Mitigating Lateral Interference: Adaptive Beam Switching for Robust Millimeter-Wave Networks

Daniel Steinmetzer
Secure Mobile Networking Lab
TU Darmstadt, Germany
dsteinmetzer@seemoo.de

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

Amanda García-García
IMDEA Networks Institute
Madrid, Spain
amanda.garcia@imdea.org

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

Matthias Hollick
Secure Mobile Networking Lab
TU Darmstadt, Germany
mhollick@seemoo.de

ABSTRACT

Putting into practice "pseudo-wire" links in wireless millimeter-wave (mm-wave) networks is challenging due to the significant side lobes of consumer-grade phased antenna arrays. Nodes should steer their beams such that they maximize the signal gain but also minimize interference from lateral directions via both their main lobe and their side lobes. Most importantly, interference can be caused by parallel operation of incompatible standards such as WiGig and IEEE 802.11ad and may change very fast. This timing requirement, prevents the use of existing beam switching solutions to mitigate interference. In this paper, we present an adaptive beam switching (ABS) mechanism that can deal with the above timescale issue in rapidly changing interference scenarios. Instead of performing a full beam sweep, the key idea is to only probe beampatterns at the receiver which are likely to avoid interference. In contrast to earlier work, our mechanism does not require any location information nor a detailed shape of the beampatterns. We exploit similarities among side lobes of beampatterns to estimate the performance of all beampatterns without sending extensive probes. To evaluate our mechanism in practice, we develop a customized research platform that allows us to control the beam-selection on low-cost IEEE 802.11ad routers. Experimental results with WiGig transceivers as interference source show that our adaptive beam switching mechanism achieves an average throughput gain of 60% and decreases the training time by 82.4% compared to the original IEEE 802.11ad behavior.

1 INTRODUCTION

Using pencil-shaped beams in millimeter-wave (mm-wave) networks to achieve "pseudo-wire" behavior is a myth. The reason is that real-world electronically-steerable phased antenna arrays suffer from significant side lobes [9]. This is particularly notable in the case of consumer-grade hardware due to its cost-efficient design. As a result, interference becomes a problem in mm-wave networks despite the use of directional communication. Faulty nodes or parallel operation of incompatible standards which distort the channel, may cause significant interference not just via main lobes but also via side lobes. That is, the best beampattern for a receiver may not be the one whose main lobe most accurately points towards the transmitter but the one whose side lobes do not capture any interference. Also, due to the sparse multi-path environment, interference is one of the main reasons for link instability in quasi-static scenarios. As such impairments are highly direction depended, mitigating interference in mm-wave networks pose a challenging task.

The straightforward approach to identify the most suitable beampattern in order to address interference is performing a receive beam sweep. However, this exhaustive search is highly inefficient. For instance, the duration of a beam sweep in the IEEE 802.11ad standard is in the order of a few milliseconds [3, 10]. Since interference may be intermittent and change very often, triggering a beam sweep each time such a change occurs would strongly impact performance—nodes would spend a long time sweeping. The key challenge when it comes to avoiding interference is the small timescale at which interference events take place. This is in contrast to other events such as blockage and device movement, which also may require beampattern adaptation but take place at the timescale of *seconds*. At present, mm-wave networks lack efficient mechanisms to address this order-of-magnitude difference regarding the timescale of interference.

In this paper, we present Adaptive Beam Switching (ABS), which is a mechanism that adapts receive beampatterns timely and efficiently to changing interference. Our mechanism does not require any location information and only needs limited information regarding beampattern shape. Specifically, the intuition behind our mechanism is as follows. Instead of performing a full beam sweep whenever interference changes, nodes only probe the beampatterns which are most likely to result in good performance. To this end, nodes keep track of the probability of interference-free transmission for each beampattern. Nodes initialize those probabilities based on beam sweeps at comparatively large, fixed intervals. Interference changes trigger the aforementioned individual beampattern probes. The key feature of our mechanism is that, whenever a node sends such an individual probe, it does not only update the probability for that beampattern, but it also updates the probability of *similar* beampatterns at *zero cost*. Basically, we update the probabilities of other beampatterns based on the correlation of their lobes with the lobes of the probed beampattern.

**Figure 1: Beampattern switch example. The beampattern in case (b) mitigates interference because the interference signal falls into a beampattern minimum.**

Figure 1 shows an example of our mechanism. In Figure 1(a), the receiver RX is receiving data from transmitter TX. To this end, it uses a receive beampattern whose main lobe points towards TX. During the communication, a nearby node starts transmitting. Although the node is not aligned with the main lobe, RX still receives interference via one of its side lobes. Our mechanism at RX detects the interference and switches to a similar beampattern as shown in Figure 1(b). While the main lobe still points roughly towards TX, the interfering signal falls into a minimum of the beampattern, thus mitigating its impact. Moreover, RX also increases the probability of using the beampattern in Figure 1(c) since its side lobes are similar to the one in Figure 1(b). All in all, the design of our mechanism has some significant advantages compared to traditional beam sweeping. With each probe, our mechanism learns about many beampatterns. This reduces the overhead dramatically. Our mechanism does not require full beampattern information but only a beampattern correlation matrix and operates in a fully distributed manner. No coordination or knowledge from other nodes in the network is needed.

We evaluate the performance of our mechanism in practice. To this end, we develop a 60 GHz research platform based on commercial off-the-shelf IEEE 802.11ad routers. This platform allows us to control the beampattern selection process while maintaining full compatibility with IEEE 802.11ad. In contrast to earlier work on 60 GHz networking, our platform provides full bandwidth, real-time operation, and is available at low cost. Specifically, our contributions are as follows:

(1) We design an algorithm that exploits similarities of side lobes in different beams to mitigate interference.
(2) We practically show that side lobes play a critical role regarding interference on commercial devices.
(3) We evaluate our algorithm in practice using our aforementioned 60 GHz research platform.

The remainder of this paper is organized as follows. We summarize related work in Section 2. In Section 3, we present our adaptive beam switching mechanism for interference avoidance, which we evaluate in Section 4. Finally, we conclude our work in Section 5.

2 RELATED WORK

Despite the theoretical assumption of pencil-shaped antenna beampatterns with negligible side lobes, current off-the-shelf devices with moderate number of antenna elements exhibit a significant amount of side lobe power [9]. As a result, these devices are more prone to interference and jamming [13]. To overcome this, smart

beam switching and training concepts [6] are important to preserve or recover the link between communicating nodes.

The simplest method to determine the best transmission beampatterns between transceivers probes all possible combinations and therewith causes high overhead. To avoid this overhead, the IEEE 802.11ad standard proposes a two-phase beam search approach [10]. First, the sector-level sweep determines an initial coarse-grained antenna sector configuration. Then, the beam refinement phase fine-tunes the selected sectors to obtain the directional beam pair with the highest channel quality. However, this approach still leads to high overhead in case of frequent beam switching. Thus, more optimizations are likely to appear as part of the IEEE 802.11ay standard [3]. To further reduce the overhead of beam search, numerical divide and conquer algorithms can efficiently narrow down the search space [7]. Similar approaches with hierarchical structures of beampatterns of different beamwidths allow to iteratively refine the beamtraining accuracy [1, 5]. Compressive path tracking approaches [8, 12] select pseudo-random beams for probing and derive the direction of the most significant signal path. Out-of-band measurement techniques can provide coarse-grained localization to support the beam training [11]. While such methods efficiently find the optimal beam using periodic sweeps, they are insufficient to handle mobility, blockage, and interference.

To achieve mobility resilience, Haider et al. adjust the data-rate and beamwidth during runtime so that low data-rates and wide beams prevent connection losses if necessary. Likewise, Sur et al. [14] study the possibility of switching to a wider sector upon blockage detection for communication link recovery. The approaches in [2, 4] suggest learning mechanisms to detect a suitable non-line-of-sight path when the line-of-sight path is blocked. Also in [15], Sanjib et al. instantaneously predict the link quality to discover an alternative beam when the primary link fails.

Unlike mobility and blockage, interference and malicious jamming is unpredictable. Widening the beam, as suggested to address mobility, potentially causes a greater level of interference. Approaches that switch to non-line-of-sight paths lack tracking of rapid changing channel distortions and, hence, are unsuitable to avoid interference. In contrast, our fully distributed algorithm learns the best beam selection and takes into account intermittent interference from various directions, as described in the following section.

3 ADAPTIVE BEAM SWITCHING (ABS)

In this section, we present an Adaptive Beam Switching (ABS) algorithm which detects and avoids interference. With neither location nor detailed beampattern information, ABS monitors the stability of the link and determines alternative beampatterns with minimal search overhead. In the following, we provide a protocol overview and describe the algorithm behind this approach in detail.

3.1 Protocol Overview

The general idea of our protocol is to continuously monitor a stability metric that represents changes of the received signal quality (i.e., the signal-to-interference-plus-noise ratio (SINR)). With low stability, our protocol probes more often the channel conditions. For probing, it probabilistically selects beampatterns that are similar to
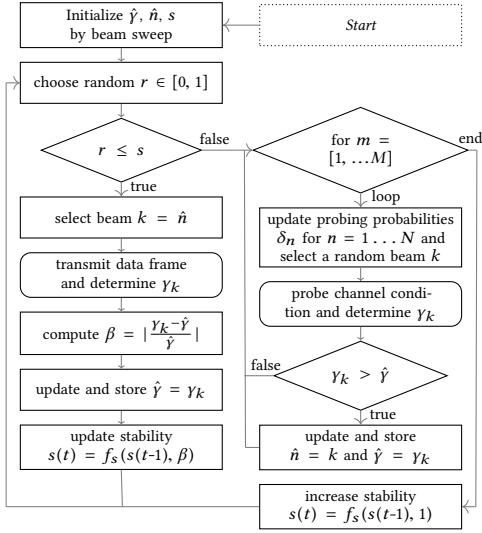
Figure 2: Flow Chart of our ABS algorithm.

the current one but exhibit different gaps in the beampattern shape. This allows to find beampatterns that mitigate interference impairments but still provide a suitable signal gain at each transceiver independently.

3.2 Protocol Specification

In the following, we provide our protocol specification as illustrated in the flow chart in Figure 2.

3.2.1 Initialization. Each transceiver has a set of N predefined receive beampatterns it can choose from to maximize the signal quality for a respective communication partner. Each beampattern $n = 1 \ldots N$ exhibits main lobes and side lobes in different directions. In ABS, we keep track of the beam $\hat{n}$ that provides the highest SINR over time. At time $t = t_0$, we initialize the selected beam $\hat{n}$ from a complete sweep operation as:

$$\hat{n} = \arg\max_{n=1\ldots N} \gamma_n. \tag{1}$$

The corresponding SINR is stored as $\hat{\gamma}$ and initialized by $\hat{\gamma} = \gamma_{\hat{n}}$ at t_0. The stability s is initialized with the maximum value $s = 1$. Note that a complete beam sweep is required only during the initialization phase (i.e., at $t = t_0$).

3.2.2 System Stability. We assume a locally time slotted system, similar to IEEE 802.11ad. At each time slot t, a transceiver either (i) probes the channel or (ii) transmits data with the best-known beampattern from the previous slot $t - 1$. To make this decision, we use a stability parameter $s \in [0, 1]$. Given a random value r from a uniform random distribution with $r \in [0, 1]$, ABS continues transmitting data (ref. Section 3.2.5) if $r \leq s$ or initiates probing (ref. Section 3.2.3) otherwise. As channel conditions change over time, the stability s is continuously updated.

3.2.3 Probing Probabilities. Each beampattern n is assigned a probing probability δ_n according to the current beampattern selection $\hat{n}$ and their similarities in beampattern shape. Let u and v be a pair of beampatterns with $u, v = 1, \ldots, N$ but $u \neq v$. Given that

every beampattern n features a unique antenna radiation pattern $W_n(\Theta)$ for $\Theta \in [-\pi, \pi]$, we determine the correlation between two beams by their cross-correlation coefficient with zero lag

$$c_{u,v} = W_u \star W_v[0] = \int_{-\pi}^{\pi} W_u^*(\Theta) \cdot W_v(\Theta) d\Theta, \tag{2}$$

where $W_u^*(\Theta)$ is the complex conjugate of $W_u(\Theta)$. Next, the algorithm also determines the correlation of the minima in beampattern shapes to find beampatterns with similar gaps as

$$\bar{c}_{u,v} = \frac{1}{1 + W_u} \star \frac{1}{1 + W_v}[0]. \tag{3}$$

Take note, that at runtime ABS does not need detailed information on the beampattern shapes. Only the correlations $c_{u,v}$ and $\bar{c}_{u,v}$ of beam pairs are required. These values can be determined in device callibration only once and stored for later use. With this approach, we also encounter for varying hardware inaccuracies in radio curcuits among different devices.

To overcome interference or signal jamming, we aim to search for an alternative beampattern that *(i)* maximizes the antenna gain towards the intended direction and *(ii)* minimizes the impact of interference. Specifically, an alternative beam should have a high correlation with the currently used beampattern but different zeros to steer the antenna away from the interference direction. Therefore, the probability of probing for beampattern n given the current beampattern $\hat{n}$ is

$$\delta_n = c_{\hat{n},n} \cdot \left(1 - \bar{c}_{\hat{n},n}\right). \tag{4}$$

3.2.4 Channel Probing. Channel probing determines if, under the current channel condition, beam switching is beneficial. It is performed in the form of a burst of M control frames. Each frame probes an individual beampattern. For each $m = 1, \ldots, M$ our algorithm selects a random beampattern k based on the corresponding δ_n for $n = 1, \ldots, N$. For the chosen beampattern, we measure the SINR γ_k and update the internal state. The measured value γ_k is compared to that of the previous used beampattern $\gamma_{\hat{n}}$. If $\gamma_k > \gamma_{\hat{n}}$, the algorithm updates the current beampattern and its corresponding SINR to $\hat{n} = k$ and $\gamma_{\hat{n}} = \gamma_k$, respectively. Doing so, the probed beampattern gets selected when it provides a better quality than the current one.

Lastly, after probing all M selected beampatterns, the stability value is increased. To update the stability, we utilize a update function $s = f_s(s, \beta)$ that updates the current stability based on the previous stability and an update parameter β. We implement this update function as moving average with:

$$f_s(s, \beta) = \alpha \cdot \beta + (1 - \alpha) \cdot s \mid \alpha = 0.1 \tag{5}$$

However, this function and especially the adjustment parameter α is replaceable for general applicability of our protocol. In preliminary experiments, we revealed that this implementation with $\alpha = 0.1$ provides suitable results. More sophisticated update strategies may take into account packet delivery and error rate to adaptively control the stability. We plan to address such strategies in future work. To increase the stability after probing, we set $\beta = 1$ and $s = f_s(s, 1)$.

3.2.5 Data Transmission Procedure. If the stability is high so that $r \leq s$, the algorithm continues transmitting data with the same beampattern used before. It selects the beampattern as $k = \hat{n}$

and transmits a data frame. Still, it updates the stability as the signal quality might change. The relative change in the channel conditions affects the stability. Specifically, the stability update is determined by $\beta = |\frac{\gamma_k - \hat{\gamma}}{\hat{\gamma}}|$, and thus updates to $s = f_s(s, \beta)$. We use the same update mechanism as for probing in Section 3.2.3. The instantaneous SINR is stored as $\gamma_{\hat{n}} = \gamma_k$.

Whenever the SINR in the current data transmission drops, the decreasing stability increases the probability of channel probing in the upcoming time slots. Still, the moving average of stability monitoring ensures that our algorithm does not overreact on rare outliers and temporary outage.

4 PERFORMANCE EVALUATION

In this section, we evaluate our ABS algorithm in practice. To this end, we extend the operation of a commercial off-the-shelf 60 GHz router to control the applied beampatterns. As a result, we obtain a first-of-its-kind IEEE 802.11ad full-bandwidth 60 GHz beamforming testbed that operates in real time. This allows us to obtain unprecedented insights compared to earlier work on 60 GHz networking. We first describe our testbed setup and, after that, analyze the performance of our interference mitigation approach.

4.1 Testbed Setup

We consider the case of interference due to parallel operation of incompatible standards. To this end, we use two TP-Link Talon AD7200 routers that use the IEEE 802.11ad protocol to communicate, and a WiHD transmitter-receiver pair that is based on WiGig. The latter does not perform clear channel assessment, and thus causes interference on the former. To control the operation of the Talon routers, we port the LEDE Project[1] to support their specific architecture. This major effort enables us to install LEDE on them, and access the firmware components that define the beampattern shape. We can thus set any arbitrary beampattern on the phased antenna array. The array consists of 32 antenna elements which are individually controllable in phase and magnitude. No earlier work in this area has achieved such control on commercial off-the-shelf 60 GHz devices. In this paper, we use the 34 pre-existing beampatterns which are defined in the interface firmware to assess the performance of our mechanism under real-world conditions. Figure 3 exemplary illustrates the shape of two of these beampatterns and how they allow for interference mitigation. All beampatterns feature different beam-widths and main-lobe directions and complement each other to provide strong gains in all directions.

Figure 4 shows our testbed setup. We place one of the Talon routers in the center of a semicircle and configure it as an IEEE 802.11ad station (STA). The second router is located on the semicircle and configured as an access point (AP). We consider two locations of the access point on the semicircle, at $+45°$ and $-45°$. Additionally, we place the WiHD transmitter on 18 evenly distributed locations on the semicircle, and the WiHD receiver in the center of the semicircle close to the STA. As a result, the STA receives both the intended transmission of the AP and the interference transmission of the WiHD transmitter from a range of different angles. For our experiment campaign, we set the beampattern of


Figure 3: Exemplary antenna patterns of Sector 22 and 24 measured in terms of SNR for a Talon AD7200 router in an anechoic environment. While both patterns exhibit similar gains in the direction of $60°$, their lobes completely differ at $0°$ and $-30°$. As a results switching between both patterns might overcome interference from that particular directions.


Figure 4: Practical experiment setup with WiHD (a) and Talon router (b). The solid line shows the location of the router at $+45°$ and the dashed line at $-45°$.

the AP to point to the center of the semicircle. This ensures that signal quality fluctuations are not due to switching beampatterns. At the receiver, we measure throughput and SNR for each possible receive beampattern and each location of the interfering WiHD transmitter on the semicircle. We perform all of our experiments in a large and empty sports hall. In each experiment, we generate TCP traffic from the AP to the STA using iperf[2].

4.2 Practical Results

In the following, we provide our practical evaluation results in terms of interference avoidance, probing time, and protocol operation.

4.2.1 Mitigating Lateral Interference. In our first experiment, we analyze whether the intuition sketched in Figure 1 holds. That is, we study whether switching to a receive beampattern that is not the one that provides the highest gain towards a transmitter helps mitigating lateral interference. The heat map in Figure 5 depicts our result. The map shows the throughput that we achieve for each possible receive beampattern and each location of the interfering WiHD node. The row marked with an arrow indicates

[1]Linux Embedded Development Environment (c.f. https://lede-project.org/)

[2]Iperf - The TCP/UDP Bandwidth Measurement Tool (c.f. https://iperf.fr/)

the best pattern when the interfering node is off, and the dashed line indicates the best pattern when the interference is switched on. The former and the latter do not match. Moreover, the latter changes for each location of the interfering node. This validates that (a) deviating from the best pattern helps in case of interference, and (b) choosing an alternative pattern is not straightforward. Since practical beampatterns are highly irregular [9], the best alternative pattern may be any of the available beampatterns, and not just the neighboring sectors. Thus, an efficient probing scheme such as ABS is needed.

Figure 6 shows the throughput gain that we achieve when selecting the sector with less interference. We compute the average over all possible locations of the interfering node. Further, we compute the gain compared to the best sector highlighted in Figure 5 and to the default behavior of the router. The latter refers to the case when we do not force the router to use a specific receive sector. Compared to the best sector, we achieve about 60% average throughput gain, but for individual locations of the interfering node we achieve gains in the range of 2× to 8×. For +45°, we observe that the gain compared to the default behavior is lower than for −45°. The reason is that the beampatterns are not symmetric. Thus, avoiding the WiHD interference is more challenging for certain alignments of the IEEE 802.11ad link than for others.

4.2.2 Probing Time. Next, we investigate the impact of the number of probes M introduced in Section 3.2.4. If M is small, ABS only probes a few sectors each time that the stability degrades. As a result, ABS requires more time to select a stable pattern. However, beyond a certain threshold of M, the probability of probing a suitable sector increases significantly, and thus ABS quickly converges. Figure 7 depicts this effect for a challenging scenario, where the interfering node is located close to the AP. In this case, the STA should set $M \geq 5$ to achieve a short stabilization time. However, if the value of M is too large, performance in terms of throughput drops since beam sweeps become inefficient. From our experiments, we observe that $M = 6$ is typically a good trade-off. In comparison to the default operation of the router, this means that the beam sweep time decreases by 82.4%. Moreover, ABS requires less beam sweeps due to its high stability. Thus, the time spent on beam training is much shorter.

4.2.3 Protocol Operation. In our last experiment, we study the protocol operation of ABS when reacting to interference. We consider intermittent interference from two WiHD transmitters located at different angles on the semicircle. When one of the interfering nodes starts transmitting, the SNR drops at the STA. As a result, the default mechanisms in IEEE 802.11ad trigger a beam sweep to find a better receive beam. However, this new receive beam is unlikely to also minimize the interference from the second WiHD node or a mobile interferer. This causes IEEE 802.11ad to trigger beam sweeps continuously, which strongly degrades performance. In contrast, the system stability s in ABS prevents such fluctuations. To validate this, we emulate the behavior of ABS and IEEE 802.11ad on the real-world traces that we collect from our practical testbed.

Figure 8 depicts two examples of the above scenario. In example (a), we place the interfering nodes at a reasonable angular distance from the AP. As expected, IEEE 802.11ad continuously triggers
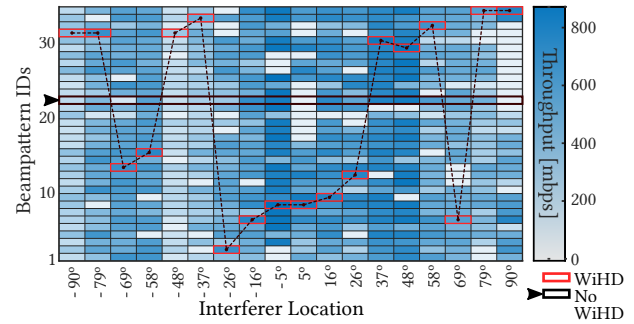

Figure 5: TCP throughput for all receive beampatterns at the STA and all locations of the interfering WiHD transmitter. The AP is located at −45°. The arrows indicate the best beampattern in case of no interference. The markings indicate the highest throughput.


Figure 6: Average throughput gain for all locations of the interfering node. The default behavior refers to the case when we do not force any specific sector.


Figure 7: Stabilization time of ABS.

beam sweeps, and switches among two different patterns. In contrast, ABS quickly finds a suitable sector and stabilizes within less than a millisecond. In Figure 8 we only show the selected beampatterns during the first 16 milliseconds of the experiment for clarity. However, the fluctuations of IEEE 802.11ad continue until the end of the experiment. The SNR graph of example (a) shows that ABS achieves a stable value whereas 802.11ad experiences frequent SNR drops. As a result, ABS improves throughput by 72%. In example (b), we place the interfering nodes close to the AP. This is a particularly challenging scenario since the phased antenna array of the routers is not capable of producing beampatterns which are narrow enough to filter the interference spatially. Figure 8 shows that this increases the stabilization time of ABS to about six milliseconds. The SNR graph shows that ABS chooses a trade-off beampattern that balances the impact of both interfering nodes. In contrast, IEEE 802.11ad achieves a higher SNR at the expense of costly beam sweeps and high instability. While there is no throughput gain,
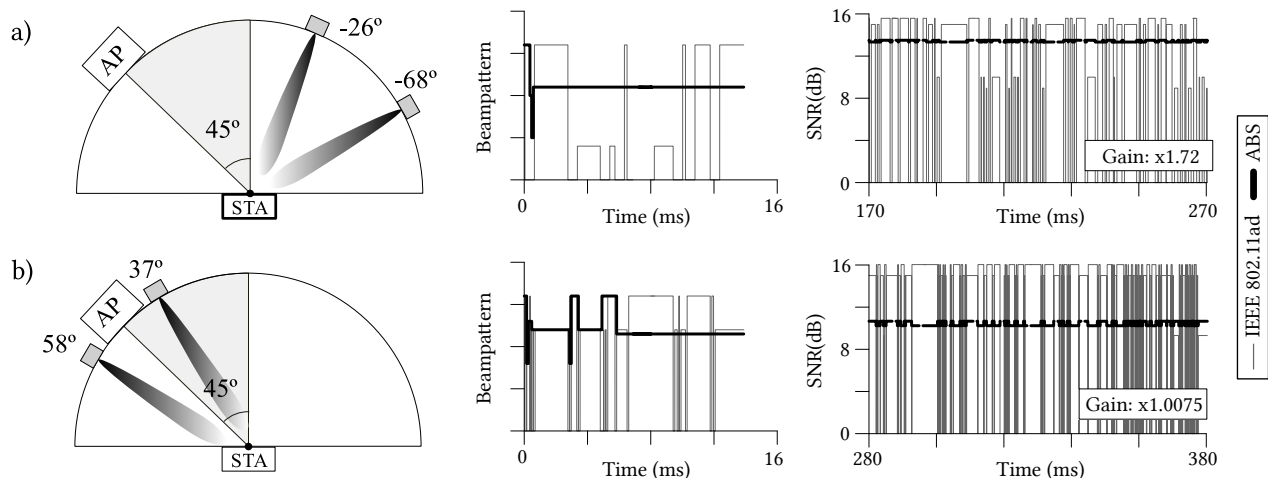
Figure 8: Protocol operation with two interfering nodes for (a) regular case and (b) challenging case.

ABS achieves a very high connection stability, which is crucial for upper-layer protocols such as TCP.

5 CONCLUSION

We present an Adaptive Beam Switching (ABS) mechanism for interference mitigation in 60 GHz wireless networks. With ABS, we significantly reduce the beam steering overhead in case of intermittent interference from neighboring links. Such interference is critical in practical deployments due to the strong side lobes of phased antenna arrays in consumer-grade 60 GHz hardware. Whenever interference occurs, our mechanism switches to an alternative beam at the receiver such that the interference falls into a minimum of the beampattern. The key difference to existing approaches is that ABS does not need to probe all possible beampatterns to mitigate interference but only the ones which are most likely to perform better than the current beampattern. We evaluate ABS in practical experiments with off-the-shelf 60 GHz devices and achieve average throughput gains of 60% and decrease the training time by 82.4% compared to conventional beam sweeping as in IEEE 802.11ad.

REFERENCES

[1] Ahmed Alkhateeb, Omar El Ayach, Geert Leus, and Robert W. Heath. 2014. Channel estimation and hybrid precoding for millimeter wave cellular systems. *IEEE Journal of Selected Topics in Signal Processing* 8, 5 (2014), 831–846.

[2] 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 *IEEE 20th International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*.

[3] Alecsander Eitan and Carlos Cordeiro. 2016. Short SSW Format for 11ay. (March 2016). IEEE 802.11-16/0416-01-00.

[4] Zulkuf Genc, Umar H. Rizvi, Ertan Onur, and Ignas Niemegeers. 2010. Robust 60 GHz Indoor Connectivity: Is It Possible with Reflections?. In *IEEE 71st Vehicular Technology Conference (VTC)*.

[5] Sooyoung Hur, Taejoon Kim, David J Love, James V Krogmeier, Timothy A Thomas, and Amitava Ghosh. 2013. Millimeter Wave Beamforming for Wireless Backhaul and Access in Small Cell Networks. *IEEE Transactions on Communications* 61, 10 (2013), 4391 – 4403.

[6] Shajahan Kutty and Debarati Sen. 2016. Beamforming for Millimeter Wave Communications: An Inclusive Survey. *IEEE Communications Surveys & Tutorials* 18, 2 (2016), 949–973.

[7] Bin Li, Zheng Zhou, Weixia Zou, Xuebin Sun, and Guanglong Du. 2013. On the Efficient Beam-Forming Training for 60GHz Wireless Personal Area Networks. *IEEE Transactions on Wireless Communications* 12, 2 (Feb 2013), 504–515.

[8] Zhinus Marzi, Dinesh Ramasamy, and Upamanyu Madhow. 2016. Compressive channel estimation and tracking for large arrays in mm-Wave picocells. *IEEE Journal of Selected Topics in Signal Processing* 10, 3 (April 2016), 514 – 527.

[9] 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 *ACM 11th Conference on Emerging Networking Experiments and Technologie (CoNext)*. ACM.

[10] Thomas Nitsche, Carlos Cordeiro, Adriana B Flores, Edward W Knightly, Eldad Perahia, and Joerg Widmer. 2014. IEEE 802.11ad: directional 60 GHz communication for multi-Gigabit-per-second Wi-Fi. *IEEE Communications Magazine* 52, 12 (Dec. 2014), 132–141.

[11] Thomas Nitsche, Adriana B. Flores, Edward W. Knightly, and Joerg Widmer. 2015. Steering with eyes closed: mm-wave beam steering without in-band measurement. In *IEEE Conference on Computer Communications (INFOCOM)*. IEEE, 2416–2424.

[12] Maryam Eslami Rasekh, Zhinus Marzi, Yanzi Zhu, Upamanyu Madhow, and Haitao Zheng. 2017. Noncoherent mmWave Path Tracking. In *ACM 18th International Workshop on Mobile Computing Systems and Applications (HotMobile)*. ACM, New York, USA, 13–18.

[13] Daniel Steinmetzer, Jiska Classen, and Matthias Hollick. 2016. Exploring Millimeter-Wave Network Scenarios with Ray-tracing based Simulations in mmTrace. In *IEEE INFOCOM 2016 Poster Presentation*.

[14] Sanjib Sur, Vignesh Venkateswaran, Xinyu Zhang, and Parmesh Ramanathan. 2015. 60 GHz Indoor Networking Through Flexible Beams: A Link-Level Profiling. *ACM International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)* (Jun 2015).

[15] Sanjib Sur, Xinyu Zhang, Parmesh Ramanathan, and Ranveer Chandra. 2016. BeamSpy: Enabling Robust 60 GHz Links Under Blockage. *13th USENIX Symposium on Networked Systems Design and Implementation (NSDI)* (2016).