Performance and Pitfalls of 60 GHz WLANs Based on Consumer-Grade Hardware

Swetank Kumar Saha, *Student Member, IEEE,* Shivang Aggarwal, *Student Member, IEEE,* Hany Assasa, *Student Member, IEEE,* Adrian Loch, Naveen Muralidhar Prakash, Roshan Shyamsunder, Daniel Steinmetzer, Dimitrios Koutsonikolas, *Senior Member, IEEE,* Joerg Widmer, *Senior Member, IEEE,* and Matthias Hollick, *Member, IEEE*

Abstract—Wireless networks operating in the 60 GHz band have the potential to provide very high throughput but face a number of challenges (e.g., high attenuation, beam training, and coping with mobility) which are widely accepted but often not well understood in practice. Understanding these challenges, and especially their actual impact on consumer-grade hardware is fundamental to fully exploit the high physical layer rates in the 60 GHz band. To this end, we perform an extensive measurement campaign using two commercial off-the-shelf 60 GHz routers in real-world environments. Our results allow us to revisit a range of issues and provide much deeper insights into the *reasons* for specific performance compared to prior work on performance characterization. Further, our study goes beyond basic link characterization and explores for the first time practical considerations such as coverage and access point deployment. While some of our observations are expected, we also obtain highly surprising insights that challenge the prevailing wisdom in the community. We derive the shortcomings of current commercial 60 GHz devices, and the fundamental problems that remain open on the way to fast and efficient 60 GHz networking.

Index Terms—60 GHz, 802.11ad, performance evaluation, COTS devices.

1 Introduction

The almost 14 GHz of unlicensed spectrum [1] centered around 60 GHz [2] have attracted ample attention from both academia and industry as a solution for providing multi-gigabit indoor WLAN connectivity. The first commodity devices operating in this band [3] were based on the WiGig standard and were introduced in the market at the end of 2013, targeting applications like wireless docking stations and wireless HDMI. Over the last few years, devices compliant with IEEE 802.11ad such as Access Points (APs) and laptops have been released commercially, and major chipset manufacturers [4], [5] provide the corresponding tri-band chipsets that support 2.4, 5, and 60 GHz. In the future, 802.11ad (and eventually 802.11ay) devices will likely become as ubiquitous as legacy WiFi.

While 60 GHz networks provide multi-gigabit rates at the physical layer, inefficient network operation can offset much of the nominal performance. The propagation characteristics at 60 GHz have led to a widely accepted set of assumptions: (i) directional communication is needed to overcome the high attenuation in the 60 GHz band, (ii) the overhead due to frequent beam training in case of blockage and mobility is prohibitively high, and (iii) at least one AP per room is required in indoor environments to provide line-of-sight (LOS) links in most cases. Researchers have devoted a significant effort to improve performance and relax some of these assumptions based on insights from experimental software-defined radio testbeds [6], [7] and non-standard-compliant commercial hardware [3]. However, the performance of consumer-grade 802.11ad devices is not well understood. This is critical, since as a community we must first understand the actual practical issues in order to focus our research efforts accordingly.

In this paper, we study 60 GHz performance using commercial off-the-shelf (COTS) devices (APs and laptops). While earlier work also considers consumer-grade hardware [8], [9], the key difference is that our devices are fully compliant with the 802.11ad standard and have been designed for the particular application of a WLAN, instead of a wireless docking scenario. We perform an extensive measurement campaign in real-world scenarios to assess the pitfalls of 802.11ad. While some of our results match the widely accepted issues of millimeter-wave networking, we obtain a number of insights that challenge the prevailing wisdom in the community: key insights are as follows:

- Certain environments such as narrow corridors facilitate propagation, enabling ranges *beyond 160 feet*.
- High rates and long ranges are feasible even if one side of the link uses highly directional beam patterns and the other side uses a quasi-omnidirectional pattern.
- Mounting 802.11ad APs on the ceiling provides the highest range and resiliency against blockage, but occasionally results in erroneous beam pattern selection.

- *Swetank Kumar Saha, Naveen Muralidhar Prakash, Roshan Shyamsunder, Shivang Aggarwal, and Dimitrios Koutsonikolas are with the CSE Department, University at Buffalo, The State University of New York. E-mail: {swetankk,shivanga,naveenmu,roshansh,dimitrio}@buffalo.edu.*
- *Hany Assasa is with IMDEA Networks Institute and Universidad Carlos III de Madrid, Spain. E-mail: hany.assasa@imdea.org.*
- *Adrian Loch and Joerg Widmer are with IMDEA Networks Institute, Madrid, Spain. E-mail: {adrian.loch,joerg.widmer}@imdea.org.*
- *Daniel Steinmetzer and Matthias Hollick are with Technische Universität Darmstadt, Germany. E-mail: {dsteinmetzer,mhollick}@seemoo.tu-darmstadt.de.*

A preliminary version of this work has appeared in the IEEE International Conference on Sensing, Communication and Networking (SECON '18).

- A single AP can serve multiple rooms in a typical home or office environment due to the low attenuation of drywalls, thus simplifying network deployment.
- The actual challenge for network deployment is self-shadowing and antenna placement on the device, which limits communication for steering angles beyond $60°$.
- Transient human blockage has low impact on 802.11ad since the cost of beam training is comparatively low.
- Node mobility is very harmful since the interaction of beam training and rate control is still unsolved.

This paper is structured as follows. In § 2 we provide background on 802.11ad and in § 3 we describe our experimental methodology. § 4, 5, 6, 7, and 8 contain our results on link performance, coverage, AP deployment, blockage, and mobility, respectively. We discuss our insights in § 9. § 10 surveys related work and § 11 concludes the paper.

2 IEEE 802.11AD PRIMER

802.11ad works similarly to earlier versions of the standard except for the additional mechanisms needed for directional communication. The channel is divided into Beacon Intervals (BIs). On our COTS hardware, the duration of a BI is 100 ms. Each BI is in turn divided into a Beacon Header Interval (BHI) and a Data Transmission Interval (DTI). The former is for control messages and the latter for data transmission. The BHI consists of three parts. During the Beacon Transmission Interval (BTI), the AP sends beacon frames on each of its sectors to announce its presence. Next, in the Association Beamform Training (A-BFT), stations and APs train their transmit sectors for the data communication in the DTI using a Sector Level Sweep (SLS). Finally, the Announcement Transmission Interval (ATI) allows stations and APs to exchange other management frames.

Our hardware does not implement the A-BFT but instead moves the SLS to the DTI. This is allowed by the standard. During the SLS, each side transmits a Sector Sweep (SSW) control message on all its sectors while the other side listens omnidirectionally. Each SSW frame is transmitted using MCS 0 to ensure the correct reception of this frame during low SNR operation. The SSW frame has a size of 26 bytes which results in a duration of 14.9 μs (including preamble and PHY header). The inter-frame spacing between two consecutive SSW frames is 1 μs which allows the device to steer its phased antenna array into a specific direction in space before frame transmission. After both sides have completed the sweep, they exchange SSW-FBCK and SSW-ACK messages to inform the other side about which sector they received best. The aforementioned frames have similar size and are also transmitted using MCS 0 which results in a duration of 18.2 μs. The inter-frame space between these two frames is 9 μs. Each TALON router transmits 36 SSWs frames during its sector sweep phase which results in a total beamforming training duration of almost 1.2 ms.

The standard does not define the periodicity of the SLS. On our hardware, nodes perform an SLS every 10 BIs if they are associated but no data are being transmitted. If data are transmitted, an SLS only occurs in case of a missing ACK. During the DTI, nodes can exchange data in a contention-based manner or in a time-division manner. Current hardware only implements the former. Also, nodes use Transmission Opportunities (TxOP) with block acknowledgments for more efficient medium usage.

3 DEVICE AND MEASUREMENT DETAILS

In this section, we first give an overview of the devices that we use and then present our experimental setup.

Devices: The TP-Link Talon AD7200 [10] and the Netgear Nighthawk X10 Smart WiFi Router [11] are the only two 802.11ad-compliant routers available on the market.

The TP-Link Talon AD7200 was the first commercially available 802.11ad router released in June 2016. It uses the 802.11ad QCA9008-SBD1 module with the QCA9500 chipset from Qualcomm, supporting single-carrier data rates up to 4.6 Gbps. The 32-element phased antenna array is located on a separate board and connected to the chipset with a MHF4 cable. The router also includes an 802.11n/ac solution from Qualcomm. Since the router only provides 1G Ethernet, maximum throughput is limited to 1 Gbps.

The Netgear Nighthawk X10 Smart WiFi Router was released around October 2016. It uses the same 802.11ad module from Qualcomm as the one used by Talon. In addition to the 1G Ethernet ports, it has a 10-Gigabit LAN SFP+ interface. Yet, we found that in practice the maximum throughput (with MCS 12) is limited to around 2.3 Gbps.

The Acer Travelmate P446-M [12] laptop, released in April 2016, has the client-version QCA9008-TBD1 of the module used in the Nighthawk and Talon routers, which includes 802.11ac, 802.11ad and Bluetooth chipsets. The host connects to the module using an M.2 slot, runs Linux OS (Fedora 24, kernel 4.x) and uses the open source wil6210 wireless driver to interface with the chipset. It comes with the same 32-element phased antenna array as the routers.

The antenna array in the Talon is placed inside one of the eight external antenna enclosures perpendicular to the router, with the front of the antenna facing away from the router (see the rectangular mark in Fig. 1(a)). This antenna placement is primarily suitable for table-top deployments. In contrast, the antenna array in the Nighthawk is rotated back at an angle of around 45^o compared to the plane of the router (Fig. 1(b)), allowing for the router to be either mounted on a wall or placed on a table. In the laptop, the antenna array is placed on the upper right corner of the laptop's LCD lid (Fig. 1(c)), facing away from the screen. In all the experiments, we keep the lid at a 90^o angle from the surface the laptop is placed on.

The laptop and the Talon router run LEDE [13], which allows us to set the chipset in monitor mode, to log MAC layer information, and to disable beamforming and fix a TX sector. The wil6210 firmware version in these two devices is the publicly available 4.1.0.55. On the other hand, the Nighthawk runs the default OEM router firmware (we were not able to install LEDE) which comes with a much older wil6210 firmware version (2.2.0.10). The different firmware version may contribute to some of the performance differences we observe with the two routers.

Measurement Methodology: For most experiments, the setup consists of one of the two routers running in AP mode attached to a high-end desktop over a 1G/10G wired link and the laptop running in client mode and associated wirelessly to the router over 802.11ad. The desktop generates

(a) Talon AD7200 (b) Nighthawk X10 (c) Acer Laptop

Fig. 1. COTS Devices and Antenna Placement. The rectangular mark shows the antenna array placement in each device and the arrow shows the direction of 0^o Orientation (see Section 4.2).

(a) Lobby. (b) Corridor.

Fig. 2. Experimental Setup.

downlink TCP traffic using iperf3 destined for the laptop. We use an additional Talon router to sniff 802.11ad control and data frames by setting the chipset to monitor mode.

The 60 GHz radios use their own rate adaptation algorithms and beamforming mechanisms. In case the link breaks due to blockage, a full sector sweep is initiated after a missing ACK, as described in Section 2, looking for an alternative path to re-establish the connection. All beams are tested one-by-one and the beam that yields the strongest SNR is selected. If the LOS path is no longer available due to blockage, the new beam (if found) will most likely result in communication via an NLOS path through a reflection. If there are no reflectors in the vicinity to enable an alternative path, the devices keeps performing sector sweeps until the blockage is removed and a path is found.

On the laptop, the wil6210 driver exports detailed connection parameters, including transmitter and receiver (Tx and Rx) MCS, MAC layer throughput, signal quality indicator (SQI), beamforming (BF) status (OK/Failed/Retrying), and sectors in use both by itself and the AP (0-63 are valid sector IDs [14], while 255 refers to cases where no valid sector was found due to low signal strength). We log all the parameters every 150 ms.

We conduct most of our experiments with both routers in a *Lobby* and a *Corridor*, with the router at a height of 6 ft and the laptop at a height of 1.5 ft, as shown in Fig. 2. The former is an open space thinly populated by some desks and chairs. The ceiling is high and thus does not serve as a viable reflector. The latter is a narrow corridor (5ft wide) with drywall on both sides. It does not contain furniture or any other objects. For comparison, we also perform measurements in an open outdoor space.

The main metric we use in this paper is the median PHY data rate calculated from the MCS logs collected from the driver, with the error bars showing the 10th/90th percentile. For the Talon, this metric better represents link performance since its 1G Ethernet interface limits overall throughput to 1 Gbps. On the Nighthawk, we verified that the TCP throughput (reported by iperf3) closely follows the trend of the PHY data rate. An example is shown in Fig. 3(a).

4 Link Performance

In this section, we explore the impact of Tx-Rx distance and relative orientation in Line-of-Sight (LOS) scenarios.

4.1 Distance

It has been a common belief that the communication range at 60 GHz is very short even in free space due to the high attenuation. As a result, commercial and proposed use of 60 GHz technology has been limited until recently to short ranges, e.g., for wireless HDMI [15], wireless docking [3], or for augmenting data center networks with high capacity wireless links [16]–[18]. Our experiments show that this is *not* the case. We first check the range for the both the routers in an outdoor setting which is devoid of multipath/reflections to get a baseline. Table 1 shows that average data rates higher than 1 Gbps (2 Gbps) can be achieved for distances of up to 85 ft (67 ft) for Nighthawk and of up to 69 ft (55 ft) in case of Talon. For longer distances, we found that the data rate drops sharply to zero.

In the indoor scenario, we observe even more encouraging performance where both routers achieve excellent range in both environments. Table 1 shows that in the lobby both routers maintain Gbps rates up to 80 ft due to reflections from walls and nearby objects that help extend the range. In the corridor, the range is even longer; rates above 2 Gbps are achieved for distances as long as 140-155 ft, possibly due to the corridor's waveguide effect [19]. Note that in both environments our maximum range is limited by the dimensions of the building and not by path loss. Unlike in the outdoor experiments, we observe no sharp drop of the throughput in the two indoor environments, which suggests that ranges can be even longer.

Fig. 3(b), 3(c) plot the median PHY data rate over distance with the two routers in the lobby and corridor, respectively. First, we observe that MCS 8 is the highest MCS selected by both routers in both environments even for SQI 100 in these experiments. We note that the current hardware does not implement MCS 5 and 9, which use a 13/16 code rate [14]. Although we did observe MCS 10-12 being used for very short intervals in other experiments (see Fig. 16), it appears that our hardware cannot operate stably at 16-QAM (the modulation used for MCS higher than 9 [14]) and it cannot sustain MCS 10-12 for an extended period of time.

Second, we observe that the performance exhibits very different trends in the two environments. The performance in the lobby exhibits a relatively smooth drop with distance in the case of Nighthawk, but remains relatively stable for the Talon. In contrast, performance differs substantially for both routers in the corridor. Although the data rate for both routers does not drop with distance until about 140-150 ft, it exhibits large variations for distances longer than 50-75 ft due to multipath effects. Interestingly, Fig. 3(b), 3(c) show that performance with Talon remains stable for longer distances than with Nighthawk in both environments. This shows the crucial role of antenna placement within the AP (the primary difference between the two routers).

TABLE 1
Range in different environments.

	Outdoor		Lobby		Corridor	
	>2 Gbps	>1 Gbps	>2 Gbps	>1 Gbps	>2 Gbps	>1 Gbps
Nighthawk	67 ft	85 ft	55 ft	80 ft	140 ft	155 ft
Talon	55 ft	69 ft	80 ft	80 ft	155 ft	155 ft

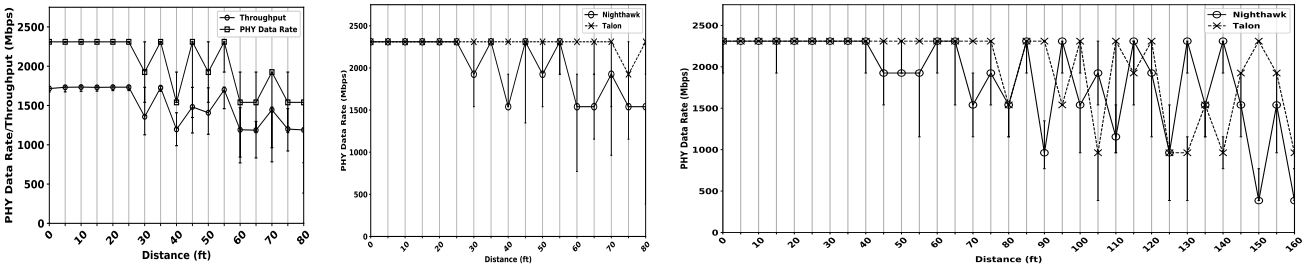

(a) Median PHY data rate and TCP throughput follow the same trend.

(b) Median PHY data rate over distance in the Lobby.

(c) Median PHY data rate over distance in the Corridor.

Fig. 3. Performance characterization with distance.

We conclude that multipath propagation and waveguide effects can boost range to levels commensurate to those achieved by legacy WiFi devices. Thus, dense deployment of 60 GHz APs may not always be required in indoor environments. On the other hand, such long ranges may result in more interference and lower spatial reuse than commonly expected for 60 GHz.

Link adaptation. We now study the two primary link adaptation mechanisms in 802.11ad: beamforming and rate adaptation. Fig. 4(a) and 4(b) show the number of Tx MCS and Tx sector transitions for each router. The Rx sector ID remained the same in all the experiments (c.f. Section 4.2). One would expect the same Tx sector to be used for all distances since the AP and client face each other. Starting from 15 ft, this is mostly the case for Nighthawk in the corridor and for Talon in the lobby. Both use sector 20, which is the default sector when two devices face one another. For shorter distances, we found that different sectors are used due to the specific geometry imposed by the placement of the client almost underneath the AP. However, in the lobby with Nighthawk and in the corridor with Talon we observe that multiple sectors are used. Such changes of the Tx sector result in throughput degradation (e.g., at 30 ft, 40 ft, 60 ft in Fig. 3(b), 4(a) for Nighthawk and 80 ft, 95 ft, 105 ft, 125 ft in Fig. 3(c), 4(b) for Talon).

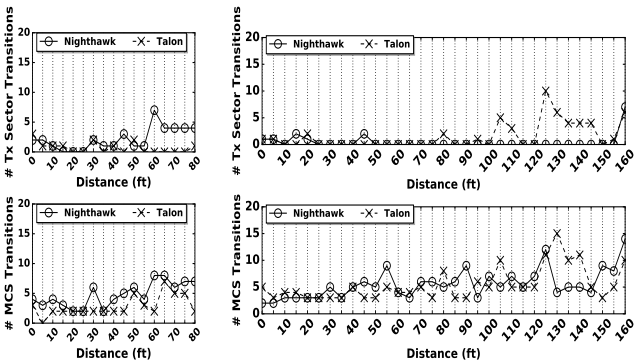

(a) Lobby.

(b) Corridor.

Fig. 4. Rate adaptation vs. beamforming.

Similarly, Fig. 4(a), 4(b) show at least 2-3 and sometimes more than 10 MCS transitions over a 10-second period, even at very short distances. Although why are not able to explain the exact cause of these transitions without access to the rate adaptation logic, we point out that we observed SQI changes at every distance. Such SQI variations may cause MCS transitions if SQI is used in the rate adaptation logic. Random packet loss, which may happen even for high signal strength and the rate adaptation algorithm probing higher rates might also result in MCS transitions.

We conclude that antenna placement and environment strongly affect the performance of link adaptation mechanisms. The beamforming mechanism performs better on Nighthawk in the corridor and on Talon in the lobby. Rate adaptation is similar on both routers in the corridor but more stable for Talon in the lobby. The number of MCS transitions is always larger than the number of Tx transitions at a given distance, suggesting that rate adaptation is triggered first, and beamforming follows when the former fails to find a suitable MCS.

Previous works 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 rate adaptation algorithms [6], [16]–[18]. Since our hardware does not provide access to SNR, in Fig. 5, here we use SQI as a proxy for SNR and investigate the relationship between SQI and MCS in the two environments. Fig. 5(a) and 5(b) plot the median MCS as well as the 10th and 90th percentile for each observed SQI value for both routers in the lobby and corridor, respectively. Fig. 5(c)-5(f) present the same information in a different way, plotting the observed SQI values (in the form of boxplots) for each used MCS. Our results in Fig. 5(a) and 5(b) show that the correlation between MCS and SQI is weak for low and medium link quality. Additionally, Fig. 5(c)-5(f) show that, for both routers in both environments, most MCS indexes are selected for a large range of SQI values and there is often a large overlap between the observed SQI values for several MCS indexes. For example, the median SQI was equal to 40 for MCS 1-4 with Nighthawk in the lobby (Fig. 5(c)), MCS 2-4 with Nighthawk in the corridor (Fig. 5(e)), and MCS 1-3 with Talon in the corridor (Fig. 5(f)). Interestingly, Fig. 5(d) shows

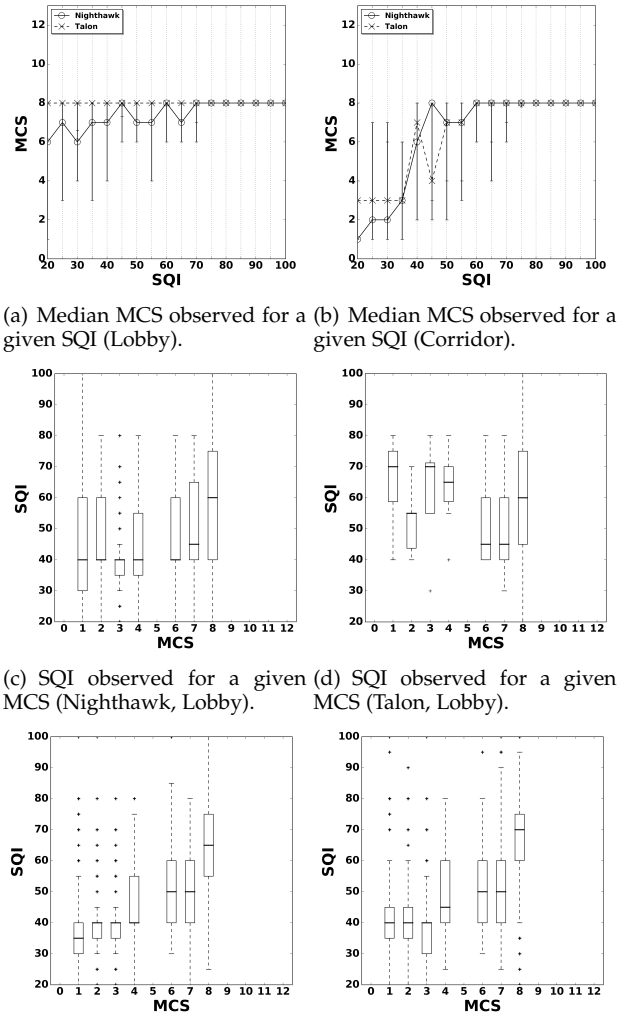

(a) Median MCS observed for a given SQI (Lobby).
(b) Median MCS observed for a given SQI (Corridor).
(c) SQI observed for a given MCS (Nighthawk, Lobby).
(d) SQI observed for a given MCS (Talon, Lobby).
(e) SQI observed for a given MCS (Nighthawk, Corridor).
(f) SQI observed for a given MCS (Talon, Corridor).

Fig. 5. Correlation between MCS and SQI.

that high median SQI values for lows MCSs (1-4) were observed in the case of Talon in the lobby (much higher than for MCS 6, 7), suggesting that the rate adaptation logic made wrong decisions dropping the MCS in spite of high link quality. We note that, in our recent study [7], we confirmed the weak correlation between MCS and SNR for indoor 60 GHz links using a software defined radio platform.

4.2 Orientation

We analyze the impact of the relative orientation between the AP and client since practical phased arrays cannot generate homogeneous beams across all directions [20]. We place the AP and client facing each other at a distance of 50 ft in the Lobby.

We first rotate the client from -90^o to $+90^o$. Surprisingly, we find that the Rx sector *never* changes, even at extreme angles, and Gbps communication is possible for Rx angles in $[-60^o, 75^o]$ with Nighthawk and $[-75^o, 75^o]$ with Talon. This suggests that the laptop uses quasi-omni beam patterns during reception. The same is true for reception with the two routers. This is even more surprising given the long communication ranges in Section 4.1 and suggests that, contrary to common belief, Gbps rates can be achieved even if high beamforming gains are only available at one side.

We then keep the client at 0^o and rotate the AP from -90^o to $+90^o$. Fig. 6(a) shows that high data rates can be sustained under large Tx angular displacement (although smaller than the Rx angular displacement), with both routers. With Nighthawk, we observe high data rates in the range $[-45^o, 30^0]$ and a gradual drop for larger angles. No connection was established for $\pm 90^o$. The performance with Talon is even better; the average data rate remains above 1 Gbps for the whole range $[-90^o, 90^0]$, with the exception of 60^0. Our results also show that performance and sector selection are asymmetric with respect to the 0^o orientation for both routers. This is caused by two factors. The beam patterns themselves are not symmetric, and the lobby is not symmetric so that reflected paths may be available on one side but not the other.

Fig. 6(b), 6(c) plot the distribution of Tx sector IDs selected by Nighthawk and Talon, respectively, for each Tx angle. We observed that both routers use different sectors at each angle, and in certain cases more than one sector for a given angle. This is especially true for Talon where up to six different sectors were selected for certain angles (-60^o, 60^o, 90^o). We study the (sub)optimality of the selected sectors for Talon by manually configuring all possible sectors, one at a time, in sequence, and measuring throughput with each one. Fig. 6(d) plots the fraction of throughput obtained by the beam-selection scheme with respect to the optimal (highest throughput) sector. Fig. 6(d) shows that the AP selects the optimal sector in almost *all* cases except for -90^o, -60^o and 45^o-75^o. While in the 45^o-75^o case this suboptimal selection has a relatively small impact on performance, as it still achieves more 75% or more of the optimal throughput, in other cases it results in significant loss (50% in case of -60^o). Fig. 6(d) also shows that, contrary to expectation, larger (extreme) angles do not necessarily result in more suboptimal sector selection. In fact, for equal amount of rotation in the clockwise vs. counter-clockwise directions can result in significantly different performance. This can be attributed to asymmetry in both the beam patterns themselves and the objects/surfaces in the environment.

5 Coverage

While in Section 4 we evaluated the impact of Tx-Rx distance and angular separation on performance separately, we now look at the impact of both factors together by evaluating the coverage in the whole lobby. We place the router (Tx) at two different locations using a different orientation for each location. The two locations (Tx1, Tx2), and their orientations are shown in Fig. 11(a). We then divide the lobby into a grid of 8×8 ft squares, and measure the throughput at the center of each square, with two different client orientations, marked as Or1, Or2, in Fig. 11(a). The results for the two routers are shown in the form of heatmaps in Fig. 7 and 8.

Fig. 7 and 8 confirm our previous findings about the range. However, the performance is the result of the combined effect of distance and relative Tx-Rx orientation. At both AP locations, when the client has orientation Or1, there are many more positions where the relative Tx-Rx angle falls within $[-60^o, 60^o]$, which is required for high data rates. In contrast, with client orientation Or2, a large part of the lobby

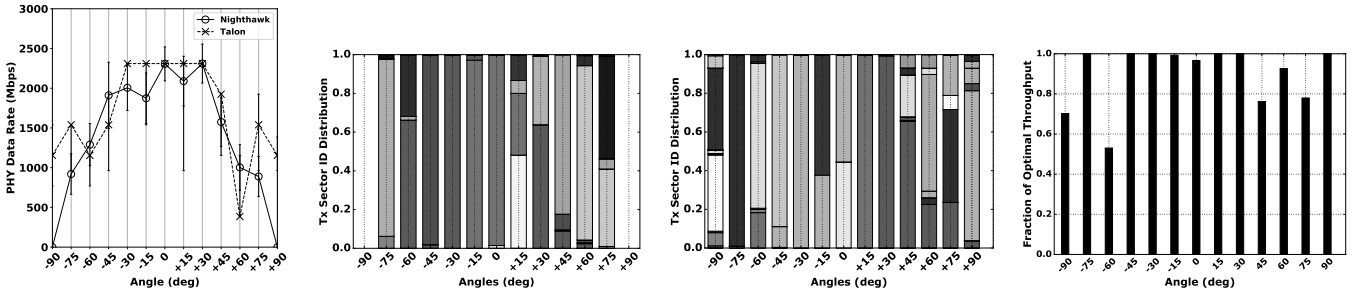

(a) Performance for different Tx angles with both routers. (b) Tx Sector ID distribution for different Tx angles with Nighthawk. (c) Tx Sector ID distribution for different Tx angles with Talon. (d) Fraction of optimal throughput achieved due to suboptimal Tx sector selection with Talon.

Fig. 6. Impact of Tx-Rx Orientation. The AP is rotated from -90^o to $+90^o$ while the client is kept at 0^o.

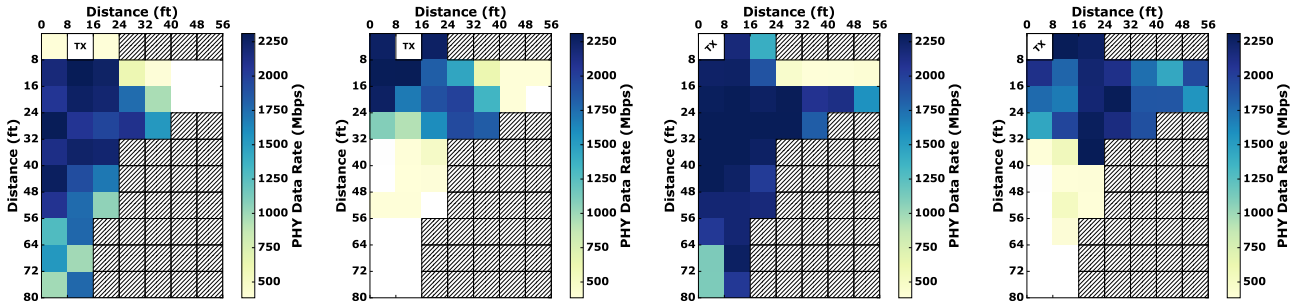

(a) AP at Tx1, client orient. 1. (b) AP at Tx1, client orient. 2. (c) AP at Tx2, client orient. 1. (d) AP at Tx2, client orient. 2.

Fig. 7. Coverage with Nighthawk.

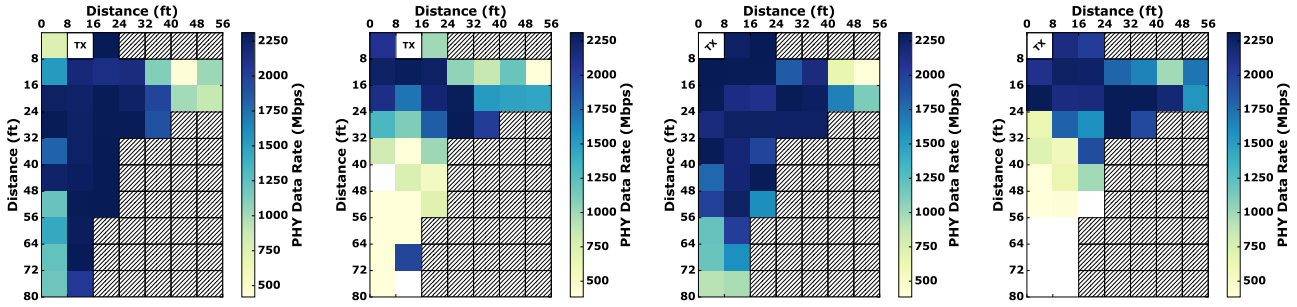

(a) AP at Tx1, client orient. 1. (b) AP at Tx1, client orient. 2. (c) AP at Tx2, client orient. 1. (d) AP at Tx2, client orient. 2.

Fig. 8. Coverage with Talon.

in not covered. We also observe a few outliers with high data rate positions surrounded by very low data rate positions, e.g., in Fig. 8(b), which are the result of wall reflections.

A comparison between Fig. 7 and Fig. 8 further reveals that coverage depends on the physical location of the antenna placement in the AP. For example, with the AP at Tx1, Talon offers much better coverage than Nighthawk for both client orientations (only 2 white squares in Fig. 8(a), 8(b) vs. 14 in Fig. 7(a), 7(b)). With the AP at Tx2, coverage is again better with Talon but there are certain locations where performance is higher with Nighthawk.

Fig. 9(a) and 9(b) plot point clouds of the IDs of the dominant sectors used by Nighthawk as a function of the Tx-Rx distance and relative angle for the two Rx orientations. Fig. 9(c) and 9(d) repeat the same plots for Talon. The angle is measured using the Tx as a reference. Further, Rx orientation is not taken into account in the angle calculation. For example, the angle for the upper left square is $+90^o$ in both Fig. 7(a) (orientation Or1) and Fig. 7(b) (orientation Or2). Each point corresponds to one driver log collected at a given square (combination of angle and distance). More points of a given color imply that the corresponding sector ID was used more often. For readability, we use randomization to avoid overlapping points. We randomly place each point within a rectangular area of $\pm 3.5^o$ on the x axis and ± 3.5 ft on the y axis.

Fig. 9(a) depicts our results for Nighthawk and Rx orientation Or1. We observe two large clusters for IDs 20 and 10. Interestingly, we observe a large degree of asymmetry. Sector 20 is mainly used for angles in $[0^o, +20^o]$, except for very short distances, where due to the specific placement of the client almost under the AP, a different sector (16) is preferred. For larger positive angles, the AP switches to sector 10. On the other hand, for negative angles, we observe a number of different sectors, among them 255 (failure to find a valid sector) corresponding to white and yellow squares in Fig. 7(a) and 7(c). In Fig. 9(b), corresponding to client orientation Or2, clusters are not as distinct as in Fig. 9(a). Still, we observe a highly asymmetric pattern, with many more sectors used for negative angles. In contrast to Fig. 9(a), this time failures (sector 255) appear for positive angles. When comparing Talon (Fig. 9(c) and 9(d)) against

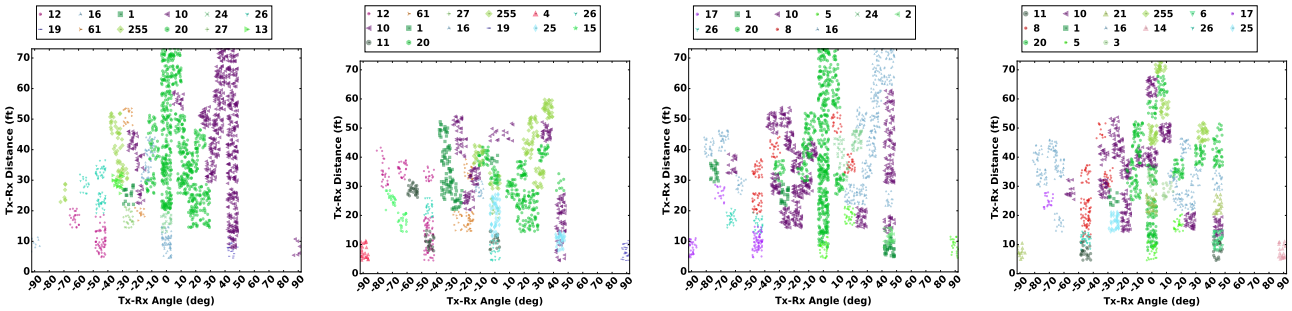

(a) Nighthawk, client orient. 1. (b) Nighthawk, client orient. 2. (c) Talon, client orient. 1. (d) Talon, client orient. 2.

Fig. 9. Tx sector IDs as a function of Tx-Rx distance and and relative angle.

Nighthawk, we observe that the two main sectors (20 and 10) are used in a similar manner. Still, Talon uses more sectors for positive angles compared to Nighthawk, and the sectors used for short distances are also different from those used by Nighthawk. Finally, note that in all four figures (Fig. 9(a)-9(d)) some sector IDs are used for multiple angles, both positive and negative ones (e.g., ID 10 in Fig. 9(a), 9(b) or ID 16 in Fig. 9(c), 9(d)). This ID distribution suggests multiple strong sidelobes (which is confirmed in [21], otherwise the mapping of IDs to angles would have been more unique.

In summary, we make the following observations: (i) For the same angle, different sectors are used for very short distances compared to longer distances, due to the specific placement where the AP needs to beamform vertically down; (ii) the selected sectors exhibit a high degree of asymmetry w.r.t. the relative angle – different sets of sectors are selected for the same positive vs. negative angle and the same distance; (iii) the two routers often select different sectors for the same distance and relative angle. Overall, it is hard to predict what sectors will be selected at a given location. This strongly depends on factors such as the radiation patterns, the antenna locations, as well as the Tx-Rx distance and relative orientation. However, wrong selection can result in weak (or loss of) coverage. Our results suggest that a careful AP placement is essential to guarantee full coverage in large spaces, typically using more than one AP despite the excellent range. E.g., our results suggest that at least 3 APs are probably required for full coverage in the lobby, placed at the top left, top right, and bottom left corner, and facing towards the center of the room.

6 Deployment Considerations

In this section, we evaluate different WLAN deployment options. We begin with different AP placements in 6.1 and then study the performance under various NLOS scenarios in 6.2.

6.1 AP Placement

We evaluate three AP deployment options, in addition to the default one used in Section 4.1: (i) table – where the AP is placed on a table at the same height as the client, a typical deployment for legacy WiFi APs in home environments, (ii) wall-mounted, at a height of 6 ft – a typical deployment for legacy WiFi APs in enterprise environments, (iii) ceiling-mounted – a less common deployment for legacy WiFi APs but not unusual. Due to space constraints, we only report results with Nighthawk.

Fig. 10(a) evaluates the data rate achieved with the four options in the lobby. We observe that the table and ceiling-mounted option yield the best performance, with a data rate of at least 2.3 Gbps up to 60 ft (a longer range than the default option) and at least 1.25 Gbps at 80 ft. In contrast, the wall-mounted option results in faster and larger performance degradation – less than 1 Gbps at distances longer than 70 ft. Interestingly, placement on the table is the only option that results in an MCS higher than 8 (10, 12) and data rates as high as 3 Gbps. However, this only happens when the AP and client are placed next to each other, which is of limited practical use. On the other hand, placement on the ceiling sustains higher throughputs at longer ranges. Additionally, this option is more resilient to human blockage.

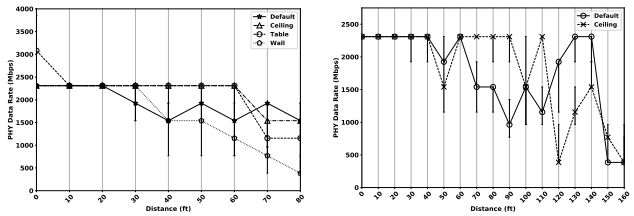

(a) Lobby (b) Corridor

Fig. 10. Evaluation of different placement options.

Fig. 10(b) compares the performance of the default and ceiling-mounted options in the corridor. We observe that both placements exhibit similar performance for up to 60 ft, the ceiling placement performs better for distances between 60-110 ft, and the default placement performs better at longer distances. Interestingly, in the case of a ceiling-mounted AP, beamforming fails to find a working sector 55% and 15% of the time at 120 ft and 160 ft, respectively, resulting in extremely low throughput. At all other distances, sector 20 is used, the same sector as with the default placement.

Overall, our results suggest that placing the AP on the ceiling generally yields high performance, especially in open spaces, but can result in outages (due to poor beamforming decisions) in narrow spaces (corridors). This shows the need for more intelligent and robust beamforming algorithms.

6.2 NLOS Performance

We study the performance of common office NLOS scenarios in a lab with three rows of desks, having metal partitions between them, and clutter such as computers, monitors, and wireless equipment, as shown in Fig. 11(b). We use four topologies, where Rx denotes the client position for all 4

topologies, and U-T, C, C-2, C-H denote the AP positions, respectively: (i) under-the-table with the AP placed under the big oval table and the client placed on the table; (ii) cubicle with one metal partition between the AP and client; (iii) cubicle-2 with one metal partition and one drawer between the AP and client; (iv) cubicle-high with the AP placed on a drawer higher than the client's position. In a fifth topology (Wall), we placed the AP and client on opposite sides of a wall making sure that the only available communication path is through the wall itself. Finally, two corner topologies are shown in Fig. 11(c); the space includes both drywall and glass walls. The performance with Nighthawk and Talon is shown in Fig. 12(a). We also plot the average percentage of outage time (fraction of 0-throughput samples over the total throughput samples) at each location in Fig. 12(b).

In contrast to common expectation, we observe that both routers work well in typical NLOS environments, including some particularly challenging ones such as "under-table" and "cubicle-2". Moreover, the outage time is below 5% in all topologies. In most cases, communication becomes feasible through reflections. Surprisingly, Talon achieves average data rates of at least 1.5 Gbps in all seven topologies. On the other hand, Nighthawk performs very well in the "cubicle" and "wall" topologies but poorly in the "under-table" and "corner" topologies. We also observed that the two routers use different sectors and different MCS for the same topology.

The result for the "wall" topology is of particular interest for practical deployment purposes. While it has been often argued that very dense deployments of 60 GHz APs (at least one AP per room) are required for WiFi-like coverage, our results show that this is not necessarily the case. Although the result in Fig. 12(a) was obtained with the AP and client very close to the wall, we also varied the distance between the two devices up to 16 ft, by moving either one or both of them away from the wall, with no performance degradation.

To further understand the feasibility of through-wall communication, we conducted a larger set of experiments. We placed the AP outside the lab (spot "Tx" in Fig. 11(b)) and ensured no reflections can be used for communication. We measured throughput at 10 different locations inside the lab. The client locations (spots 1 to 10) and orientation are also shown in Fig. 11(b). Performance is far from uniform (Fig. 12(c)), with average data rates at different locations varying from 0 (location 2 for Nighthawk and 10 for both routers) to as high as 2.3 Gbps (location 7 for Talon and 6 for both routers). A careful inspection of the floorplan reveals that the performance is a combined effect of distance, orientation, and blockage, confirming our conclusions in Sections 4 and 5.

Further, we again observe that Talon (the older of the two devices) has a favorable antenna placement and provides better coverage and performance. To the best of our knowledge, this is the first work to explore in detail through-wall communication in 60 GHz using COTS 802.11ad hardware and discuss implications in WLAN deployments.

Remarks. Overall, our results in Sections 4, 5, 6 reveal two interesting, and potentially conflicting, factors that affect indoor 60 GHz performance and can provide useful deployment guidelines for 60 GHz WLANs. On one hand, performance depends heavily on the Tx-Rx angular separation due to the fact that practical phased arrays cannot generate homogeneous beams across all directions; this dependence limits the tolerable misalignment between a pair of devices and implies that more than one APs might be required to provide full coverage in large rooms. On the other hand, Gbps communication is feasible through drywall (a typical material for office walls at least in the US), suggesting that one AP can provide coverage in more than one rooms. For example, one might be able to cover a blind spot in a room (due to large angular separation) with an AP in a neighboring room or in the corridor. At the same time, through-wall communication can potentially increase the interference footprint and limit spatial reuse – making the deployment in large enterprise environments challenging, as the standard only supports 4 orthogonal channels [22]. In addition, performance heavily depends on the antenna placement on the AP. Together, our results indicate that WLAN deployment in 60 GHz is much more challenging that in legacy WiFi and show the need for intelligent placement algorithms, taking into account all these factors.

7 Blockage

We study the robustness to blockage in two different scenarios: (i) mobile client scenario, where the client moves behind a blocking object and (ii) static client scenario, where a human moves into the LOS between the client and the AP.

7.1 Mobile Client, Wall Blockage

We repeated the following experiment five times with each router: starting from a LOS position, the client moves at a constant speed (3 ft/s) behind a wall until the link breaks, stands for 1 s at the point where connectivity is lost, and then moves back at the same speed until the link is re-established. Fig. 13 shows the MAC throughput (obtained from the driver), Tx sector ID, and MCS timelines for one run with each router. The other four runs gave similar results.

We focus on Nighthawk. Fig. 13(a) shows that throughput drops to 0 after 6 s and it recovers after 18 s. Fig. 13(c) shows that the driver starts reporting sector ID 255 and status "Retrying" at the 6th second. At the 13th second, the connection is completely lost; and at the 14th second, the client starts moving back towards the AP. Nonetheless, no sector ID is reported until the 21st second. Similarly, MCS remains at 1 during the interval 5-13 s, and no MCS is reported during 13-21 s. Finally, note that even though a valid sector is found at the 21st second, it takes 3 more seconds for communication to be fully established (non-zero throughput only starts at the 24th second). We observe a similar behavior for Talon in Fig. 13(b), 13(d), with the exception that, in one of the 5 runs, the link was never re-established even though the driver reported a working sector after 12 s.

We used a second Talon router in monitor mode, placed on the floor between the client and the AP with its antenna oriented towards the client, to investigate the reason behind the long period of outage. From the traces, we noticed that after several attempts of unsuccessful SLS execution, both sides halt their data transmission and their beamforming training attempts roughly for 10-12 s, followed by repeated

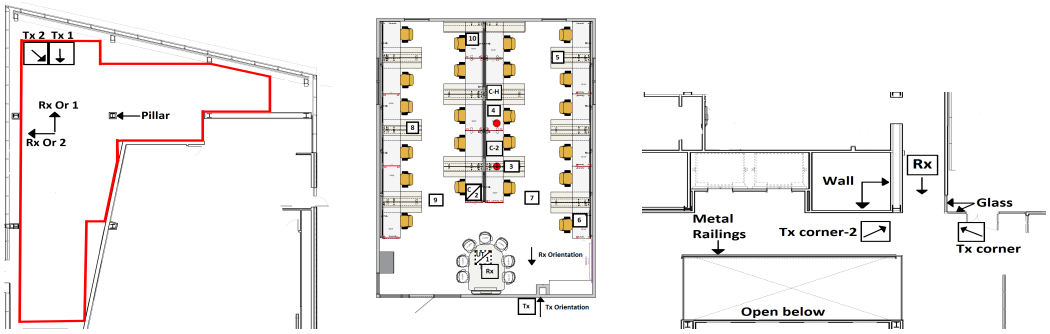

(a) Lobby floorplan. (b) Lab floorplan. (c) Corner scenario floorplan.

Fig. 11. Coverage and NLOS topologies.

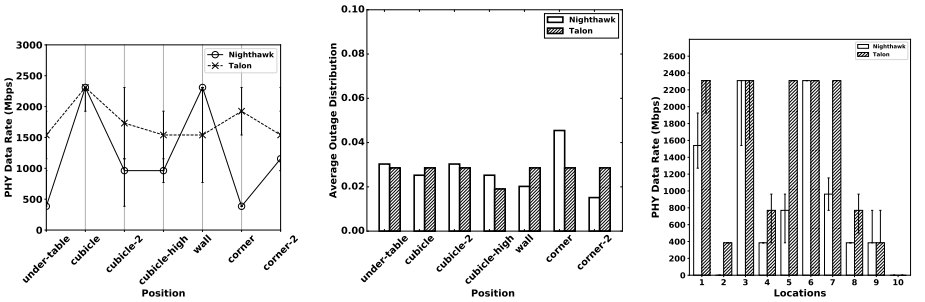

(a) Median PHY Data Rate. (b) Percentage of outage time. (c) Performance through wall.

Fig. 12. Evaluation in various NLOS topologies.

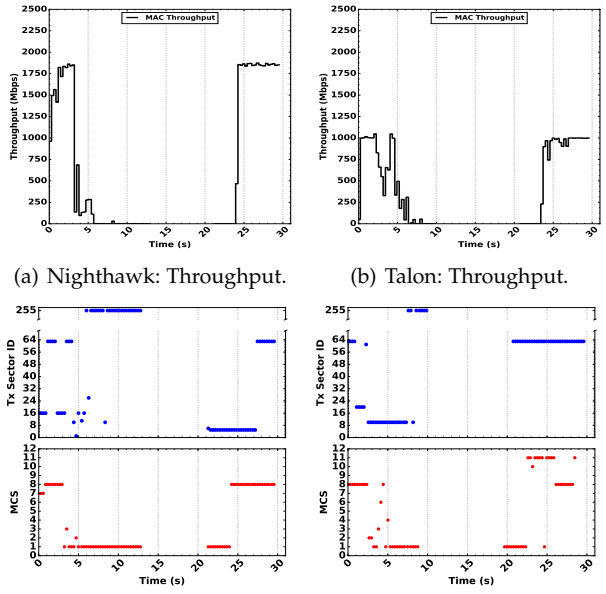

(a) Nighthawk: Throughput. (b) Talon: Throughput.

(c) Nighthawk: Tx sector ID, MCS. (d) Talon: Tx sector ID, MCS.

Fig. 13. Mobile client, static blockage – Timelines of MAC throughput, Tx sector ID, and MCS.

re-association attempts spanning an extra 1-2 s. During this period, the client considers the AP is not reachable and disassociates itself internally from the AP. After that, the client initiates a successful SLS phase. As a result, the client sends a Probe Request frame to the AP with wild card SSID to represent all SSIDs. The AP replies back with a Probe Response frame indicating the SSID it broadcasts. Then, the client waits for almost 5-6 BIs and associates again with the AP. However, the AP still considers the client as associated and, as a response to the association request, it disassociates the client by sending a Disassociation frame to the client. Again, this results in client waiting for some BIs to reassociate again with the AP. After this association attempt, data communication finally resumes between the client and the AP. This behavior can be due to improper implementation of the association state-machine in the driver and the firmware which causes this high latency in re-establishing the communication link.

7.2 Static Client, Human Blockage

7.2.1 Quasi-Static Blockage

We keep the AP and the client in LOS, with the AP placed at a height of 6 ft and the client at a height of 1.5 ft, and we initiate a 30 s TCP iperf session. 5 s after the iperf session starts, a human moves into the LOS between the AP and the client, stands at the same position for 20 s, and then moves out. We consider three different blockage positions: near the Tx, in the middle, and near the Rx. We repeated each experiment five times. Fig. 14 presents boxplots of the MAC layer goodput during the 20 s of blockage for both routers in each scenario.

Our experiments reveal different behavior for the two routers. The performance of Nighthawk is severely affected at 20 ft for all three blockage positions, with the median PHY data rate during blockage dropping below 1 Gbps. Different from the case of mobile blockage, here the AP tried 4-5 alternative sectors but none of them resulted in better performance. On the other hand, the impact was less severe for longer distances (40 ft and 60 ft) with a median PHY data rate above 1.9 Gbps; in particular, in the case of blockage in the middle, the performance remained largely unaffected at 40 ft and 60 ft. In these cases, the AP was able to find alternative paths through reflections using the same subset of sectors as in the case of mobile blockage.

The performance of Talon was less affected at 20 ft, with the median PHY data rate during blockage equal to 1.2 Gbps

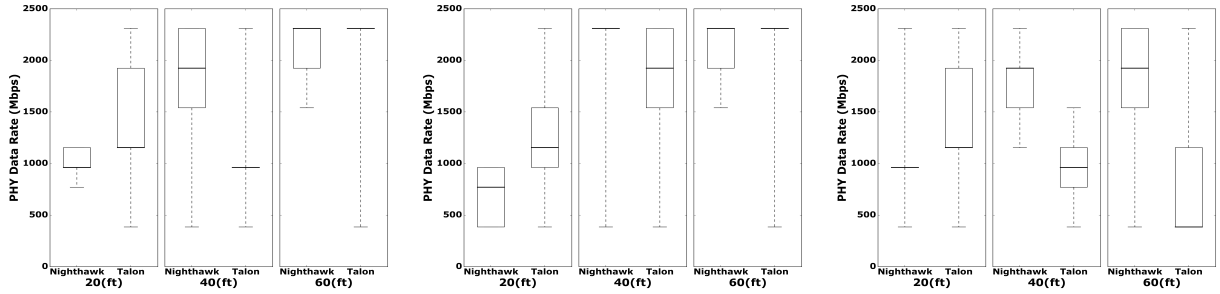

(a) Near the receiver. (b) Middle. (c) Near the transmitter.

Fig. 14. PHY data rate during quasi-static human blockage.

for all three blockage positions. At longer distances, Talon surprisingly maintained excellent performance when the blockage was in the middle, without switching to a different sector after blockage. On the other hand, the performance dropped significantly in the case of blockage near the Tx or Rx at 40 ft (with a median data rate of 962 Mbps) and in the case of blockage near the Tx only at 60 ft (with a median rate of only 385 Mbps and outage intervals as long as 9 s).

7.2.2 Mobile Blockage

We use the same setup and location as in Section 7.2.1. We initiate a 30-s TCP iperf session during which a human moves continuously in and out of the LOS, blocking the link at different random positions. We consider three different AP-client distances: 20 ft, 40 ft, and 60 ft. For each distance, we repeated the experiment five times with each router. Our results reveal that both routers are resilient to mobile human blockage and throughput never dropped to zero for any of the three distances. Nonetheless, we observe significant differences in the performance across different distances and across the two routers for the same distance.

Fig. 15 show an example of the PHY data rate and Tx sector ID timelines for one run for each of the distances with Nighthawk and Talon, respectively. The remaining 4 runs with each router are similar. The data rate with Nighthawk dropped significantly at many blockage events at 20 ft, but always recovered within 1-2 s. The AP used a single sector ID (20) during almost the whole 30-s duration. It is likely that due to the short AP-client distance, alternative paths via reflections are hard to find, and the beamforming algorithm ends up back to the same sector after every beam training event. The performance was much better at 40 ft and 60 ft, where the AP switched multiple times between 2-3 and 3-4 different sectors, respectively, maintaining a high quality link, with no/few drops of the data rate below 1 Gbps.

Surprisingly, the results with Talon were quite different. The performance was best at 60 ft, followed by 20 ft, and then by 40 ft. Interestingly, at 20 ft, Talon was able to discover alternative working paths via different sectors (16, 63) more often than Nighthawk. On the other hand, at 40 ft, it used one extra sector compared to Nighthawk, which proved a poor choice most of the time. Note that this diversity in performance between devices and sector selection under both quasi-static and mobile human blockage is heavily dependent on the phased antenna array used by the devices. Despite Nighthawk and Talon having the same chipset, the antenna integration in the device itself has a strong impact on the beam pattern shape and thus, on which reflections are available in each scenario.

Remarks. Overall, we conclude that the current beamforming algorithm suffers from two issues: (i) in the case of mobile clients moving behind blocking objects and back, connectivity cannot typically be re-established at the point where the link initially broke, and (ii) in the case of human blockage, sometimes beamforming fails with detrimental effect to performance. Nonetheless, the results in the case of human blockage are in general encouraging, showing overall better performance than in previous studies. Overall, our results show the need for more intelligent beamforming algorithms (e.g., maintaining a number of failover paths [23] or incorporating more advanced techniques [24], [25]) that will be able to discover existing alternative paths without resorting to beam training at all.

8 Mobility

We consider two simple cases of mobility: (i) Moving towards the AP: The client starts at a distance of 180 ft away from the AP and moves towards the AP at a speed of 3 ft/s; (ii) Moving away from the AP: The client starts in front of the AP and moves away up to a distance of 180 ft at the same speed. In both cases, a TCP iperf session is run for the duration of the motion. We repeated five runs for each type of motion.

Fig. 16, plotting the average PHY data rate over distance for both motion types with each router, shows that Talon is much more robust to motion compared to Nighthawk. When the client moves away from the AP, Talon maintains an average PHY data rate of at least 2.3 Gbps up to 67 ft, while Nighthawk's data rate drops below 2.3 Gbps at 37 ft. For longer distances, Talon's average data rate remains higher than 1 Gbps while Nighthawk's rate drops to much lower levels. Similarly, when the client moves towards the AP, Talon's average rate remains above 1 Gbps for distances shorter than 157 ft, while Nighthawk sustains rates above 1 Gbps only for distances shorter than 90 ft. We also found that the average fraction of outage time (due to beamforming failures) remains below 30% with Talon and drops to zero for distances shorter than 110 ft (Fig. 17(c), 17(d)), while it reaches up to 85% with Nighthawk and only drops to 0 for distances shorter than 70 ft (Fig. 17(a), 17(b)).

Overall, we observe that beamforming failures can significantly hurt the performance in case of mobility, and different devices can exhibit very different levels of robustness. Note that in these two simple cases of motion, where the client moves on a straight line always facing

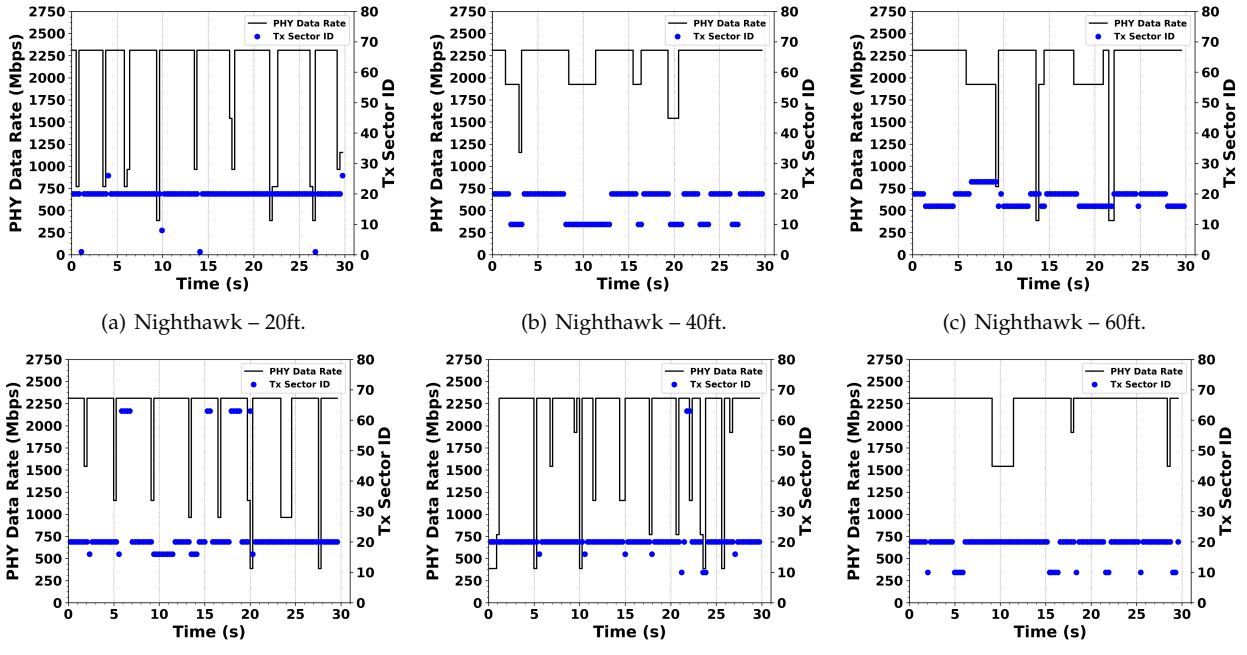

(a) Nighthawk – 20ft. (b) Nighthawk – 40ft. (c) Nighthawk – 60ft.

(d) Talon – 20ft. (e) Talon – 40ft. (f) Talon – 60ft.

Fig. 15. PHY data rate and Tx sector ID timelines under mobile blockage with Nighthawk and Talon.

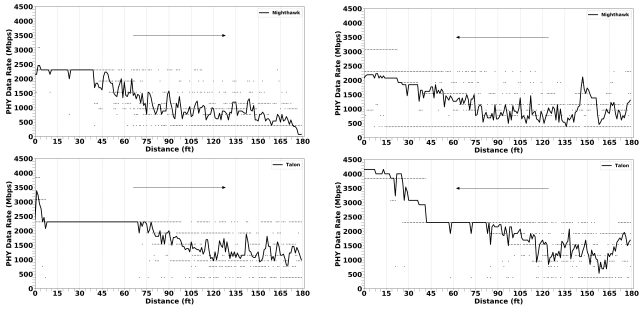

(a) Away from AP. (b) Towards the AP.

Fig. 16. Timeline of average goodput during client mobility. The dots show the values for each of the 5 runs.

the AP, intuitively the same sector should always work; there is no need for beamforming at all. Unfortunately, RSS changes due to mobility trigger beam training, and when the messages exchanged during training are lost, the beamforming algorithm fails. Note also that a simple memory-based algorithm (maintaining a set of previously working sectors and trying them before performing training again) would also trivially solve the problem in this case. The challenge here is that the AP does not know the type of motion or the cause of RSS drop (e.g., in case of blockage, beam training might indeed be required) and always resorts to beam training even when this is not necessary. This shows again the need for intelligent adaptation algorithms that can react differently in different scenarios.

9 Discussion

We show that the *actual* challenges of practical consumer-grade 60 GHz networks are not always in line with the issues one would expect based on the propagation characteristics at such frequencies. We find unexpected issues but also issues which in practice are not as critical as earlier work assumes.

Non-critical challenges. Transient human blockage causes link degradation but has a smaller impact than reported in related work [26]. The underlying reason is that 802.11ad APs perform much more frequent sector sweeps than earlier hardware [3]. Still, we achieve very high data rates. That is, the cost of frequent sector sweeps is lower than suggested in the literature. Our packet traces show that a sector sweep takes less than 1 ms, whereas blockage typically occurs at the timescale of 100s of ms. Sweeps are short due to the devices not training their receive sectors, but using a quasi-omni pattern for reception. Our measurements reveal that this does not have a strong impact on performance even for long links. This shows that highly directional communication is not as critical as predicted in earlier work. Also, wall attenuation is limited, enabling an AP to serve clients in different rooms despite the use of wide beam patterns.

Unexpected challenges. The interaction of beam training and rate control plays a much more important role than the literature suggests. Current hardware takes wrong beam and rate decisions even in very simple scenarios such as a LOS link in a static environment. The impact of such errors propagates through the protocol stack, having a massive impact on upper layers such as TCP. Further, beam steering accuracy strongly degrades at angles beyond 60°, which limits the coverage area of an AP. That is, multiple APs may be needed within a room but not due to attenuation or blockage, but due to the limited steering capabilities of 60 GHz devices. This is partially due to the device casing causing self-shadowing. The Talon router often performs much better than Nighthawk just because its antenna array is much more exposed. Thus, the network deployment is closely related AP's form factor.

Expected challenges. As discussed in earlier work, movement and rotation have a strong impact on performance. Mobile scenarios perform particularly poorly in our experiments since link adaptation often fails. To improve this, 60

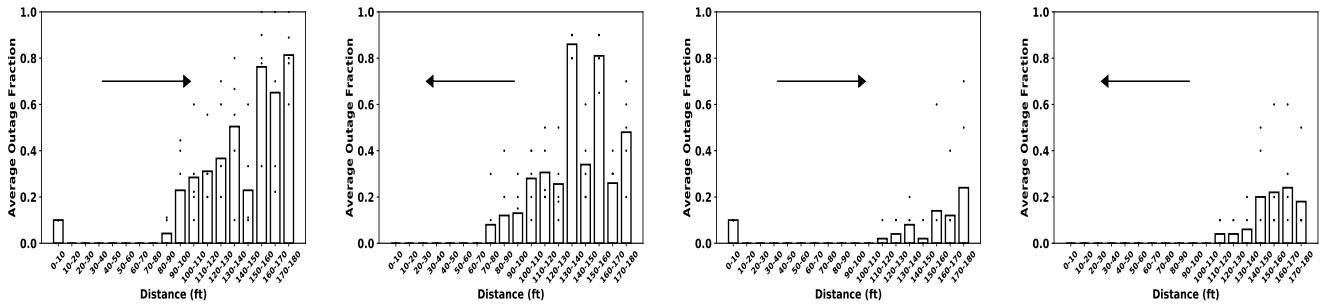

(a) Nighthawk, average outage moving away from AP.
(b) Nighthawk, average outage moving towards AP.
(c) Talon, average outage moving away from AP.
(d) Talon, average outage moving toward AP.

Fig. 17. Average fraction of outage time during mobility. The dots show the values for each of the 5 runs.

GHz networks need better control algorithms that use historic information or are able to interpret SNR drops better. For instance, if an SNR drop occurs because the link length increases, rate control should handle the issue. However, in current hardware this often triggers beam training. Also, rate control often takes place based on link quality indicators, which for 60 GHz networks are even more unrelated to the actual state of the channel than at lower frequencies. As a result, performance is highly unpredictable.

10 Related Work

Channel Measurements and Link Characterization. An enormous amount of work has focused on characterizing 60 GHz for indoor/outdoor channels using dedicated channel sounding hardware (e.g., [19], [27]–[34]). Some works focused on modeling human blockage impact on the performance of 60 GHz links [35], [36].

SDR Approaches. The platform of choice for academic research until now, for varied reasons, has been primarily an SDR for baseband generation (USRP, WARP, etc.) coupled with upconverters/down converters and horn antennas [6], [23], [24], [37]–[41]. Such setups typically face the following limitations: (i) absence of MAC and higher layers, (ii) baseband limited to few hundreds of MHz and (iii) use of mechanically steerable horn antennas. With these limitations, it becomes hard to say whether experimental results obtained from such platforms can account for the often complex interactions between PHY, MAC, and upper layers of the stack, the wider bandwidth used by 802.11ad, and the non-uniformity and imperfectness of beams formed by commercial phased antenna arrays. While recent testbeds, e.g., [7] address some of these limitations (wideband transmission and use of phased arrays), the main challenge of non-standard compliant PHY/MAC implementations remains.

Practical Work on 60 GHz COTS Devices. A number of recent works [8], [9], [42], [43] have conducted experimental studies using WiGig COTS hardware [3]. While these devices allow researchers to study performance across multiple layers of the protocol stack and can provide insights into the operation of antenna arrays, they suffer from a number of limitations: (i) they are based on WiGig and implement a proprietary association protocol, not fully standard-compliant, (ii) they are targeted towards short-range, LOS, semi-static P2P link use cases, rather than a WLAN scenario, and hence, they are not designed to deal efficiently with blockage or client mobility, and (iii) they export only limited lower layer information to the user offering limited insights. While some researchers [8] have managed to obtain a deeper understanding of lower layer operations of these devices using a signal analyzer, all the works based on this hardware are primarily focused on the performance of a single link. In contrast, we use 802.11ad-compliant COTS routers in our study. Our access to a richer set of link parameters allows us to obtain much deeper insights into the reasons for specific performance results. Further, our study goes beyond basic link characterization and explores for the first time practical considerations in WLANs such as coverage and AP deployment. A few more recent studies [25], [44]–[47] conduct experiments using 802.11ad compliant hardware. However, they focus on performance comparison between 60 GHz and legacy WiFi [44], multi-AP coordination [25], [47], 802.11ad power consumption [45], and transport protocol aspects [46], and hence, are complementary to this work.

Our results on the communication range are in sharp contrast with the results reported in most of the works using WiGig hardware [8], [9], [18] (a range of 70 ft for MCS 1) but closer to the results reported in a few more recent studies [42], [43]. We also note that previous works using either proprietary channel sounding hardware (e.g., [48]) or narrowband SDRs [6] have reported that drywall only induces a 2-3 dB loss and measurements with pre-802.11ad hardware [42], [43], [49] have shown that Gbps communication is possible through a wall. Nonetheless, to our best knowledge, this is the first work to explore in detail range and through-wall communication in 60 GHz using COTS 802.11ad hardware and discuss potential implications in WLAN deployments.

11 Conclusions

We analyze the performance of COTS consumer-grade 802.11ad hardware. In contrast to earlier work in this area, our hardware fully complies with the standard, and we focus on deployment aspects such as indoor coverage, AP orientation, and impact of antenna placement. While our insights partially match the prevailing wisdom in the 60 GHz community, our measurements also reveal both unexpected challenges and challenges which are not as critical as suggested in the literature. The former includes steering accuracy and device casing self-shadowing, whereas the latter includes range, transient blockages, and beam sweep overhead. We provide a detailed study of these issues, which is crucial to enable researchers in the field to focus on the most relevant practical problems.

12 Acknowledgements

This work has been supported in part by NSF grant CNS-1553447, the ERC project SEARCHLIGHT grant no. 617721, the Ramon y Cajal grant RYC-2012-10788, the Madrid Regional Government through the TIGRE5-CM program (S2013/ICE-2919), the German Research Foundation (DFG) in the Collaborative Research Center (SFB) 1053 MAKI, by the German Federal Ministry of Education and Research (BMBF) and the State of Hesse within CRISP-DA, and the Hessian LOEWE excellence initiative within NICER.

References

[1] IEEE, "FCC Promotes Higher Frequency Spectrum for Future Wireless Technology," https://www.fcc.gov/document/fcc-promotes-higher-frequency-spectrum-future-wireless-technology, 2015.

[2] T. Nitsche, C. Cordeiro, A. B. Flores, E. W. Knightly, E. Perahia, and J. C. Widmer, "IEEE 802.11ad: directional 60 GHz communication for multi-Gigabit-per-second Wi-Fi," *IEEE Communications Magazine*, vol. 52, no. 12, 2014.

[3] *Dell ships its WiGig-based Wireless Dock, gives your Latitude a home base for $249*. [Online]. Available: http://www.engadget.com/2013/02/25/dell-ships-its-wigig-based-wireless-dock-for-latitude/

[4] *Intel Tri-Band Wireless-AC 18265*. [Online]. Available: https://www.intel.com/content/www/us/en/wireless-products/tri-band-wireless-ac-18265.html

[5] Qualcomm, "Qualcomm Technologies' Tri-Band Solution," https://www.qualcomm.com/products/features/80211ad, 2017.

[6] S. Sur, V. Venkateswaran, X. Zhang, and P. Ramanathan, "60 GHz Indoor Networking through Flexible Beams: A Link-Level Profiling," in *Proc. of ACM SIGMETRICS*, 2015.

[7] S. K. Saha *et al.*, "X60: A programmable testbed for wideband 60 ghz wlans with phased arrays," *Elsevier Computer Communications (COMCOM)*, vol. 133, pp. 77–88, 2019.

[8] T. Nitsche, G. Bielsa, I. Tejado, A. Loch, and J. Widmer, "Boon and Bane of 60 GHz Networks: Practical Insights into Beamforming, Interference, and Frame Level Operation," in *Proc. of the 11th ACM CoNEXT*, December 2015.

[9] Y. Zhu, Z. Zhang, Z. Marzi, C. Nelson, U. Madhow, B. Y. Zhao, and H. Zheng, "Demystifying 60ghz outdoor picocells," in *Proc of ACM MobiCom*, 2014.

[10] *TP-Link Talon AD7200 Multi-Band Wi-Fi Router*. [Online]. Available: http://www.tp-link.com/us/products/details/cat-5506_AD7200.html

[11] *Netgear Nighthawk X10 Smart WiFi Router*. [Online]. Available: https://www.netgear.com/landings/ad7200/

[12] *Acer TravelMate P446-M*. [Online]. Available: https://www.acer.com/ac/en/US/content/professional-series/travelmatep4

[13] *LEDE Project*. [Online]. Available: https://lede-project.org/

[14] I. . T. G. AD, "IEEE 802.11ad, Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band," 2012.

[15] *Wireless HD*. [Online]. Available: http://www.wirelesshd.org/

[16] D. Halperin, S. Kandula, J. Padhye, P. Bahl, and D. Wetherall, "Augmenting data center networks with multi-gigabit wireless links," in *Proc. of ACM SIGCOMM*, 2011.

[17] X. Zhou, Z. Zhang, Y. Zhu, Y. Li, S. Kumar, A. Vahdat, B. Y. Zhao, and H. Zheng, "Mirror Mirror on the Ceiling: Flexible Wireless Links for Data Centers," in *Proc. of ACM SIGCOMM*, 2012.

[18] Y. Zhu, X. Zhou, Z. Zhang, L. Zhou, A. Vahdat, B. Y. Zhao, and H. Zheng, "Cutting the Cord: a Robust Wireless Facilities Network for Data Centers," in *Proc. of ACM MobiCom*, 2014.

[19] P. F. M. Smulders, "Statistical characterization of 60-ghz indoor radio channels," *IEEE Transactions on Antennas and Propagation*, vol. 57, no. 10, pp. 2820–2829, October 2009.

[20] T. S. Rappaport, R. W. H. Jr., R. C. Daniels, and J. N. Murdock, *Millimeter Wave Wireless Communications*. Prentice Hall, 2014.

[21] D. Steinmetzer, D. Wegemer, M. Schulz, J. Widmer, and M. Hollick, "Compressive Millimeter-Wave Sector Selection in Off-the-Shelf IEEE 802.11 ad Devices," in *Proc. of the 11th ACM CoNEXT*, 2017.

[22] IEEE, "IEEE Standards 802.11ad-2012, Amendment 3: Enhancements for Very High Throughput in the 60 GHz Band," https://standards.ieee.org/standard/802_11ad-2012.html, 2012.

[23] M. K. Haider and E. W. Knightly, "Mobility Resilience and Overhead Constrained Adaptation in Directional 60 GHz WLANs: Protocol Design and System Implementation," in *Proc. of ACM MobiHoc*, 2016.

[24] S. Sur, X. Zhang, P. Ramanathan, and R. Chandra, "BeamSpy: Enabling Robust 60 GHz Links Under Blockage," in *Proc. of USENIX NSDI*, 2016.

[25] T. Wei and X. Zhang, "Pose Information Assisted 60 GHz Networks: Towards Seamless Coverage and Mobility Support," in *Proc. of ACM MobiCom*, 2017.

[26] A. Loch, I. Tejado, and J. Widmer, "Potholes ahead: Impact of transient link blockage on beam steering in practical mm-wave systems," in *European Wireless Conference*, 2016.

[27] A. Maltsev, R. Maslennikov, A. Sevastyanov, A. Khoryaev, and A. Lomayev, "Experimental investigations of 60 GHz WLAN systems in office environment," *IEEE Journal on Selected Areas in Communications (JSAC)*, vol. 27, no. 8, pp. 1488–1499, October 2009.

[28] C. R. Anderson and T. S. Rappaport, "In-building Wideband Partition Loss Measurements at 2.5 and 60 GHz," *IEEE Transactions on Wireless Communications*, vol. 3, no. 3, pp. 922–928, 2004.

[29] J. S. Lu, P. Cabrol, D. Steinbach, and R. V. Pragada, "Measurement and characterization of various outdoor 60 ghz diffracted and scattered paths," in *IEEE MILCOM*, 2013.

[30] H. Xu, V. Kukshya, and T. S. Rappaport, "Spatial and temporal characteristics of 60-ghz indoor channels," *IEEE Journal on Selected Areas in Communications (JSAC)*, vol. 20, no. 3, pp. 620–630, April 2002.

[31] P. F. M. Smulders and L. M. Correia, "Characterisation of propagation in 60 GHz radio channels," *Electronics & Communication Engineering Journal*, vol. 9, no. 2, pp. 73–80, April 1997.

[32] A. I. Sulyman, A. Alwarafy, G. R. MacCartney, T. S. Rappaport, and A. Alsanie, "Directional radio propagation path loss models for millimeter-wave wireless networks in the 28-, 60-, and 73-ghz bands," *IEEE Transactions on Wireless Communications*, vol. 15, no. 10, pp. 6939–6947, Oct 2016.

[33] E. Ben-Dor, T. S. Rappaport, Y. Qiao, and S. J. Lauffenburger, "Millimeter-wave 60 ghz outdoor and vehicle aoa propagation measurements using a broadband channel sounder," in *IEEE GLOBECOM*, 2011.

[34] R. C. Daniels, J. N. Murdock, T. S. Rappaport, and R. W. Heath, "60 ghz wireless: Up close and personal," *IEEE Microwave Magazine*, vol. 11, no. 7, pp. 44–50, Dec 2010.

[35] M. Jacob, S. Priebe, T. Krner, M. Peter, M. Wisotzki, R. Felbecker, and W. Keusgen, "Fundamental analyses of 60 ghz human blockage," in *Proc. of EuCAP*, 2013.

[36] ——, "Extension and validation of the ieee 802.11ad 60 ghz human blockage model," in *Proc. of EuCAP*, 2013.

[37] T. Nitsche, A. B. Flores, E. W. Knightly, and J. Widmer, "Steering with Eyes Closed: mm-Wave Beam Steering without In-Band Measurement," in *Proc. of IEEE INFOCOM*, 2015.

[38] S. Naribole and E. Knightly, "Scalable Multicast in Highly-Directional 60 GHz WLANs," in *Proc. of IEEE SECON*, 2016.

[39] A. Zhou, X. Zhang, and H. Ma, "Beam-forecast: Facilitating Mobile 60 GHz Networks via Model-driven Beam Steering," in *Proc. of IEEE INFOCOM*, 2017.

[40] T. Wei, A. Zhou, and X. Zhang, "Facilitating Robust 60 GHz Network Deployment By Sensing Ambient Reflectors," in *Proc. of USENIX NSDI*, 2017.

[41] H. Hassanieh, O. Abari, M. Rodriguez, M. Abdelghany, D. Katabi, and P. Indyk, "Fast Millimeter Wave Beam Alignment," in *Proc. of ACM SIGCOMM*, 2018.

[42] S. K. Saha, V. V. Vira, A. Garg, and D. Koutsonikolas, "Multi-Gigabit Indoor WLANs: Looking Beyond 2.4/5 GHz," in *Proc. of IEEE ICC*, 2016.

[43] ——, "A Feasibility Study of 60 GHz Indoor WLANs," in *Proc. of IEEE ICCCN*, 2016.

[44] S. Sur, I. Pefkianakis, X. Zhang, and K.-H. Kim, "WiFi-Assisted 60 GHz Networks," in *Proc. of ACM MobiCom*, 2017.

[45] S. K. Saha, T. Siddiqui, D. Koutsonikolas, A. Loch, J. Widmer, and R. Sridhar, "A Detailed Look into Power Consumption of Commodity 60 GHz Devices," in *Proc. of IEEE WoWMoM*, 2017.

[46] H. Assasa, S. K. Saha, A. Loch, D. Koutsonikolas, and J. Widmer, "Medium Access and Transport Protocol Aspects in Practical 802.11ad Networks," in *Proc. of IEEE WoWMoM*, 2018.

[47] S. Sur, I. Pefkianakis, X. Zhang, and K.-H. Kim, "Towards Scalable and Ubiquitous Millimeter-Wave Wireless Networks," in *Proc. of ACM MobiCom*, 2018.

[48] B. Langen, G. Lober, and W. Herzig, "Reflection and transmission behavior of building materials at 60 GHz," in *Proc. of IEEE PIMRC*, 1994.
[49] X. Tie, K. Ramachandran, and R. Mahindra, "On 60 GHz wireless link performance in indoor environments," in *Proc of PAM*, 2012.

Swetank Kumar Saha received the B. Tech. degree in Computer Science and Engineering (CSE) from IIIT-Delhi, New Delhi, India, in May 2013. He is currently working towards a PhD in CSE at the University at Buffalo, The State University of New York. His research interests include experimental wireless networking and mobile computing. He is a student member of the IEEE.

Shivang Aggarwal received his Bachelor of Technology (B.Tech.) degree in Computer Science and Engineering from Manipal Institute of Technology, Karnataka, India in July 2016. He also received his Master of Science (M.S.) degree in Computer Science and Engineering from the University at Buffalo, The State University of New York, USA in December 2017. He is currently pursuing his Ph.D. in Computer Science at the University at Buffalo. His research interests include experimental wireless networking, specifically in the millimeter wave band. He is a student member of the IEEE and SIGMOBILE.

Hany Assasa Hany Assasa is currently a post-doctoral researcher at IMDEA Networks, Madrid, Spain. He obtained his Ph.D. from Universidad Carlos III de Madrid in July 2019. In his Ph.D. studies, Hany focused on building efficient, robust, and reliable millimeter-wave wireless networks both in theory and real-world testbed deployments. He is the lead developer of the open-source IEEE 802.11ad/ay project in network simulator ns-3, which is widely used by the research community. During his Ph.D., Hany carrier out internships at NEC Laboratories Europe, Heidelberg, Germany on 5G multi-RAT mm-wave access and at imec/PharrowTech, Leuven, Belgium to utilize their 60 GHz hybrid MU-MIMO testbed. Before his Ph.D., Hany has obtained a double degree master in Information and Communication Technologies (ICT) from Politecnico di Torino (Italy) and KTH Royal Institute of Technology (Sweden). In addition, Hany spent one year at Ericsson Research AB (Sweden).

Adrian Loch is currently a post-doc researcher at IMDEA Networks in Madrid, Spain. He graduated in Electrical Engineering from Universidad Politécnica de Madrid and Technische Universität Darmstadt in 2011 after completing an international double degree program. After that, he obtained a PhD in Computer Science from Technische Universität Darmstadt in March 2015. During his PhD, he was a research associate at the Secure Mobile Networking Lab. His main areas of interest lie in cooperative communications for both wireless access and wireless multihop networks, including routing issues as well as practical validation on wireless testbeds. Currently, he focuses on millimeter-wave communications and, in particular, wireless LANs such as in the 802.11ad standard.

Naveen Muralidhar Prakash received his B.E. degree in Computer Science and Engineering(CSE) from Anna University, Chennai, India. He graduated with a Master's degree in CSE from the University at Buffalo, The State University of New York, in Feb 2018. During his Master's, he worked in the UB WiNS Lab on experimental wireless networks, with a primary focus on millimeter wave networking.

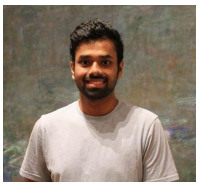

Roshan Shyamsunder received his B.E degree from Sri Venkateswara College of Engineering, Anna University, Chennai, India in April 2016. He also received his M.S degree from the University at Buffalo, The State University of New York in June 2018. He is currently pursuing his doctoral studies in Computer Science and Engineering at Northwestern University with research interests in content distribution and prototype design and implementation of protocols and algorithms for the internet.

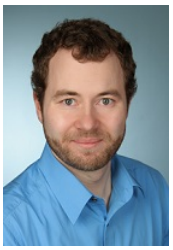

Daniel Steinmetzer Daniel Steinmetzer received the bachelors and masters degrees in information system technology from TU Darmstadt, Germany, in 2012 and 2014, respectively. He joined Secure Mobile Networking Lab in 2014, and currently focuses his research on secure and robust millimeter-wave communication systems. His other interests are physical layer security and wireless communications in general.

Dimitrios Koutsonikolas received the PhD degree in electrical and computer engineering from Purdue University, in 2010. He worked as a post-doctoral researcher at Purdue University from September to December 2010. He is currently an associate professor of computer science and engineering at the University at Buffalo, The State University of New York. His research interests are broadly in experimental wireless networking and mobile computing. He received an IEEE Region 1 Technological Innovation (Academic) Award in 2019, the UB Teaching Innovation Award in 2018, the UB SEAS Senior Teacher of the Year Award in 2017, the NSF CAREER Award in 2016, the UB SEAS Early Career Researcher of the Year Award in 2015, and Best Paper Awards from mmNets 2019, WCNC 2017, and SENSORCOMM 2007. He is a senior member of the IEEE and ACM and a member of USENIX.

Joerg Widmer Joerg Widmer is Research Professor and Research Director of IMDEA Networks in Madrid, Spain. Before, he held positions at DOCOMO Euro- Labs in Munich, Germany and EPFL, Switzerland. His research focuses on wireless networks, ranging from extremely high frequency millimeter-wave communication and MAC layer design to mobile network architectures. Joerg Widmer authored more than 150 conference and journal papers and three IETF RFCs, and holds 13 patents. He was awarded an ERC consolidator grant, the Friedrich Wilhelm Bessel Research Award of the Alexander von Humboldt Foundation, a Spanish Ramon y Cajal grant, as well as best paper awards at IEEE ICC, IEEE PMIRC, IEEE WoWMoM, ICST WICON, IEEE MediaWiN, NGC, and the IEEE Communications Society Best Tutorial Paper Award. He is Senior Member of IEEE and Distinguished Member of ACM.

Matthias Hollick Matthias Hollick received the Ph.D. degree from TU Darmstadt, Germany, in 2004, where he is currently heading Secure Mobile Networking Lab, Computer Science Department. He has been researching and teaching at TU Darmstadt, Universidad Carlos III de Madrid, and the University of Illinois at UrbanaChampaign. His research focus is on resilient, secure, privacy-preserving, and quality-of-service-aware communication for mobile and wireless systems and networks.