20], interference [20, 34], transport layer aspects [4], energy consumption [24], etc. Our work is complementary to all these studies, exploring aspects of link adaptation that have been largely neglected by previous works.

3 SETUP AND METHODOLOGY

3.1 X60 Testbed

We conducted our measurements using X60 [23] instead of COTS standard-compliant devices, because its SDR architecture provides us with full access to the PHY and MAC layers. X60 is the only 60 GHz SDR-based testbed that combines fully programmable PHY and MAC layers, multi-Gbps data rates, and practical reconfigurable phased arrays. While the testbed does not support the 802.11ad standard, many of its features resemble those of 802.11ad.

(a) Two X60 nodes.

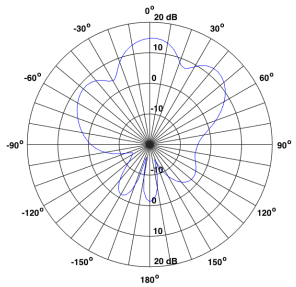

(b) Example azimuth beam pattern.

Figure 1: The X60 testbed.

Each X60 node (Fig. 1a) consists of a NI mmWave transceiver system [13] and a user-configurable phased antenna array from SiBeam. Transmissions take place over a 2 GHz wide channel, same as in 802.11ad. The PHY reference implementation enables 7 Single Carrier (SC) MCSs resulting in data rates from 300 Mbps to 4.75 Gbps, similar to those supported by the SC 802.11ad PHY layer. In contrast to COTS 802.11ad radios that use CSMA, X60 uses TDMA with 10 ms frames divided into 100 slots of 100 μs each. A slot

Table 1: Summary of datasets – # of measurement positions.

	Main Dataset					Test Dataset		
	Total	*Lobby*	*Lab*	*Conf. Room*	*Corridors*	*Total*	*Building 1*	*Building 2*
Displacement	94	22	13	10	49	34	23	11
Blockage	12	4	1	2	5	4	2	2
Interference	12	4	1	2	5	4	2	2
Overall	118	30	15	14	59	42	27	15

consists of 92 codewords, each of which has an attached CRC block. For our study, which only includes single-link experiments and hidden terminal scenarios, where CSMA is not useful, the use of TDMA instead of CSMA does not affect our results. Also, note that structure of an X60 frame resembles an 802.11 AMPDU, consisting of multiple packets, each with its own CRC.

The in-built phased array has 24 elements; 12 each for TX and RX. SiBeam's reference codebook defines 25 beam patterns that can be steered in real-time (electronic switching in $< 1\ \mu s$). The beams are spaced roughly 5° apart in their main lobe, thus spanning around 120° in the azimuth, from -60° to 60°. The 3 dB beamwidth ranges from 25° to 35°, hence, each beam's main lobe overlaps with several neighboring beams. The beam patterns (one example is shown in Fig. 1b) feature large side lobes in addition to the central main lobe, similar to the beam patterns in COTS 60 GHz devices [26].

3.2 Environments

We collected a *main* dataset by taking measurements in multiple locations within a campus building (referred to as Main Building). **Lobby.** This is a large open space with glass panels covering the upper part and metallic sheets covering the lower part of one side and a wall on the other side. It is shown in Fig. 2a along with the various TX and RX positions we used for our measurements. The TX and RX antennas are kept at a height of 1.4 m. **Lab.** This is an $11.8 \times 9.2 \times 3.4$ m^3 space with 4 rows of desks surrounded by metallic storage cabinets and white boards (Fig. 2b). The TX antenna is placed at a height of 2.05 m and the RX antenna at a height of 1.25 m. **Conference Room.** This is a $10.4 \times 6.8 \times 3.2$ m^3 space with a large white board covering one of the walls (Fig. 2c). There are metallic cabinets, a large desk in the center of the room, and many chairs. The TX and RX antennas are placed at a height of 1.4 m. **Corridors.** We performed measurements in 3 corridors of width 1.74 m, 3.2 m, and 6.2 m, with the TX and RX antennas at a height of 1.4 m.

To test if our findings hold across different environments, we also collected two smaller *test* datasets in two different campus buildings. In Building 1, we conducted measurements in a corridor of width 2.5 m with the RX at several distances away from the TX. This building is much older than the Main Building, with walls of different material, and fewer reflective surfaces. In Building 2, we conducted measurements in a wide open area, much larger than the lobby in Fig. 2a. A summary of all the datasets is shown in Table 1.

3.3 Scenarios

We consider three typical scenarios capturing all the factors that can trigger BA or RA due to a drop in channel quality.

Linear and/or angular displacement. In each environment, we fixed the TX position and orientation, we selected an initial position for the RX, and then moved or rotated the RX to cause different levels of signal attenuation due to increased distance, TX/RX misalignment, or both. In all the rotation experiments, we rotated the RX from 0° to -90° and from 0° to 90° in steps of 15° where 0° is the initial orientation at each position.

Lobby: We first fixed the TX at position TX_1 facing west and the RX at position 0 facing east (Fig. 2a). We then moved the RX while keeping the orientation fixed (east) along three directions – backward, lateral, and diagonal – and took measurements at multiple positions in each direction. We also took measurements by rotating the RX at positions 2 and 19. We then took a second set of measurements with the TX fixed at position TX_2 facing west and the RX at 9 positions always facing the TX. *Lab:* The TX was fixed facing west and the RX was placed at 10 different positions, always facing east, starting at position 0. We also took measurements by rotating the RX at positions 2, 5, and 8. *Conference room:* The TX was fixed facing west and the RX was placed at different positions and orientations around the table, shown in Fig. 2c, starting at position 0. Note that for positions 4, 5, 6 and 7, the RX is facing in the same direction as the TX and communication is enabled through reflections. We also performed rotations at positions 0 and 4. *Corridors:* In the narrow corridor, we performed measurements at 17 different RX positions starting from a distance of 2.5 m away from the TX and moving back in steps of 1.25 m, with the TX and RX always facing each other. In the 2 wider corridors, we fixed the TX at two positions and performed measurements at 10 different RX positions at steps of 1.25 m with the TX and RX facing each other at all times. We also performed rotations 5 m, 10 m, and 15 m away from the TX.

Blockage. We hand picked a few representative positions from all the positions described above and repeated the measurements by introducing human blockage on the LOS path between the TX and the RX 1) in the middle between the TX and RX, 2) near the TX and 3) near the RX. Interestingly, we found that link quality degradation is not strongly correlated with the blockage position. For example, the SNR drop caused by blockage in the middle, near the Tx, and near the Rx varies from 1.5-15 dB, 2-12 dB, and 5-12.5 dB, respectively, with median values of 8 dB, 6.5 dB, and 7.5 dB.

Interference. We created hidden terminal scenarios by placing a TP-Link Talon AD7200 router and an Acer P446-M laptop at different positions to create 3 types of interference on the X60 receiver: 1) High interference: the throughput of the X60 link drops by ~80%, 2) Low interference: throughput drops by ~20%, and 3) Medium interference: throughput drops by ~50%. We performed these measurements at the same positions as the blockage experiments.

3.4 Methodology

We use the term *state* to describe every position, orientation, and the presence/absence of blockage or interference. We define the initial state as: the Rx position closest to Tx for each displacement scenario in the lobby, lab, and corridors; Rx position 0 for each displacement scenario in the conference room (Fig. 2c); the 0° Rx orientation for each rotation scenario; and the state before the introduction of blockage or interference for each blockage and interference scenario. All the other states, at which either the RX coordinates, or the RX orientation, or the blockage or interference status are different from the initial state are called *new states*; these are the states where RA or BA is needed to repair the link.

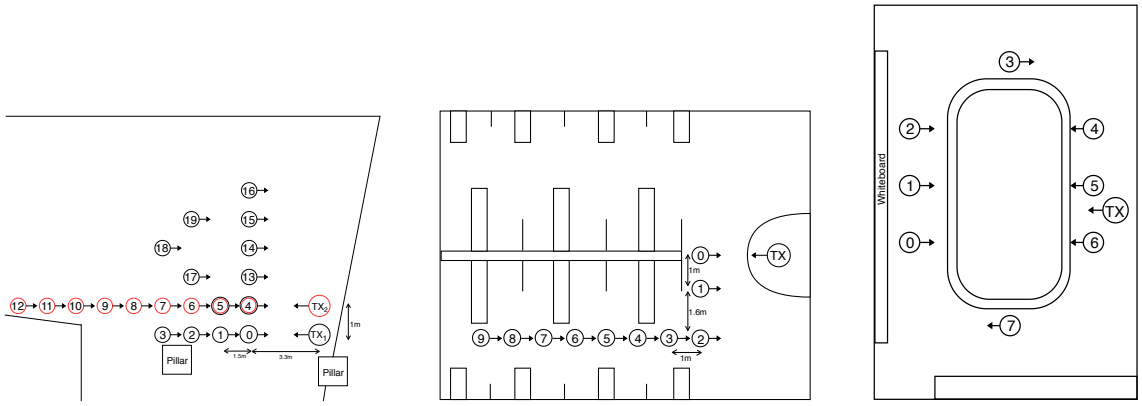

(a) Lobby. (b) Lab. (c) Conference Room.

Figure 2: Environments.

At each state, *we first performed an SLS to collect SNR measurements for all 625 (25 × 25) possible beam pairs and select the best beam pair based on SNR. This process emulates BA using the naive $O(N^2)$ algorithm* described in §2. Then, for the best beam pair, we collected three 1 s PHY layer traces (SNR, Noise level, power delay profile (PDP), codeword delivery ratio (CDR)) and MAC throughput traces for each of the 7 supported MCSs. X60 logs all these metrics for every frame. We also measured offline the time-of-flight (ToF) for the chosen beam pairs at all positions. For all states except for the initial ones, we also collected PHY and throughput traces and ToF values for the beam pair that was the best at the corresponding initial state. *Searching over all the MCSs with the best current/initial beam pair and selecting the one with the highest throughput emulates RA after/before BA at the new state.*

4 EFFECTIVENESS OF LINK ADAPTATION MECHANISMS

In this section, we study the effectiveness of the two adaptation mechanisms under various types of link impairments. We assume that both mechanisms perform exhaustive search and always find the optimal MCS or beam pair, and we defer the study of specific practical approaches to RA and BA in §5 and §6, respectively. We consider large data transfers and evaluate the effectiveness of the two mechanisms in the steady state, ignoring the search overhead.

We assume that after BA, the TX uses the best initial MCS (another option would be to always start with the highest MCS), i.e., the best MCS at the initial state, before the link impairment that triggered BA. However, we found that in many cases, the throughput with the new best beam pair using the best initial MCS is zero, but lowering the MCS by 1 or 2 levels results in high throughput. Hence, we also evaluate the combined effect of triggering BA immediately followed by RA, denoted as BA+RA. Note that, since we are ignoring the overhead here, this approach is always guaranteed to find the optimal configuration.

4.1 Static Clients

Fig. 3a shows boxplots of the throughput achieved with RA, BA, and combined BA+RA in all the scenarios included in the main dataset. Figs. 3b, 3c, 3d break down this result by comparing the performance of the two mechanisms in scenarios involving different types of link impairments: displacement, blockage, and interference, respectively. In all four figures, we also plot the throughput when no adaptation is performed (denoted as NA) as a baseline.

Fig. 3a shows that, without adaptation (NA), the link cannot be restored in the median case, and, with the exception of a few outliers, the throughput is always lower than 500 Mbps. While BA+RA, RA alone, and BA alone achieve a similar maximum throughput (higher than 4 Gbps), the median throughput with BA+RA is much higher than with BA or RA alone (~2 Gbps vs. ~1.1 Gbps and ~0.9 Gbps, respectively). Additionally, BA and RA alone result in a non-working link 25% of the time. We conclude that the performance with either of the two link adaptation mechanisms alone is very far from the optimal.

Fig. 3b shows that RA performs particularly poorly in displacement scenarios, resulting in a median throughput of only 100 Mbps. This is because many of the displacement cases involve angular displacement between the TX and RX and hence, the initial beam pair can no longer work. The performance of BA and BA+RA in such scenarios is similar to the overall performance in Fig. 3a.

Fig. 3c shows that, in case of blockage, neither BA nor RA works alone in the median case; both mechanisms yield a median throughput lower than 100 Mbps and a combination of the two (BA+RA) is needed to restore the link (median throughput of ~1.2 Gbps). When the LOS path is blocked, BA is clearly required to find an alternative path between the TX and RX but because the new best path is an NLOS path through a reflection, the received signal is weaker and hence there is often a need to drop the MCS as well.

In contrast, Fig. 3d shows that in the presence of interference, both mechanisms work well alone in the median case. In particular, RA yields a median throughput of ~2 Gbps, similar to BA+RA (although the 75th percentile is higher for BA+RA compared to RA). Surprisingly, NA also performs well here, yielding a median throughput of ~1.7 Gbps similar to BA. Overall, depending on the amount of interference and the environment, either adaptation mechanism alone can be effective and often there is no need for adaptation at all.

Remark: With the exception of interference scenarios, neither of the two link adaptation mechanisms alone is sufficient to achieve optimal performance.

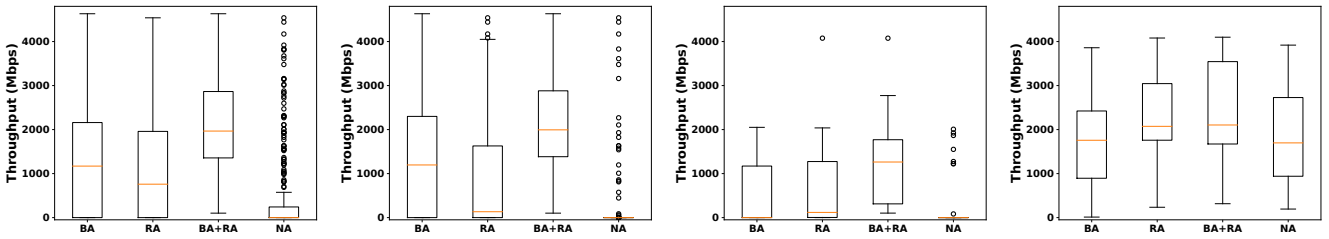

(a) Overall. (b) Displacement. (c) Blockage. (d) Interference.

Figure 3: Performance of link adaptation mechanisms in different scenarios.

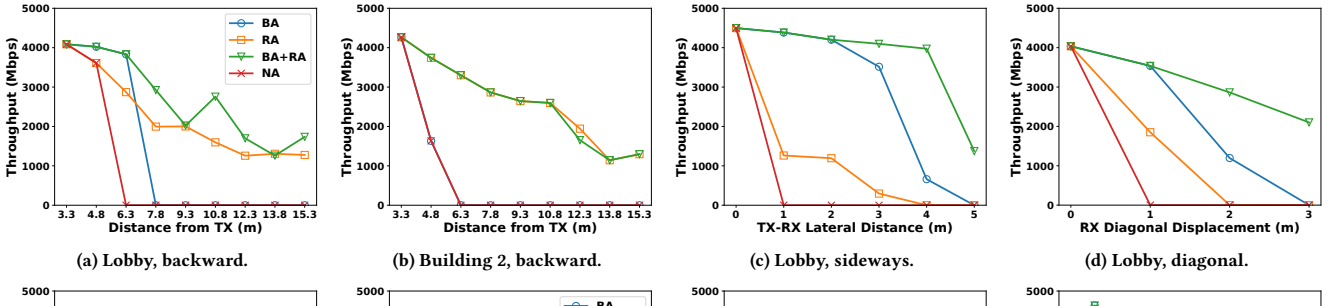

(a) Lobby, backward. (b) Building 2, backward. (c) Lobby, sideways. (d) Lobby, diagonal.

(e) Building 2, diagonal. (f) Narrow corridor, backward. (g) Building 1, backward. (h) Wide corridor, backward.

Figure 4: Performance of link adaptation mechanisms under linear motion.

4.2 Mobile Clients

We emulate linear motion and rotation by considering consecutive RX positions on a straight line and consecutive angular displacements at a fixed position, respectively.

Linear motion. Figs. 4a and 4b show cases where the RX is moving backwards in a wide open space, in different buildings. We observe that, in both buildings, NA and BA stop working after ~6 m whereas RA remains close to the optimal (BA+RA). The signal strength drops while the RX is moving away from the TX, and hence there is a need to drop the MCS to retain a working link. On the other hand, BA selects the same beam pair in this case but uses the initial MCS, which stops working beyond a certain distance.

Fig. 4c shows a case where the RX is moving sideways with respect to the TX ($0 \rightarrow 4 \rightarrow \ldots \rightarrow 16$ in Fig. 2a). Here, the throughput with RA drops by more than 3 Gbps with a motion of just 1 m. On the other hand, BA retains performance close to the optimal till 3 m and then experiences a significant drop because the initial MCS no longer works and hence BA+RA is needed. We make similar observations for the case of moving diagonally ($0 \rightarrow 5 \rightarrow 17 \rightarrow 18$ in Fig. 2a) in Figs. 4d and 4e; again, BA can retain optimal performance for up to 1 m, but as the RX moves further away from the TX, both the beam pair and the MCS need to be adjusted because of the increasing distance and angular separation.

In Figs. 4f, 4g, and 4h, we look at performance of the link adaptation mechanisms while moving backwards in different corridor environments. In all cases, RA performs very close to the optimal, while BA and NA sustain a working link only till 7.5 m. In the narrow corridor, we can see that all link adaptation mechanisms exhibit an inconsistent behavior in terms of their performance as a function of the distance, which is likely due to the waveguide effects (also reported in [22, 28]). This behavior is much less evident in the other 2 wider corridors in the Main Building and Building 1.

Rotation. In Figs. 5a and 5c, the RX is placed in front of the TX in a wide open space in the Main Building and Building 2, respectively. We observe that a rotation of more than 15° breaks the link (NA throughput is 0) and adaptation is required to restore it. RA is effective only for small amounts of rotation; throughput drops below 1 Gbps in the Main Building and to almost 0 in Building 2 for angles greater than 30°. For larger angles, BA is required to maintain the link and retains performance close to the optimal for angles up to 60°. Finally, when the RX rotates 60° or more, BA throughput also drops to zero and BA+RA is needed to sustain a working link.

In Fig. 5b, we see the effect of placing the RX at an angle with respect to the TX. When we rotate the RX in the clockwise (positive, towards TX) direction, BA almost always yields optimal performance whereas, in the anti-clockwise direction, BA stops working

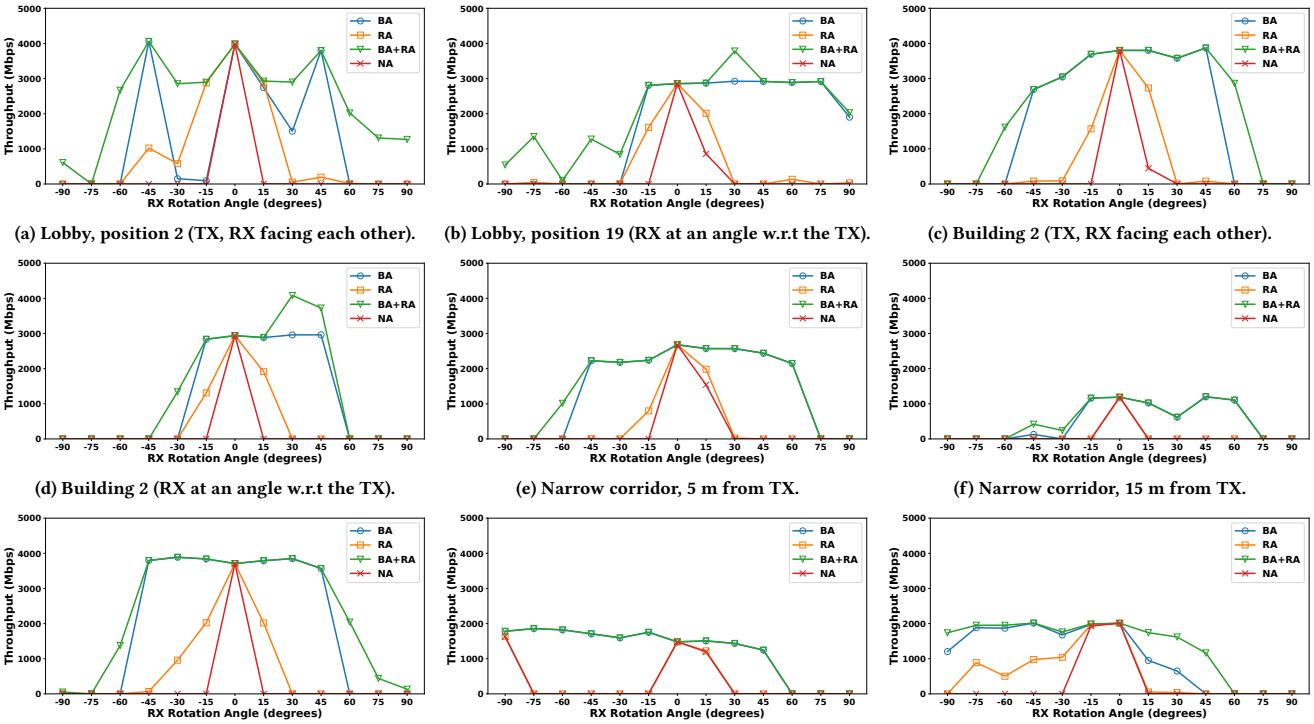

(a) Lobby, position 2 (TX, RX facing each other).
(b) Lobby, position 19 (RX at an angle w.r.t the TX).
(c) Building 2 (TX, RX facing each other).
(d) Building 2 (RX at an angle w.r.t the TX).
(e) Narrow corridor, 5 m from TX.
(f) Narrow corridor, 15 m from TX.
(g) Building 1, 5m from TX.
(h) Conference room, position 4.
(i) Lab, position 5.

Figure 5: Performance of link adaptation mechanisms under rotation.

after a 30° rotation. We also see that BA+RA achieves around 1 Gbps throughput at certain negative angles due to reflections from the wall. We make similar observations regarding the performance of BA in Building 2 (Fig. 5d). However, since there is no wall next to the RX in this case, BA+RA also fails to achieve non-zero throughput in the anti-clockwise direction for angles greater than 45°.

In the lab (Fig. 5i), BA performs nearly optimally but surprisingly RA also achieves Gbps performance even when the RX is rotated by 75° in the anti-clockwise (towards TX) direction. We believe this to be the result of many reflective surfaces present in the lab environment. In the corridors (Figs. 5e, 5f and 5g), RA and NA again do not work when the RX rotates more than 15° while BA results in optimal performance in nearly all cases. Finally, Fig. 5h shows an example (position 4 in the conference room, where the RX faces away from the TX) where none of the two mechanisms alone works even for small angles (their performance is identical to that of NA).

Remark: Overall, with the exception of backward motion with the RX and TX facing each other (i.e., absence of angular displacement), RA alone is ineffective in mobility scenarios. In contrast, BA alone often yields close to optimal performance; a combination of both mechanisms is needed only for very large TX-RX misalignment – in cases of very large angular displacement or rotation angles.

5 RATE ADAPTATION

We explore two approaches to RA: an SNR-based approach, which has been proposed by recent works, and an ML-based approach using PHY layer information.

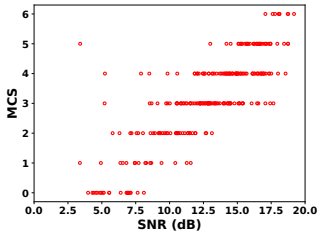

	SNR (dB)	
MCS	*SNR-A*	*SNR-C*
0	<4.9	<8.1
1	4.9 - 5.8	8.1 - 11.6
2	5.8 - 8.5	11.6 - 13.1
3	8.5 - 9.8	13.1 - 17.7
4	9.8 - 13	17.7 - 18.6
5	13 - 17.1	18.6 - 18.7
6	>17.1	>18.7

Figure 6 & Table 2: SNR-MCS relationship.

5.1 SNR-based RA.

Recent experimental works [8, 10, 28, 36, 37] have argued that 60 GHz links are much more stable than legacy WiFi links due to the high directionality, and suggested the use of simple SNR-based RA algorithms via a direct SNR-MCS mapping. However, in our previous work [23], we showed that such a direct SNR-MCS mapping is difficult to find in several displacement scenarios.

In this section, we explore the effectiveness of this approach using a much larger dataset compared to the one in [23], which includes more environments as well as blockage and interference scenarios, in addition to displacement scenarios. Fig. 6 plots for each state in the main dataset the measured SNR on the x-axis and the optimal MCS on the y-axis. We observe that under very low SNR (< 5 dB), MCS 0 is typically the only working MCS (except for a few outliers). However, for SNR values higher than 5 dB, no unique optimal MCS can be selected based on SNR and up to 3 MCSs need to be probed for various SNR ranges (MCS 0-2 for 5-8 dB, MCS 1-3 for 8-11 dB, etc.).

Since there is no direct SNR-MCS mapping, we consider two approaches to build an SNR-MCS table: a *conservative* approach (SNR-C), which sets the threshold for MCS_i equal to the highest observed SNR for MCS_{i-1}, and an *aggressive* one (SNR-A), which sets the threshold for MCS_i equal to the minimum SNR observed for MCS_i. As an example, using Fig. 6, MCS 2 will be selected when the SNR lies between 11.6 - 13.1 dB with the conservative approach and between 5.8 - 8.5 dB with the aggressive approach. Note that for both approaches, we remove outliers (shown in Fig. 6) when we build the mapping table.

5.2 ML-based RA

We explore for first time the use of PHY layer information and ML to predict the best MCS. We consider a number of PHY layer metrics, described below. Note that all these metrics are accessible in today's COTS WiFi drivers [9, 27] and we expect this trend to continue for the 60 GHz drivers.

SNR, Noise level: For each of these two metrics, we consider the difference between the value at the initial state and the value at the new state.

Time of Flight (ToF) Difference. We consider the difference between the ToF at the initial and the current state. ToF increases with the distance and hence, one can expect a non-zero ToF difference when linear displacement is involved (backward, lateral, or diagonal motion) and zero difference in cases involving only rotation.

Multipath-related Metrics. Researchers have proposed the use of Channel State Information (CSI) to guide RA in legacy OFDM-based WiFi systems and have shown that CSI-based algorithms [9, 27] outperform traditional algorithms. To our best knowledge, the use of CSI in RA for 60 GHz WLANs has not been explored before. Since X60 does not support OFDM, we cannot measure directly CSI. Instead, we use PDP, a metric that also captures the impact of multipath propagation but in the time domain instead of the frequency domain. We also use the FFT of the PDP and as an estimate of CSI. Following [27], for each of these metrics, we calculate the similarity between the two instances of the metric (at the initial and current state) in the form of the Pearson correlation coefficient.

Error/Delivery Rate. State-of-the-art RA algorithms for legacy WiFi systems are often based on the subframe error rate (SFER), the fraction of successful MPDUs inside an AMPDU. Since the X60 PHY does not support frame aggregation, we use CDR, the fraction of successful codewords in a 10 ms X60 frame to approximate SFER in WiFi. Note that the length of an X60 frame is same as the maximum allowed AMPDU length in 802.11n/ac and the codeword size (180-1080 bytes for different MCSs) is similar to an MPDU size.

An important parameter here is the *length of the observation window* for each metric at the initial and new state. A longer window can potentially yield more robust estimation of a metric, at the cost of delayed decisions. Since X60 samples each metric every 10 ms, the shortest possible window is 20 ms (1 sample before and 1 sample after introducing a link impairment). We tried different window lengths varying from 20 ms to 2 s and found that the accuracy remains similar for any length of at least 40 ms. In the following, we use a window of 40 ms, which means that the algorithm makes a new decision every 20 ms (or 2 frames in X60). We note that this duration can get shorter depending on the protocol frame duration; e.g., in 802.11ad, where the max frame duration is 2 ms, the algorithm would make decisions every 4 ms.

Table 3: Cross-Validation Results.

	Decision Tree	Random Forest	SVM
Accuracy	0.57	0.69	0.58
F1 Score	0.57	0.68	0.57

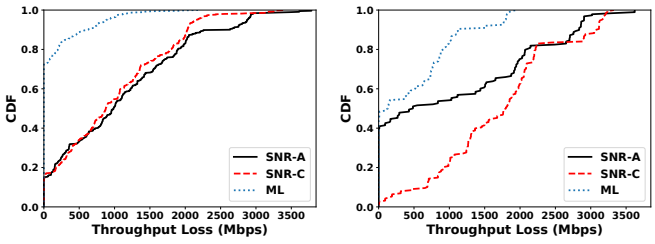

(a) Main Dataset. (b) Testing Dataset.

Figure 7: Performance comparison of the two RA approaches.

We tried 3 popular ML models for the problem of predicting the best MCS using the aforementioned metrics as features: decision trees, random forests, and support vector machines (SVMs). We ran a stratified 5-fold cross validation using the dataset from the Main Building and calculated the accuracy and the weighted F1 score. We averaged the results over 500 random splits of the dataset. The results (with the best combination of the various model parameters, such as the impurity measure, tree depths, regularization parameters, etc.) are shown in Table 3. We observe that none of the three models yields particularly high accuracy; the best accuracy, obtained with Random Forest, is only 0.69.

5.3 Performance Comparison

While in §5.1, §5.2 we saw that none of the two approaches is able to predict the optimal MCS with high accuracy, it is not clear which of the two approaches performs best in terms of throughput. Fig. 7a plots the CDF of the throughput loss (with respect to the optimal MCS) with the SNR-based schemes and the ML-based scheme for the main dataset. We observe that both the aggressive and the conservative SNR-based schemes perform poorly, yielding optimal throughput in only 15-17% of all the cases. In the median case, both schemes result in a throughput loss of ~1 Gbps while the throughput loss is close to 2 Gbps at the 90th percentile. In contrast, the ML-based scheme performs much better and results in optimal throughput in around 70% of the cases. At the 95th percentile, we observe a throughput loss of about 1 Gbps, and the max loss is ~1.8 Gbps, much lower compared to the SNR-based schemes.

Several of the PHY layer metrics capture properties of the multipath channel structure, which is heavily affected by the environment. To evaluate the impact of the environment, we now build the SNR-MCS tables and train the ML model using the main dataset and evaluate the throughput loss using the testing dataset, in Fig. 7b. SNR-C results in optimal throughput only in 2% of the cases; the median throughput loss is ~1.8 Gbps, which is significantly worse than in Fig. 7a. SNR-A surprisingly results in optimal throughput in 40% of the cases for the testing dataset, which is a significant

Table 4: Results after testing in the 2 new buildings.

	Decision Tree	Random Forest	SVM
Accuracy	0.31	0.49	0.48
F1 Score	0.31	0.49	0.46

improvement compared to Fig. 7a. However, the throughput loss reaches almost 1 Gbps around the 50th percentile and 2 Gbps around the 70th percentile. As far as the ML-based scheme is concerned, the prediction accuracy drops further with all three ML models, as shown in Table 4. The drop in the accuracy results in a significant drop in performance compared to Fig. 7a; the loss reaches 1 Gbps around the 85th percentile and close to 2 Gbps in the 95th percentile. However, its performance is optimal in 50% of the cases and remains much better compared to the SNR-based schemes.

Remark: In contrast to conjectures made by previous works, SNR-based RA does not work well in practice in 60 GHz WLANs. On the other hand, the ML-based scheme performs much better, especially when it is trained and tested in the same environment. Although its performance drops significantly when trained in a different environment, we still believe that it is a promising approach, especially when combined with online training algorithms. We plan to further explore this approach as part of our future work.

6 BEAM ADAPTATION

In this section, we turn our attention to BA and explore the effectiveness of a simple heuristic – maintaining a few strong beam pairs as backup pairs. If such an approach works, it can be applied to different link adaptation schemes (e.g., [3, 8]) to reduce the beam training overhead and speed up link recovery.

To evaluate the potential of backup beam pairs to restore a link, we use the Beam Index Difference (BID) metric from [23]. BID for two beam pairs, with indices (TX_1,RX_1) and (TX_2,RX_2), is defined as $BID = |TX_2 - TX_1| + |RX_2 - RX_1|$. We also define the Beam Index Difference between two TX or RX beams as $TXBID = |TX_2 - TX_1|$ and $RXBID = |RX_2 - RX_1|$, respectively.

Fig. 8a plots the CDF of the BID between the new best beam pair (after displacement, blockage, or interference) and the best, 2nd best, and 3rd best beam pair before introducing a link impairment (denoted as BID_1, BID_2, BID_3, respectively) for all the positions in the main dataset. In the median case, the new best beam pair has a BID of 8 from the top 3 initial pairs, and the 90th percentile is more than 16. Further, the new best beam pair is the same as the 2nd or 3rd initial best beam pair less than 5% of the time. Figs. 8b and 8c, which plot the TXBID and RXBID for the same beam pairs, show that both the TX and the RX contribute to this large BID, and the TX contribution is larger. The median and 90th percentile values are 3-4 and 13, respectively, for the TXBID, and 2 and 11, respectively, for the RXBID.

While the BID is large, it is possible that the SNR difference between the new best beam pair and the 2nd or 3rd old best beam pairs is still small, if there are many beam pairs with similar SNR [23]. If that is true, then maintaining the 2nd and 3rd best beam pairs as backup pairs can still work in practice. However, Fig. 8d, which plots the CDF of the SNR difference between the new best beam pair and the initial best, 2nd best, and 3rd best beam pair (denoted as $SNRD_1$, $SNRD_2$, $SNRD_3$, respectively), shows that this is not the case. The median SNR difference is 4.5 dB, which results in 1-2 levels of MCS drop (Fig.6). Further, in about 25% of the cases, the difference is more than 15 dB, which means that selecting one of the top 3 initial best beams after a link impairment would completely fail to restore the link.

We conclude that the simple heuristic of maintaining the second and third best beam pairs as fail-over pairs, which has been proposed in previous works, e.g., [8], does not appear to work in practice. The reason is the fact that a large number of beams around the strongest beam have similar SNRs due to the imperfect beam patterns in COTS devices [23]. This is also implied by Fig. 8a; all three CDFs are very close to each other, suggesting that in most cases, the 2nd and 3rd best beam pair are very close to the best pair, and any link impairment has a similar impact on all three pairs.

There is one caveat to the above result. Even if the 2nd and 3rd best beam pairs were far from the best pair, if the RX rotates or moves in such a way that the relative angle between the TX and RX changes, there is no guarantee that the initial strong beam pairs will still remain strong. In other words, there is no relationship between the best beam pairs before and after such a motion. Hence, in the remainder of this section, we focus on cases where the relative TX-RX angle remains the same and this heuristic might still be useful. To see if that is the case, we plot the BID (Fig. 9) and SNR difference (Fig. 10) between the new best beam and the old top 3 beams for three different types of link impairment: non-angular displacement, blockage, and interference.

In the case of non-angular displacement, the BID between the new best beam and the old top 3 beams is 6-7 in the median case (Fig. 9a), only slightly lower than in Fig. 8a. However, as shown in Fig. 10a, the SNR difference between the old best beam and the new best beam is just 1 dB in the median case and 3 dB in the 95th percentile. In such cases, as we also saw in §4, there is often no need for BA (even if a new better beam pair is indeed discovered, the performance improvement is typically small) and RA alone is able to repair the link.

In blockage scenarios (Figs. 9b and 10b), we observe a median BID of 5 for all 3 old best beams, slightly lower compared to the displacement scenarios. However, selecting one of old three top beam pairs after blockage results in a moderate SNR loss of 2-3 dB in the median case, showing that the heuristic of maintaining the top three sectors can often work in practice. Nonetheless, unlike in the displacement scenario, in 10% of the blockage cases the SNR drop is as high as 15-20 dB resulting in a complete link failure.

Finally, when interference is introduced (Figs. 9c and 10c), the BID varies from 2 to 5 in the median case but the SNR difference is less than 2 dB for any of the old 3 best beams in the median case and around 4 dB in the 90th percentile, suggesting that using one of the old three top beam pairs can work well in most cases.

Number of backup beams. Next, we evaluate the number of the top beam pairs one has to maintain as backup in practice in order to maintain satisfactory performance while keeping the beam searching overhead low. Fig. 11 plots the SNR difference between the best beam pair at the new state (discovered via a full SLS) and the beam pair out of the 'k' old best beam pairs that has the highest SNR at the new state.

In the case of non-angular displacement (Fig. 11a), we observe again that reusing the old best beam pair (k=1) results in a loss of

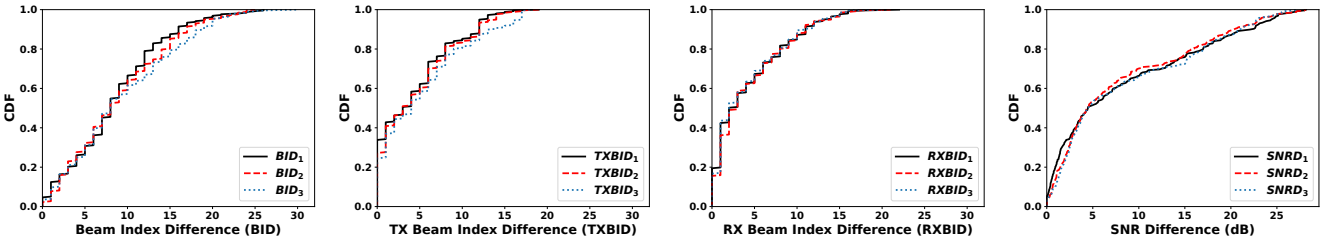

(a) Beam Index Difference (b) TX Beam Index Difference (c) RX Beam Index Difference (d) SNR Difference

Figure 8: Comparison between the new best and old top 3 beam pairs in terms of different metrics.

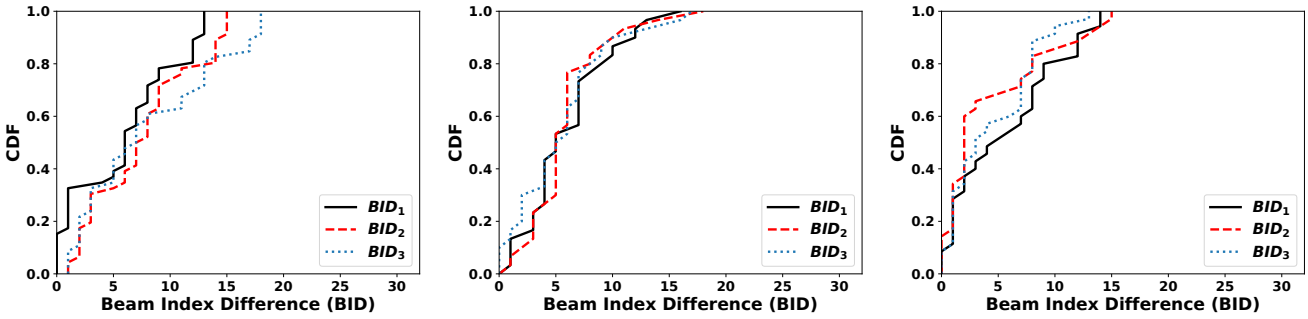

(a) Non-angular displacement. (b) Blockage. (c) Interference.

Figure 9: Beam Index Difference (BID) for different types of link impairment.

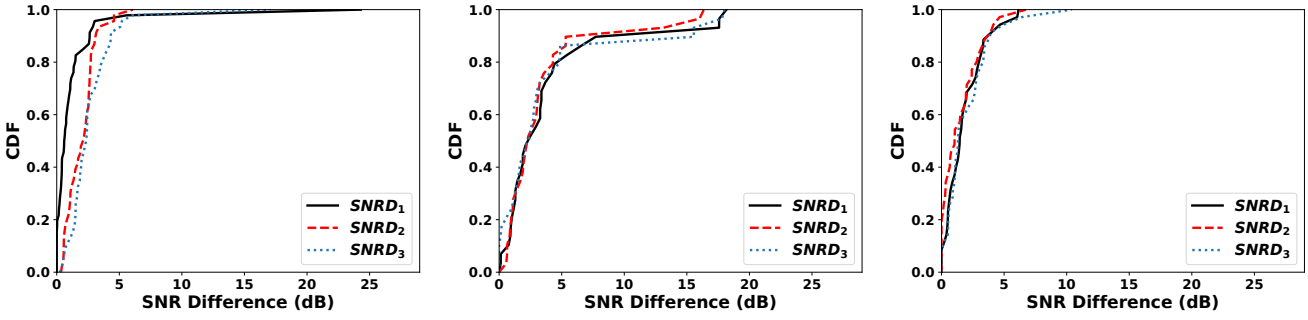

(a) Non-angular displacement. (b) Blockage. (c) Interference.

Figure 10: SNR Difference for different types of link impairment.

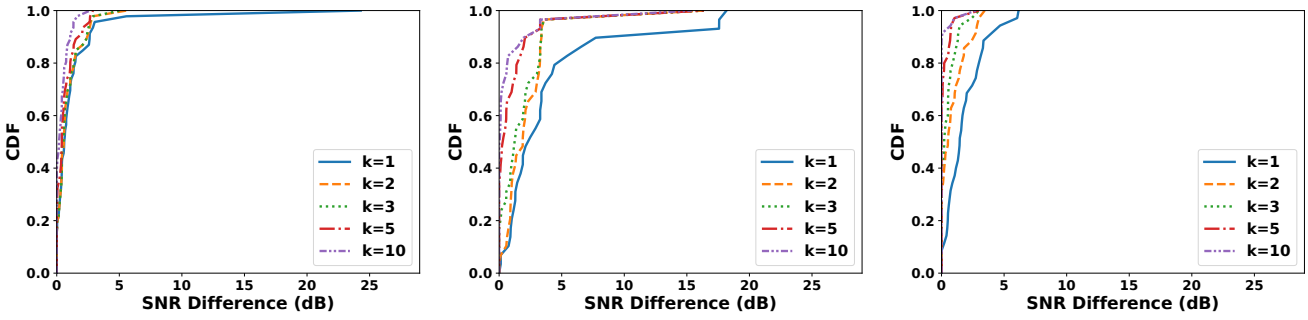

(a) Non-angular displacement. (b) Blockage. (c) Interference.

Figure 11: SNR Difference (k backup beams).

only 1 dB in the median case and 2-3 dB at the 95th percentile. In the case of blockage (Fig. 11b), the loss in the median case is 2-3 dB when the same beam pair is used before and after blockage (k=1), suggesting that, depending on the type of blockage, BA may not be required. However, in 20% of the cases the loss is more than 5 dB which could result in an MCS drop by multiple levels. Interestingly,

maintaining one backup beam (k=2) results in an SNR loss lower than 4 dB in 95% of the cases. Finally, in the case of interference (Fig. 11c), the SNR loss is less than 2 dB in the median case and less than 4 dB in the 90th percentile when k=1, suggesting again that BA is often not required to maintain good performance. Further, with k=2, the SNR loss at the 90th percentile is less than 2 dB suggesting that choosing among the top two beam pairs almost always results in near-optimal performance.

Remark: In all types of link impairment that do not involve angular displacement, maintaining the top 2 beam pairs as backup pairs almost always results in an SNR loss within 2-3 dB, while saving the overhead of performing a full beam search.

7 CONCLUSION

We conducted an extensive experimental study of rate adaptation and beam adaptation, the two primary link adaptation mechanisms in 60 GHz WLANs. We compared the effectiveness of the two mechanisms in a variety of indoor environments and scenarios, including linear and angular displacement, mobility, blockage, and interference. We found that none of the two mechanisms performs consistently well alone and often a combination of the two is required for optimal performance. Next, we studied the effectiveness of two rate adaptation approaches – SNR-based rate adaptation, which has been proposed in recent works, and a learning-based approach using PHY layer information, and showed that the latter is promising, especially when combined with online training. Finally, we explored the effectiveness of maintaining backup beams to speedup link recovery and reduce the beam training overhead. We found that this heuristic is quite effective in most scenarios that do not involve angular displacement.

ACKNOWLEDGMENTS

This work was supported in part by the NSF grants CNS-1553447 and CNS-1801903.

REFERENCES

[1] [n.d.]. Minstrel Rate Adaptation Algorithm Documentation. http://madwifiproject.org/browser/madwifi/trunk/ath_rate/minstrel/minstrel.txt
[2] Qualcomm Inc. Amichai Sanderovich. 2016. Techniques For Rate Selection In Millimeter-Wave Communication Systems. US Patent 9,312,985.
[3] Xueli An, Chin-Sean Sum, R. Venkatesha Prasad, Junyi Wang, Zhou Lan, Jing Wang, Ramin Hekmat, Hiroshi Harada, and Ignas Niemegeers. 2009. Beam Switching Support to Resolve Link-Blockage Problem in 60 GHz WPANs. In *Proc. of IEEE PIMRC*.
[4] Hany Assasa, Swetank Kumar Saha, Adrian Loch, Dimitrios Koutsonikolas, and Joerg Widmer. 2018. Medium Access and Transport Protocol Aspects in Practical 802.11ad Networks. In *Proc. of IEEE WoWMoM*.
[5] John C. Bicket. 2005. *Bit-rate Selection in Wireless Networks*. Master's thesis. MASSACHUSETTS INSTITUTE OF TECHNOLOGY.
[6] Joseph Camp and Edward Knightly. 2010. Modulation Rate Adaptation in Urban and Vehicular Environments: Cross-Layer Implementation and Experimental Evaluation. *IEEE/ACM Transactions on Networking* 18, 6 (2010), 1949 – 1962.
[7] Felix Fietkau. 2010. Minstrel_ht: new rate control module for 802.11n. https://lwn.net/Articles/376765/
[8] Muhammad Kumail Haider and Edward W. Knightly. 2016. Mobility Resilience and Overhead Constrained Adaptation in Directional 60 GHz WLANs: Protocol Design and System Implementation. In *Proc. of ACM MobiHoc*.
[9] Daniel Halperin, Wenjun Hu, Anmol Sheth, and David Wetherall. 2010. Predictable 802.11 packet delivery from wireless channel measurements. In *Proc. of ACM SIGCOMM*.
[10] Daniel Halperin, Srikanth Kandula, Jitendra Padhye, Paramvir Bahl, and David Wetherall. 2011. Augmenting data center networks with multi-gigabit wireless links. In *Proc. of ACM SIGCOMM*.
[11] Haitham Hassanieh, Omid Abari, Michael Rodriguez, Mohammed Abdelghany, Dina Katabi, and Piotr Indyk. 2018. Fast Millimeter Wave Beam Alignment. In *Proc. of ACM SIGCOMM*.
[12] IEEE 802.11 Working Group. 2012. IEEE 802.11ad, Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band. (2012).
[13] National Instruments. 2017. Introduction to the NI mmWave Transceiver System Hardware - National Instruments. http://www.ni.com/white-paper/53095/en/
[14] Tarun Joshi, Disha Ahuja, Damanjit Singh, and Dharma P. Agrawal. 2008. SARA: Stochastic Automata Rate Adaptation for IEEE 802.11 Networks. *IEEE Transactions on IEEE Transactions on Parallel and Distributed Systems* 19, 11 (2008), 1579–1590.
[15] Glenn Judd, Xiaohui Wang, and Peter Steenkiste. 2008. Efficient Channel-aware Rate Adaptation in Dynamic Environments. In *Proc. of ACM MobiSys*.
[16] Ad Kamerman and Leo Monteban. 1997. WaveLAN II: A high-performance wireless LAN for the unlicensed band. *Bell Labs Technical Journal* (1997).
[17] Ernest Kurniawan, Lin Zhiwei, and Sumei Sun. 2017. Machine Learning-Based Channel Classification and Its Application to IEEE 802.11ad Communications. In *Proc. of IEEE GLOBECOM*.
[18] Chi-Yu Li, Syuan-Cheng Chen, Chien-Ting Kuo, and Chui-Hao Chiu. 2020. Practical Machine Learning-based Rate Adaptation Solution for Wi-Fi NICs: IEEE 802.11ac as a Case Study. *IEEE Transactions on Vehicular Technology* (2020).
[19] D. Nguyen and J. J. Garcia-Luna-Aceves. 2011. A practical approach to rate adaptation for multi-antenna systems. In *Proc. of IEEE ICNP*.
[20] Thomas Nitsche, Guillermo Bielsa, Irene Tejado, Adrian Loch, and Joerg Widmer. 2015. Boon and Bane of 60 GHz Networks: Practical Insights into Beamforming, Interference, and Frame Level Operation. In *Proc. of ACM CoNEXT*.
[21] Ioannis Pefkianakis, Yun Hu, Starsky H.Y. Wong, Hao Yang, and Songwu Lu. 2010. MIMO Rate Adaptation in 802.11N Wireless Networks. In *Proc. of ACM MobiCom*.
[22] S. K. Saha, H. Assasa, A. Loch, N. M. Prakash, R. Shyamsunder, S. Aggarwal, D. Steinmetzer, D. Koutsonikolas, J. Widmer, and M. Hollick. 2018. Fast and Infuriating: Performance and Pitfalls of 60 GHz WLANs Based on Consumer-Grade Hardware. In *Proc. of IEEE SECON*.
[23] Swetank Kumar Saha, Yasaman Ghasempour, Muhammad Kumail Haider, Tariq Siddiqui, Paulo De Melo, Neerad Somanchi, Luke Zakrajsek, Arjun Singh, Roshan Shyamsunder, Owen Torres, Daniel Uvaydov, Josep Miquel Jornet, Edward Knightly, Dimitrios Koutsonikolas, Dimitris Pados, Zhi Sun, and Ngwe Thawdar. 2019. X60: A Programmable Testbed for Wideband 60 GHz WLANs with Phased Arrays. *Computer Communications* (2019).
[24] Swetank Kumar Saha, Tariq Siddiqui, Dimitrios Koutsonikolas, Adrian Loch, Joerg Widmer, and Ramalingam Sridhar. 2017. A Detailed Look into Power Consumption of Commodity 60 GHz Devices. In *Proc. of IEEE WoWMoM*.
[25] Souvik Sen, Naveen Santhapuri, Romit Roy Choudhury, and Srihari Nelakuditi. 2010. AccuRate: Constellation Based Rate Estimation in Wireless Networks. In *Proc. of USENIX NSDI*.
[26] Daniel Steinmetzer, Daniel Wegemer, Matthias Schulz, Joerg Widmer, and Matthias Hollick. 2017. Compressive Millimeter-Wave Sector Selection in Off-the-Shelf IEEE 802.11 ad Devices. In *Proc. of the 11th ACM CoNEXT*.
[27] Li Sun, Souvik Sen, and Dimitrios Koutsonikolas. 2014. Bringing Mobility-Awareness to WLANs using PHY layer information. In *Proceedings of the ACM CoNEXT*.
[28] Sanjib Sur, Vignesh Venkateswaran, Xinyu Zhang, and Parmesh Ramanathan. 2015. 60 GHz Indoor Networking Through Flexible Beams: A Link-Level Profiling. In *Proc. of ACM SIGMETRICS*.
[29] Sanjib Sur, Xinyu Zhang, Parmesh Ramanathan, and Ranveer Chandra. 2016. BeamSpy: Enabling Robust 60 GHz Links Under Blockage. In *Proc. of USENIX NSDI*.
[30] Mythili Vutukuru, Hari Balakrishnan, and Kyle Jamieson. 2009. Cross-Layer Wireless Bit Rate Adaptation. In *Proc. of ACM SIGCOMM*.
[31] Chiapin Wang. 2013. Dynamic ARF for Throughput Improvement in 802.11 WLAN via a Machine-Learning Approach. *Elsevier Journal of Network and Computer Applications* 36, 2 (2013), 667–676.
[32] Michael S. Y. Wong, Jeffrey M. Gilbert, and Craig H. Barratt. 2008. Wireless LAN using RSSI and BER parameters for transmission rate adaptation. US Patent 7,369,510.
[33] Starsky H. Y. Wong, Hao Yang, Songwu Lu, and Vaduvur Bharghavan. 2006. Robust Rate Adaptation for 802.11 Wireless Networks. In *Proc. of ACM MobiCom*.
[34] Ding Zhang, Panneer Selvam Santhalingam, Parth Pathak, and Zizhan Zheng. 2019. Characterizing Interference Mitigation Techniques in Dense 60 GHz mmWave WLANs. In *Proc. of ICCCN*.
[35] A. Zhou, X. Zhang, and H. Ma. 2017. Beam-forecast: Facilitating mobile 60 GHz networks via model-driven beam steering. In *Proc. of IEEE INFOCOM*.
[36] Xia Zhou, Zengbin Zhang, Yibo Zhu, Yubo Li, Saipriya Kumar, Amin Vahdat, Ben Y. Zhao, and Haitao Zheng. 2012. Mirror Mirror on the Ceiling: Flexible Wireless Links for Data Centers. In *Proc. of ACM SIGCOMM*.
[37] Yibo Zhu, Xia Zhou, Zengbin Zhang, Lin Zhou, Amin Vahdat, Ben Y. Zhao, and Haitao Zheng. 2014. Cutting the Cord: a Robust Wireless Facilities Network for Data Centers. In *Proc. of ACM MobiCom*.