


# 60 GHz indoor WLANs: insights into performance and power consumption

Swetank Kumar Saha<sup>1</sup>  · Darshan Godabanahal Malleshappa<sup>1</sup> · Avinash Palamanda<sup>1</sup> · Viral Vijay Vira<sup>2</sup> · Anuj Garg<sup>3</sup> · Dimitrios Koutsonikolas<sup>1</sup>

Published online: 4 March 2017  
© Springer Science+Business Media New York 2017

**Abstract** This paper presents a feasibility study of 60 GHz indoor WLANs. We evaluate 60 GHz performance in a typical academic office building under the primary assumption that 60 GHz WLAN APs and clients will be equipped with relatively wide-beam antennas to cope with client mobility. In contrast to previous works which measured performance at a single layer using custom, non-standard compliant hardware, we investigate performance *across multiple layers* using primarily 802.11ad-compliant wide-beam COTS devices. Our study shows that the large number of reflective surfaces in typical indoor WLAN environments combined with wider beams makes performance highly unpredictable and invalidates several assumptions that hold true in static, narrow-beam, Line-Of-Sight scenarios. Additionally, we present the first measurements, to our best knowledge, of power consumption

of an 802.11ad NIC and examine the impact of a number of factors on power consumption.

**Keywords** 60 GHz · WLAN · 802.11ad · Performance · Power consumption

## 1 Introduction

The use of millimeter-wave (mmWave) radios in the unlicensed 57–64 GHz spectrum (colloquially known as the 60 GHz band), which is supported by IEEE 802.11ad [47], has recently emerged as an alternative to the traditional 2.4/5 GHz WiFi, promising multi-Gigabit throughput. 802.11ad defines three 2.16 GHz channels and offers bitrates between 385 Mbps and 6.76 Gbps. However, since free-space loss scales up with the square of the carrier frequency, the propagation loss at 60 GHz is 21.6 dB worse than at 5 GHz. Further, due to the short wavelength, 60 GHz signals are easily blocked by obstacles such as walls or humans. To overcome these challenges, 60 GHz radios are typically highly directional, introducing new challenges in scenarios involving device mobility.

Due to these characteristics, until recently, the use of the 60 GHz technology had been limited to static, short-range, LOS scenarios, e.g., for wireless docking or for augmenting data center networks with high capacity wireless links [11, 44, 46]. Signal propagation is easy to model in these scenarios as it exhibits near-free space propagation properties [11, 29, 44]. However, the true potential of the mmWave technology cannot be realized if its use is limited to such scenarios. Recent work [45] demonstrated the feasibility of 60 GHz outdoor picocells. Another scenario of increasing interest is the use of the 60 GHz technology for building *multi-gigabit indoor WLANs* [32, 35]. The

---

✉ Swetank Kumar Saha  
swetankk@buffalo.edu

Darshan Godabanahal Malleshappa  
darshang@buffalo.edu

Avinash Palamanda  
apalaman@buffalo.edu

Viral Vijay Vira  
viralvvira@gmail.com

Anuj Garg  
anujgarg1208@gmail.com

Dimitrios Koutsonikolas  
dimitrio@buffalo.edu

<sup>1</sup> University at Buffalo, The State University of New York, Buffalo, NY 14260-2500, USA  
<sup>2</sup> Amazon LLC, 1915 Terry Avenue, Seattle, WA 98101, USA  
<sup>3</sup> FactSet Research Systems Inc., Merritt 7 Corporate Park, 601 Merritt 7, Norwalk, CT 06851, USA

typical indoor enterprise WLAN environment is highly complex, with many objects/surfaces/moving humans that can attenuate, completely block, or reflect the signal, making it harder to predict link behavior.

Additionally, these scenarios imply that battery-powered mobile devices will be the next target for the mmWave technology. Recently SiBeam announced the first 802.11ad equipped smartphone [25]. A study from ABI Research predicts that smartphones will account for nearly half of 802.11ad chipset shipments in 2018 [2]. However, improved communication speeds generally come at the cost of higher power consumption. Studies using 802.11n/ac chipsets and smartphones have shown that power increases with PHY data rate [10, 22, 38] and channel width [22, 41], as well as with application layer throughput [12, 31]. 802.11ad offers much higher data rates compared to 802.11ac and an order of magnitude wider channels, which can result in significantly increased power. Hence, it becomes essential for chip designers to understand the factors that affect power consumption.

### 1.1 Contributions

This paper is one of the first to investigate the feasibility of building general-purpose 60 GHz indoor WLANs by evaluating 60 GHz client performance and power consumption in a typical academic office building. We make the following contributions. First, *in contrast to previous works* [9, 11, 19, 32, 35, 44], *which measured performance at a single layer (typically PHY)* [9, 19, 32, 35, 44] *often using custom, non-standard compliant hardware and mechanically steerable horn antennas* [9, 11, 32, 35, 44], *we investigate performance across multiple layers of the protocol stack using (primarily) 802.11ad-compliant COTS devices equipped with electronically steerable phased antenna arrays.* Although the use of commodity hardware limits our access to certain PHY layer information (e.g., RSS) and our control on certain protocol features (e.g., rate adaptation, beamforming), it allows us to obtain a much better idea of how the 802.11ad technology would behave when used to build full-scale WLANs. To enhance our understanding of signal propagation in indoor WLAN environments, we repeat a subset of our measurements using non-standard compliant hardware and horn antennas of various beamwidths. Second, recent work [32] has shown that, although narrow-beam antennas can greatly extend range, they yield poor performance in scenarios involving client mobility and human blockage. Hence, our study assumes that *60 GHz WLAN APs and clients will use relatively wide-beam antennas to cope with mobility.* Our results differ significantly in several aspects from both theoretical modeling and experimental results reported in the past, suggesting that *findings from using narrow-beam*

*antennas in simple, reflection-free environments, or based on simple theoretical models do not always apply in more complex environments.* Third, *we present what we believe to be the first measurements of power consumption of an 802.11ad NIC.*

### 1.2 Summary of findings

Our study adds to the findings of recent experimental studies in indoor 60 GHz WLAN environments by answering the following questions.

(1) *What is the expected performance in different indoor WLAN environments?* Our results confirm that high-throughput 60 GHz communication is feasible with COTS devices equipped with wide-beam antennas at various setups typical of an indoor WLAN environment (corridors, halls, labs, through walls or glass).

(2) *What is the impact of antenna orientation on performance in these environments?* We found that communication is possible in the case of antenna array misalignment, either via beamsteering or through an NLOS link via reflection, but the performance can vary significantly depending on the location and Tx-Rx distance and can be much lower than in the case of optimal orientation.

(3) *How does Tx/RX distance affect performance at different layers?* We found that communication is possible at distances longer than 100 ft but performance is generally unpredictable and highly dependent on the environment (type and number of reflective surfaces). Further, signal propagation in the case of wide-beam antennas in indoor WLANs cannot be characterized by simple log-distance path loss models, which have been extensively used in 802.11ad simulators [11, 35, 44, 46].

(4) *Is there any correlation among performance metrics at different layers?* We found that RSSI can only serve as a weak indicator of PHY data rate and TCP throughput and only at certain locations, but not across locations. Further, PHY data rate is not always a good indicator of higher layer performance. Hence, translating signal strength to PHY data rate or PHY data rate to higher layer performance, a common practice in recent measurement studies [32, 45] due to limitations of available 60 GHz hardware,<sup>1</sup> can yield inaccurate results in typical indoor WLAN environments. Similarly, signal-strength based rate adaptation algorithms which have been used in recent simulation studies [11, 35, 44, 46] may not perform well in such environments.

(5) *What is the impact of human blockage on the performance of 60 GHz links in indoor WLAN environments?*

<sup>1</sup> Commercial hardware often does not report PHY data rate and custom hardware [39] does not provide throughput commensurate with 802.11ad rates.

Our results show that human blockage, especially by humans standing near the Tx or Rx, presents a major challenge for 60 GHz links in indoor environments. However, a client served by two APs simultaneously can maintain 100% uptime and high throughput.

(6) *What is the power consumption of an 802.11ad NIC compared to that of an 802.11ac NIC? What is the impact of factors such as signal strength, PHY data rate, and packet size, which are known to affect WiFi power consumption, on the 802.11ad power consumption?* We found that an 802.11ad NIC consumes much higher power than legacy WiFi (802.11n/ac) NICs but its much higher data rates makes it significantly more energy efficient. Interestingly, the average 802.11ad power consumption is not affected by Tx-Rx distance, PHY data rate, or RSSI and only slightly increases with packet size.

(7) *Recent studies have shown that the overhead of the 802.11ad's beam searching process may be prohibitively high, potentially nullifying the benefits of electronically steerable antenna arrays. What is the impact of this process on power consumption?* We found that the beam searching algorithm after a link outage can incur a significant amount of power consumption, in addition to the performance penalty which was observed by previous studies.

Some areas that this paper does not investigate, because they were studied in previous works, include mobility [9, 32, 45], interference patterns and spatial reuse [19, 32, 45], communication through reflections [19], and the benefits from using relays [23].

## 2 Related work

### 2.1 60 GHz in indoor environments

Our work is not the first to investigate the feasibility of 60 GHz technology in indoor WLANs. Initial experimental works focused on measuring channel propagation characteristics using dedicated channel sounding hardware (e.g., [4, 15, 29, 30, 40]). A few works also measured [7], simulated [24], or studied analytically [8] the impact of human blockage.

More recent works measured upper layer performance. Tie et al. [35] studied IP-over-wireless-HDMI performance of 60 GHz links with respect to blockage and antenna orientation using custom designed non-802.11ad hardware. Sur et al. [32] conducted a link-level profiling of indoor 60 GHz links, using a custom software defined radio platform (WiMi) [39], based on WARP [37] and the Vubiq 60 GHz development system [36]. The same platform was used in [34] to study the impact of human blockage and in [43] to explore beam steering. WiMi uses a small channel

width of only 245 MHz and thus, it cannot achieve Gbps data rates. Hence, throughput results in these works are extrapolated from RSS and noise floor measurements in narrow channels and they may not reflect the behavior of real 802.11ad links. Haider et al. [9] used a similar testbed with an even smaller channel width (20 MHz) and presented only RSSI measurements. Additionally, Vubiq transceivers use horn antennas with fixed beamwidth. In contrast to these works, we are using COTS 802.11ad-compliant hardware equipped with phased array antennas and measure performance *across different layers* of the protocol stack via TCP data transfers.

The only works that used phased antenna arrays in their evaluation are [1, 19, 42]. OpenMili [42] is an improved version of WiMi using rudimentary 4-element phased array antennas (each antenna element's weight can only be either  $e^0$  or  $e^{j\pi}$ ). The use of FPGAs in OpenMili allows for Gbps rates. However, its supported channel width of only 1 GHz is still much smaller than 802.11ad's channel width. Abari et al. [1] used a proprietary phased array testbed operating at 24 GHz instead of 60 GHz. Nitsche et al. [19] studied interference, beamforming, and frame aggregation using the same COTS hardware we use in this study. Since their focus was on studying beam patterns, most of their results were obtained by analyzing signal power traces obtained with an oscilloscope. Our study, focusing on performance across layers and the interaction among layers, is complementary to theirs.

Recent work also has argued for the use of 60 GHz technology to augment data centers [11, 44, 46] and demonstrated the feasibility of this approach using both expensive proprietary devices [11, 44, 46] and the same COTS hardware we use in this paper [46]. The data center environment, with static LOS links established on top of TOR switches (free from reflections), is very different from the complex indoor enterprise WLAN environment, and several of our findings are very different from the findings of these works.

### 2.2 60 GHz in outdoor environments

Channel sounding measurements have also been conducted in outdoor environments [14, 20, 30] and a few works simulated the impact of human blockage [3]. Several works have considered the use of 60 GHz technology for building outdoor mesh networks or backhauling, e.g., in 5G networks [18, 26–28, 33]; all these proposals have only been evaluated in simulations. More recently, Zhu et al. experimentally demonstrated the feasibility of 60 GHz-based outdoor picocells [16, 45] using both COTS 802.11ad and proprietary non-802.11ad hardware. The outdoor picocell scenario differs significantly from the one we are

concerned with, as also pointed out in [32], and several observations reported in that work do not hold for our use-case.

### 3 Measurement methodology

#### 3.1 Hardware

Our 802.11ad link setup (Fig. 1a) consists of two COTS devices: a Dell Latitude E420 laptop equipped with a Wilocity wil6210 802.11ad radio and a Dell Wireless Dock D5000. The dock has an 802.11ad WNIC and acts as an AP. Another laptop is connected to the dock through a Gigabit Ethernet interface to generate/receive TCP/UDP traffic. The use of the Ethernet interface limits the throughput in our experiments to 1 Gbps. The Wilocity radios are equipped with  $2 \times 8$  phased array antennas with relatively wide main beams ( $30^\circ - 40^\circ$ ) [32, 46], have lower EIRP (23 dBm [45]) than the maximum allowed by FCC (40 dBm), and support PHY data rates in the range 385–3850 Mbps. They export to the user-space the current


PHY data rate and an RSSI value between 0 and 100. They do not allow us to control the PHY data rate and use their own rate adaptation algorithm and an in-built beamforming mechanism to control beam properties. In case the link is blocked, the radios automatically search for an alternative NLOS path to re-establish the connection.

In order to better understand some of the results, we repeated a subset of our experiments with a pair of HXI Gigalink 6651 radios [13], each equipped with a horn antenna (Fig. 1b). A laptop is connected to each radio through a Gigabit Ethernet interface to generate/receive TCP/UDP traffic. An older model of the same hardware (6451) has been used in a number of measurement studies in data center and outdoor picocell environments [11, 44–46]. Unlike electronically steerable antenna arrays, a horn antenna points always towards a fixed direction and has a fixed beamwidth. We experimented with 3 different antenna beamwidths:  $8^\circ$  (24 dBi gain),  $12^\circ$  (24 dBi gain), and  $23^\circ$  (15 dBi gain), in both symmetric (same Tx and Rx beamwidth) and asymmetric (different Tx and Rx beamwidth) configurations. With a transmission power of 10 dBm, the EIRP is 25 dBm in the case of  $23^\circ$  beamwidth and 34 dBi in the case of  $8^\circ/12^\circ$  beamwidth. Note that the HXI radios operate on a proprietary, non-802.11ad configuration and do not support rate adaptation. A simple OOK modulation results in a transmission rate of 1.25 Gbps (although the use of the Ethernet interface limits again the throughput to 1 Gbps). Also, note that these radios are originally designed for outdoor, LOS communication, and they have no in-built mechanisms to combat multipath fading.

Finally, we repeated a smaller subset of the experiments with WiFi using two Dell Inspiron M5030 laptops, one configured as AP and the other one as client. Each laptop is equipped with a Mini PCI-e 802.11n/ac WiFi adapter featuring the Qualcomm-Atheros QCA9880 Version 2 chipset and controlled by the open source ath10k [5] driver. The card supports  $3 \times 3$  MIMO operation, channel widths of 20 MHz, 40 MHz, and 80 MHz, and PHY data rates in the range 6.5–1170 Mbps.

#### 3.2 Locations

For most of our experiments, we chose two locations inside an academic building in order to capture the diverse scenarios that are likely to occur in a office environment. The first location is an open *Hall* thinly populated by some desks and chairs and a staircase from the floor above. It offers better conditions for emulating near-free space propagation, reducing multipath effects. The ceiling height is rather high and thus it does not serve as a viable reflector. The second location is a rather narrow *Corridor* (5.08 ft wide) offering ample opportunities for reflection/multipath



**Fig. 1** Measurement Setup. **a** Dell Dock D5000. **b** HXI Gigalink



from the walls on the side, in addition to the floor and the ceiling.

Apart from these two locations, we also studied 60 GHz link behavior through different commonly found materials and some typical setups inside the building (Table 1), with different Tx/Rx orientations. We followed a methodology very similar to that in [35] and performed multiple experiments at each of the locations with same 16 different Tx/Rx orientations as in [35] (see Table 1 and Fig. 4b in [35]).

### 3.3 Methodology

We used iperf3 to generate traffic. Each experiment consists of a 10-second TCP or UDP session. All the results are the average of 10 sessions. All experiments were performed late night to remove the possibility of human blockage. The experiment environment consisted of only static objects present in the building.

We measured the power consumption of the wireless NIC, which comes in a Half-mini PCIE form factor, by plugging it to a PEX1-MINI-E PCI EXPRESS X1 to PCI Express Mini interface adapter [6] which can be powered from an external source. Since the client (Laptop) only exposes a Mini PCIE interface, we used a Mini PCIE to PCIE Express X1 Riser Card along with a high speed extender cable to connect the adapter to the the laptop’s Mini PCIE slot. Finally, we used a Monsoon Power Monitor [17] to supply power to the setup and record the power consumed. The setup is shown in Fig. 2. Due to limitations with our hardware (the Dock is a sealed box and does not expose any interface where our setup could be attached), we were only able to measure the receive power consumption. Based on previous studies on WiFi power consumption, we expect that the transmit power will be even higher.




Fig. 2 Power measurement setup

## 4 Performance measurements

### 4.1 Transmitter height

We begin by investigating the effect of dock/transmitter (Tx) height on performance as measured at the transport layer. We pick two different distances (8 and 16 ft) at each of the two locations and measure TCP and UDP throughput with different Tx height. To decide upon the optimal height for the Tx, we fix the height of the receiver (Rx) to 2’6” and vary the Tx height from 2’6” to 6’6” at the interval of 1’.

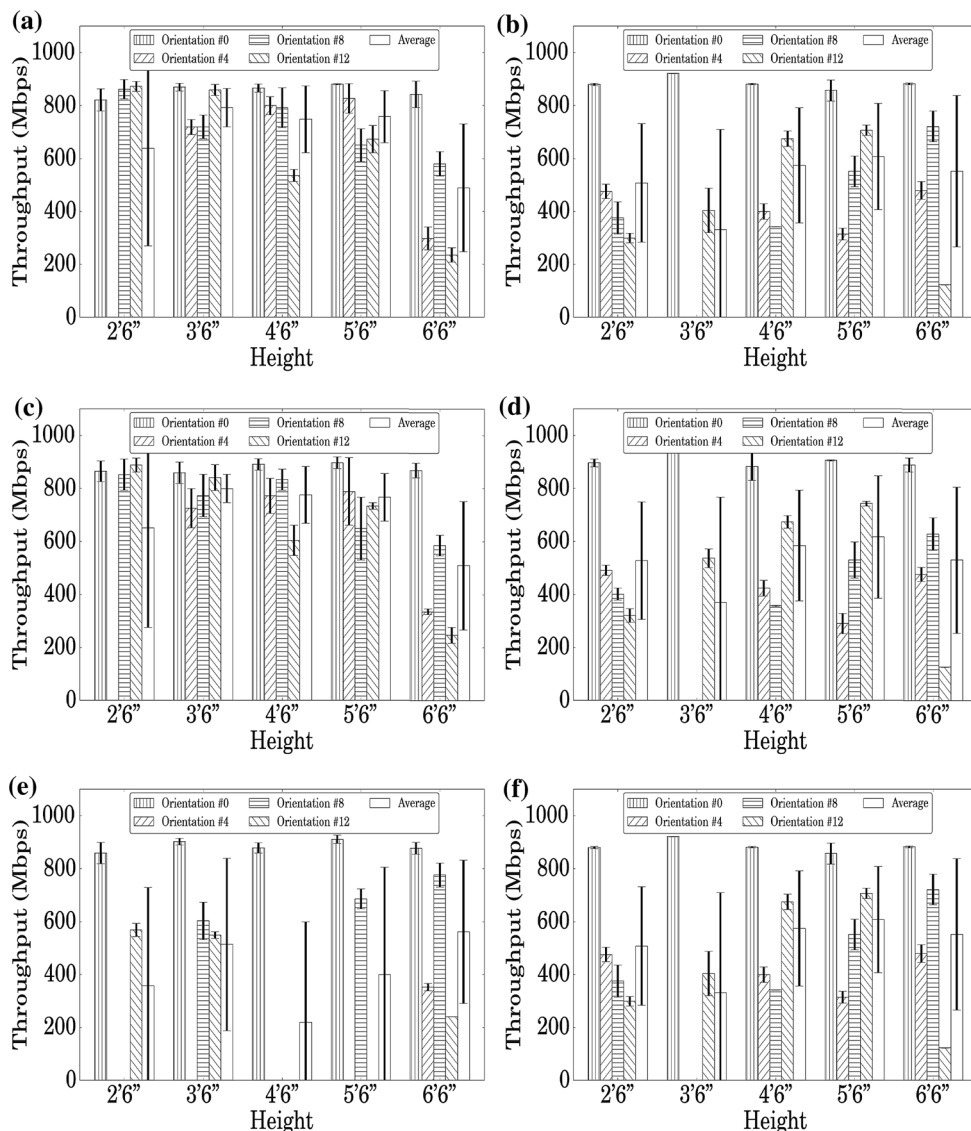
In a real deployed WLAN, the AP’s location and orientation will be fixed. However, mobile clients can be assumed to have any possible orientation with respect to the AP. We will discuss in detail in Sect. 4.3 how the relative orientation of the Tx and Rx antenna array is an important factor in determining the link performance. For this experiment, we chose four representative orientations: 0, 4, 8 and 12 (see Table 1) for the Rx and repeat our experiments at each one of them. Orientation #0 is the optimal case when both Tx and Rx antenna array are aligned. Orientations #4, 8 and 12 are obtained by rotating the Rx by 90°, 180° and 270°, respectively, while the AP/dock orientation remains fixed.

Figures 3a, c plot the TCP and UDP throughput achieved in the Hall for each of the four orientations and an average across all orientations for different Tx height. When both Tx and Rx are at the same height (2’6”), orientation #4 gives zero throughput, resulting in low average with a large standard deviation. As we increase the height of Tx to 3’6”, 4’6” and 5’6”, the average for all orientations approaches 800 Mbps, indicating that larger Tx height is more favorable to link performance. When the height is increased further to 6’6”, the throughput for all orientations deteriorates significantly, giving an average even lower

Table 1 Measurement Locations and Orientations

L#	Dist	Desc	Orientation					
			O#	Rx	Tx	O#	Rx	Tx
0	8’6”	Open Space	0	→	←	8	←	←
1	16’	Open Space	0	→	←	8	←	←
2	8’6”	Corridor/Symmetric	1	→	↓	9	←	↓
3	8’6”	Corridor/Asymmetric	2	→	→	10	←	→
4	16’	Corridor/Asymmetric	3	→	↑	11	←	↑
5	8’6”	Wall	4	↑	←	12	↓	←
6	8’6”	Glass	5	↑	↓	13	↓	↓
7	8’6”	Corner	6	↑	→	14	↓	→
8	8’6”	Lab	7	↑	↑	15	↓	↑
9	24’	Lab						

**Fig. 3** Variation of TCP/UDP throughput with Tx height. **a** Hall (8 ft): TCP Throughput. **b** Corridor (8 ft): TCP Throughput. **c** Hall (8 ft): UDP Throughput. **d** Corridor (8 ft): UDP Throughput. **e** Hall (16 ft): TCP Throughput. **f** Corridor (16 ft): TCP Throughput



than 2'6". However, in contrast to 2'6" Tx height, note that 6'6" Tx height can support all four orientations.

Figures 3b, d plot transport layer throughput for experiments done in the Corridor. Here, the average throughput increases for heights 4'6" and 5'6" but it reduces slightly for 6'6". The 3'6" height, which performed similarly to 4'6" and 5'6" in the hall, fails to provide connectivity at orientations #4 and #8. In this case, 5'6" seems to be optimal height providing the highest average throughput.

To further ensure that we choose the best possible Tx height, experiments at both locations were repeated with an increased distance of 16' between the Tx and Rx. Figures 3e, f plot the TCP throughput for Hall and Corridor, respectively (we omit UDP results which are very similar to TCP in the interest of space). In the Hall, 3'6" provides a throughput 100 Mbps higher than 5'6" but in the Corridor, 5'6" outperforms 3'6" by more than 200 Mbps. Lastly, in

the Corridor environment, 5'6" can support all four orientations while 3'6" can give non-zero throughput only for two orientations.

A general observation from these experiments is that *performance over 60 GHz indoor links is heavily affected by a number of factors – location, orientation, Tx height, Tx-Rx distance – and it is highly unpredictable; e.g., changing the Tx height by just 1' can improve or deteriorate TCP/UDP throughput by several hundreds of Mbps.* Hence, it is very hard to pick an optimal height for all possible locations, antenna orientations, and link distances. Overall, our experiments suggest that 5'6" performs better in general, although in some cases 3'6" shows better performance. We decided in favor of 5'6" since previous studies have concluded that larger Tx height results in propagation characteristics closer to the Friis model [32] as it reduces the chances of multipath caused by reflections

from the floor. Also, from a perspective of deployment of WLANs, a larger height for the AP is desirable for a larger coverage distance/area. For the rest of the experiments, unless stated otherwise, we use 5’6” as the height of the transmitter.

### 4.2 Performance across locations

In this section, we study the impact of location on the performance of 60 GHz links. Figures 4a, c plot the average RSSI, the selected PHY rates, and the average TCP throughput at each of the 10 locations in Table 1. We consider both orientation #0, which represents the case when both the Tx and Rx antenna arrays are fully aligned, and the average across 16 orientations.

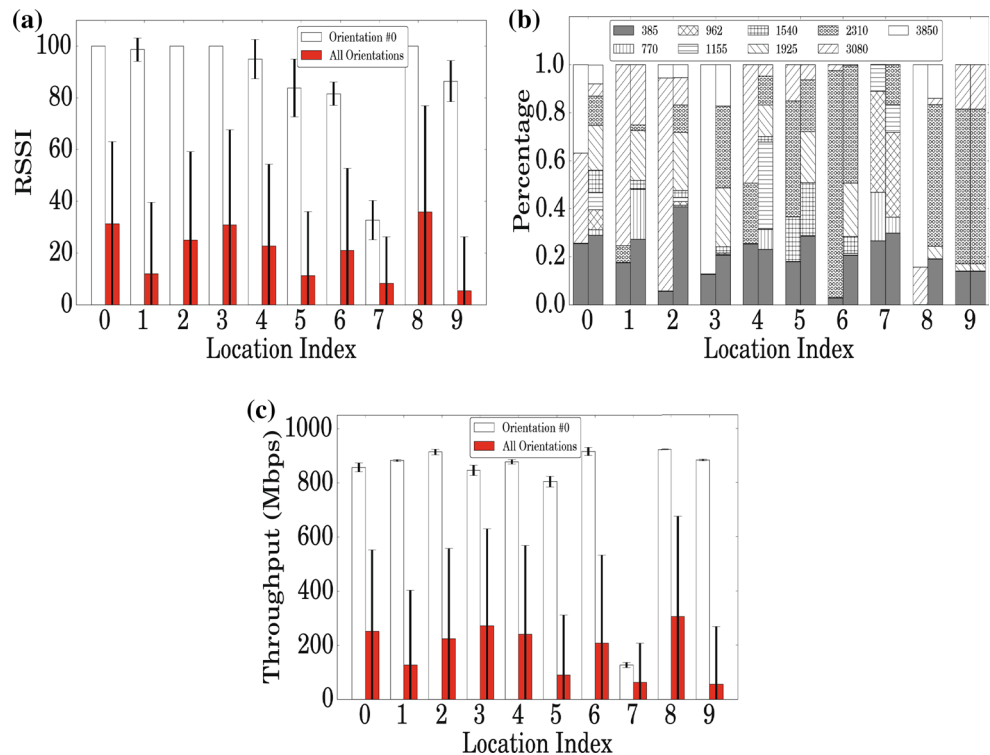
Figures 4a, c show that orientation #0 provides near best possible performance (RSSI between 80 and 100, TCP throughput between 800 and 900 Mbps, very low standard deviations) at all locations except Location #7. Location #7 is a rather special case, where the Tx and Rx are placed around the edges of a corner, in a manner that there is no LOS path possible between them. We conclude that *high-throughput 60 GHz links can be established through materials such as drywall or glass* even with wide-beam antennas. Although the signal attenuates when it passes through such materials, Fig. 4b shows that, in the case of optimal antenna orientation, a NLOS link through drywall was able to sustain rates of 1540–3080 Mbps 80% of the

time and a NLOS link through glass was able to sustain a rate of 2310 Mbps 95% of the time.

In contrast, the performance averaged across all orientations is much lower than for Orientation #0; RSSI (Fig. 4a) and TCP throughput (Fig. 4c) never cross their halfway mark (50 or 400 Mbps, respectively). Further, the extremely large standard deviations suggest very large performance variation at a given location for different orientations. In fact, some orientations resulted in zero throughput, not even allowing a connection establishment between the sender and receiver. For example, in the presence of a wall or a corner between the sender and the receiver, non-zero throughput was achieved only at 3 orientations. Even worse, in the case of Location #9 (a relatively long link in a lab filled with “clutter” [4], i.e., objects that do not directly block the LOS between the Tx and the Rx, such as office furniture, soft partitions that do not extend to the ceiling, and lab equipment), Orientation #0 was able to sustain high data rates (1925 Mbps or higher for 85% of the time) and high throughput, but no link was established for any of the remaining 15 orientations. Although [4] found that attenuation due to clutter decreases as we move from 2.5 to 60 GHz, our results show that clutter can have a severe impact on 60 GHz performance, except in the case of very short distances or perfect antenna orientation.

**Remarks** Overall, these results confirm that even wide-beam antennas can provide high-throughput 60 GHz communication at various locations typical of an indoor

**Fig. 4** 802.11ad performance across different locations. **a** RSSI. **b** PHY data rate distribution. **c** TCP throughput



WLAN environment. Further, although communication is possible in the case of antenna array misalignment, either via beamsteering or through an NLOS link via reflection, the performance can be much lower than in the case of optimal orientation. This strong dependence on the relative orientation between the Tx and Rx antenna arrays argues in favor of wide-beam antennas to reduce the beam steering overhead.

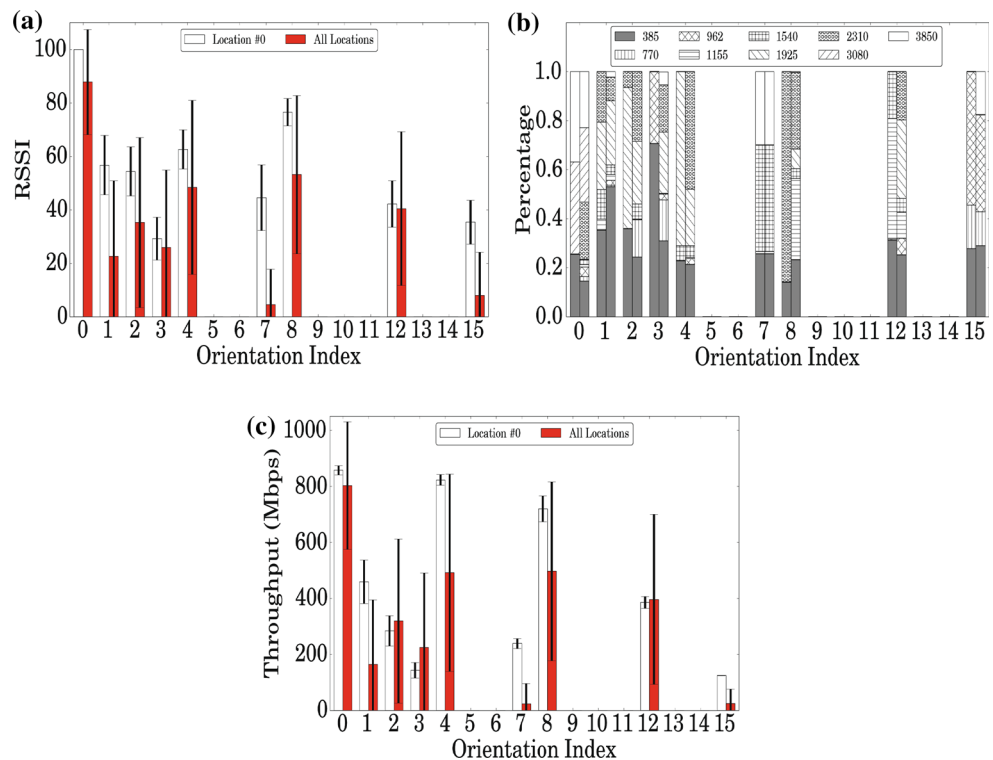
### 4.3 Impact of Tx/Rx orientation

We now take a closer look at the performance of each of the 16 orientations. Figures 5a, b, c plot the the average RSSI, the selected PHY rates, and the average TCP throughput at each orientation, averaged across all locations. We also consider separately Location#0, which represents the best-case scenario (LOS over 8 ft). These figures show that orientation #0, as expected, performs the best among all orientations. Interestingly, orientations #4, #8, and #12, i.e., cases where the Tx points directly towards the Rx location, as in the best scenario, but the Rx is rotated by 90°, 180°, or 270°, give very similar and significantly higher throughput than all other orientations, indicating that the Tx position is more critical to performance. On the other hand, orientations #1, #2, and #3 where the Rx is fixed facing the Tx location and the Tx is

rotated in 90° intervals are characterized by throughputs lower than 450 Mbps, and rather large standard deviations. Figure 5b shows that the PHY data rates are lower for these three orientations, compared to orientation #0, e.g., the lowest rate of 385 Mbps is used up to 70% of the time. Even worse, for any given Tx orientation except the one directly facing the Rx location (#0, #4, #8, #12), all Rx orientations except the one directly facing the Tx location give extremely low or zero performance. E.g., consider orientations #1, #5, #9 and #13, where Tx orientation is fixed, in Figs. 5a, c; among them, only orientation #1 gives non-zero RSSI/throughput. The heavy impact of antenna orientation may initially sound counter-intuitive for Wilocity radios, since they are equipped with electronically steerable antenna arrays. However, practical 802.11ad phased antenna arrays cannot generate homogeneous beams across all directions [21]; this has been recently verified experimentally in [32].

*Orientation/link asymmetry* One additional observation from Figs. 5a, b, c is that, at Location #0 (Hall), orientations #4 and #12, which are symmetric w.r.t the Tx position, do not give similar throughput. Similarly, the PHY data rates used with these two orientations are very different even at Location #0 (Fig. 5b). The same observation can be made about orientations #1 and #3, which are symmetric w.r.t the Rx position. To eliminate the impact of

**Fig. 5** Performance across different orientations. **a** RSSI. **b** PHY data rate distribution. **c** TCP throughput





environmental asymmetry (there are still walls in the Hall although far from the Tx-Rx link, as well as furniture), we looked at the results at Location #2 (a Corridor with walls of the same material on both sides). The result was similar (we omit it here due to space limitation). Further, [32] showed that 60 GHz links exhibit link asymmetry (down-link and uplink throughput are different when the Tx and Rx use different beamwidths). These orientation and link asymmetries make it extremely hard to predict and/or accurately model performance in indoor environments.

**Outdoor experiments** In case of Tx/Rx antenna array misalignment (orientations other than #0), communication can be achieved either via beamsteering which allows for realignment of the main beam or over an NLOS link through reflection. Since the Wilocity radios neither allow us to control the beam steering process nor provide any information about it, we resorted to outdoor experiments in order to obtain a better idea about the factor that allows communication at different orientations. Specifically, we repeated the experiments for all 16 orientations at an outdoor open space devoid of reflective and/or obstructive surfaces and objects. Only four orientations (0, 4, 8 & 12), where the Tx faces the Rx gave non-zero throughput.

**Remarks** The results show that COTS 802.11ad devices can establish Gigabit links even with imperfect Tx/Rx orientation. In indoor environments, the presence of multiple reflective surfaces creates additional opportunities via NLOS links, although the throughput of such NLOS links is typically much lower compared to the throughput of LOS links. We conclude that Tx orientation is more important in determining the performance and the possibility of a connection. On the other hand, orientations where neither of the Tx or Rx antenna point toward the other's location, are not suitable for communication at all. Further, due to the complex indoor environment, symmetric properties cannot be assumed.

#### 4.4 Impact of distance

In Sect. 4.3, we only focused on the relative orientation of Tx and Rx antenna arrays and hence most experiments were conducted with Tx and Rx very close to each other, to avoid the effects of severe signal attenuation of the 60 GHz radio signal. However, for the mmWave technology to be a viable option for building WLANs, it is necessary to examine how channel quality indicators, RSSI and PHY data rate, and the corresponding transport layer throughput, vary with the distance between the transmitter and the receiver. In this section, we study the impact of the Tx/Rx distance in LOS scenarios using both the COTS 802.11ad-compliant Wilocity radios and the proprietary HXI radios.

##### 4.4.1 Wilocity radios

We begin our study with the Wilocity radios. Figure 6 plots the RSSI, the PHY rate distribution, and the TCP throughput over distance at the two main locations.

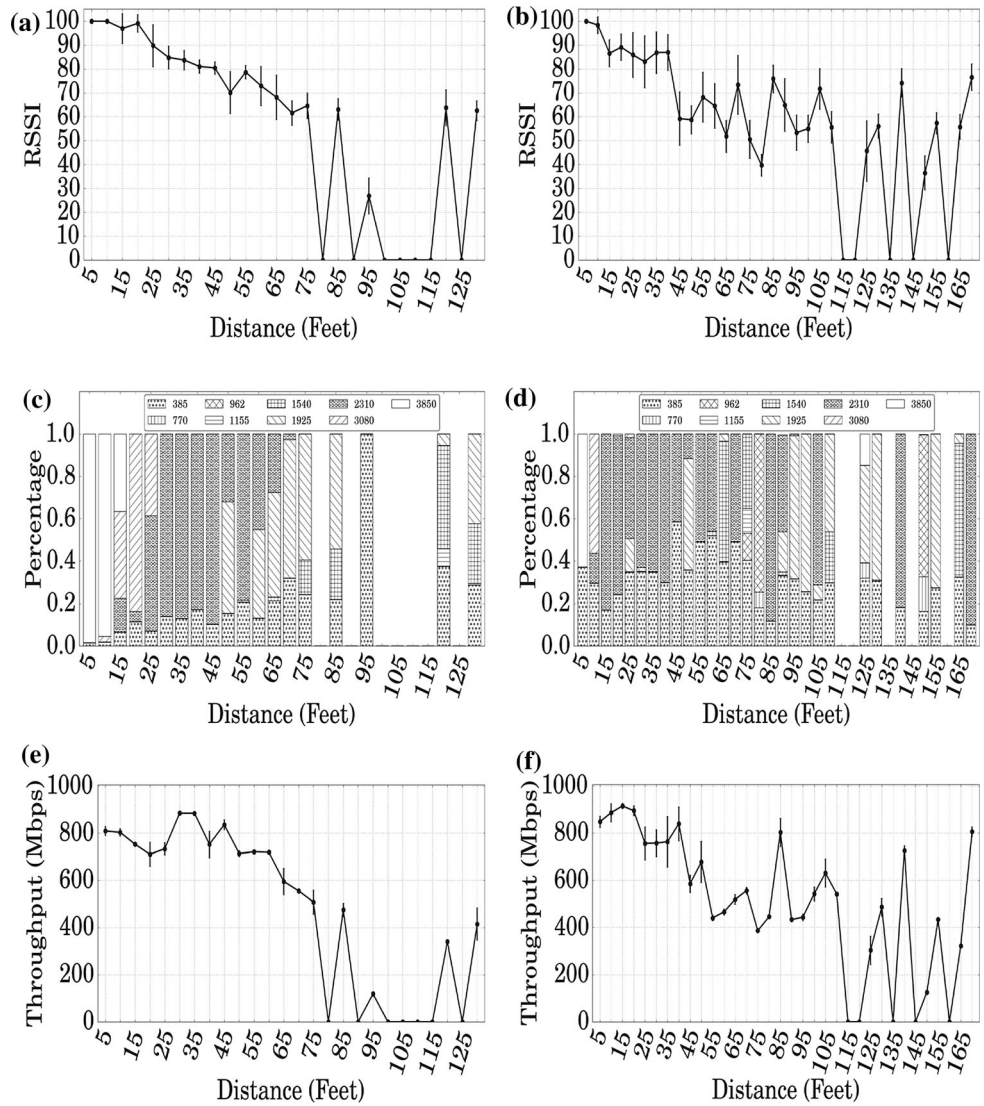
**Range** Fig. 6 shows that long ranges indoor can be achieved even with radios designed for short-range applications, which use relatively wide beams and have lower EIRP than the maximum allowed by FCC. The Corridor measurements show that RSSI exhibits large oscillations (due to a phenomenon known as *waveguide effect* [29]) but does not drop with distance beyond 40 ft (Fig. 6b) and a PHY data rate of 2310 Mbps can be supported at a distance of 170 ft (Fig. 6d). The Hall measurements show a different picture, closer to what one would expect, with RSSI gradually dropping with distance up to 75 ft (Fig. 6a) but even in this case, the link was able to support a rate of 1540 Mbps or 1925 Mbps roughly 70% of the time at a distance of 130 ft (Fig. 6c). These ranges are much longer than the values reported recently with the same hardware (770 Mbps at 72 ft in a data center [46], 385 Mbps at 72 ft and 2310 Mbps at only 33 ft in an outdoor environment [45]).

**RSSI vs. distance** Recent experimental work [11, 44–46] observed that the attenuation of 60 GHz signals with distance follows closely the Friis model in LOS scenarios, both in stable data center and outdoor picocell environments. Figure 6 shows that this assumption does not hold true in the case of wide beams in indoor environments.


In Fig. 6a (Hall), the distance axis can be divided in 3 distinct regions. For distances up to 20 ft, RSSI remains close/equal to 100 and the link can sustain the two highest PHY data rates at least 80% of the time. The next region is between 25 and 75 ft where RSSI decreases with distance. Lastly, distances between 80 and 130 ft are characterized by extremely large RSSI oscillations; RSSI drops to zero at several distances and then rises again, often to high levels. Although we cannot confirm it, we believe *these link outages are the result of multipath*. We also hypothesize that such “dead zones” might have led researchers previously [45, 46] to conclude a much shorter range for the Wilocity radios. It is possible that narrower beams can eliminate dead zones at the cost of higher vulnerability to blockage and mobility [32].

In Fig. 6b (Corridor), we observe 4 distinct zones. RSSI shows a decreasing trend with distance only for very short distances (5–10 ft), remains almost stable for distances of 15–40 ft, exhibits very large variations and non-monotonic behavior but non-zero values for longer distances of 40–110 ft, and finally exhibits “dead zones” at distances longer than 110 ft. For comparison, Fig. 7a plots the RSS (in dBm) over distance for 802.11ac in the Corridor. Although the effects of multipath are still visible, interestingly, signal strength shows a more clear decreasing trend with distance compared

**Fig. 6** 802.11ad performance as a function of distance in two locations. **a** RSSI (Hall). **b** RSSI (Corridor). **c** PHY data rate distribution (Hall). **d** PHY data rate distribution (Corridor). **e** TCP throughput (Hall). **f** TCP throughput (Corridor)



**Fig. 7** 802.11ac performance as a function of distance in the Corridor. **a** RSS. **b** TCP throughput



to Fig. 6b, despite the fact that the 802.11ac cards are equipped with omni-directional antennas.

*PHY data rate vs. distance* Fig. 6c, d show that for most distances there are 2 or 3 dominant data rates, and the lowest rate of 385 Mbps is always used at least 10% and up to 60% of the time, even in the case of very short

distances/high RSSI (with the exception of very short distances in the Hall). *This observation suggests highly time-varying channels and/or inability of the rate adaptation algorithm to converge to a single rate.* In the Hall (Fig. 6c), we still observe a monotonic decrease with distance and RSSI; lower data rates dominate at longer

distances/lower RSSI values. In contrast, there is no such monotonicity in the Corridor (Fig. 6d).

The assumption of the validity of the Friis propagation model (or more generally a log-distance path loss model) in LOS scenarios has led to the use of simple RSS-based rate adaptation algorithms in 802.11ad simulators [11, 35, 44, 46] and the use of RSS as a direct indicator of the PHY data rate [32, 45]. Our results in Fig. 6a, b clearly showed that propagation in indoor WLANs when radios are equipped with wide-beam antennas does not follow the Friis model since RSSI does not decrease monotonically with distance. Since we cannot directly compare the measured propagation characteristics with those of the Friis model due to the fact that our cards report RSSI instead of the actual received signal strength (RSS), we attempt an indirect comparison via the supported PHY data rates. Specifically, for each distance  $R$ , we calculate a theoretical RSS value  $P_{RX}(R)$  (in dBm) based on the commonly used log-distance path loss model adjusted to account for shadowing fading [29] and potential losses due to reflections in case of NLOS links [45].

$$P_{RX}(R) = EIRP + G_{RX} - L_{ploss}(R) + X_{\Omega} - L_{margin} \quad (1)$$

$$L_{ploss}(R) = 10 \log_{10} \frac{16\pi^2 R^2}{\lambda^n} \quad (2)$$

where  $\lambda$  is the wavelength,  $EIRP$  is equal to  $23\text{dBm}$  for Wilocity Radios [45],  $G_{RX} = 10 \log_{10} N_{RX}$  is the receiver antenna gain (in dBi) as a function of the antenna elements  $N_{RX} = 16$  [45],  $n$  is the path loss exponent (we use different values for Corridor and Hall based on [29]) and  $X_{\Omega}$  represents a shadowing component (zero mean Gaussian random variable with standard deviation values also obtained from [29] for different environments). In [45],  $L_{margin}$  is taken equal to 15 dB although most materials typically lead to 6–7 dB loss. In our case, we consider three different values: 0, 7, and 15 dB. We then use the rate-sensitivity table for 802.11ad (Table 2 in [32]) to convert  $P_{RX}$  to a PHY data rate.

Figures 8a, b compare the measured dominant rates against the theoretically computed rates from (1). We observe that the conservative models which account for reflection losses significantly underestimate the data rate; if we assume a 15/7 dB loss, only the control data rate (27.5) Mbps can be supported for distances longer than 30/65 ft. On the other hand, assuming zero loss due to reflections results in overestimation of the data rate for short distances (up to 30 ft) and underestimation for long distances in both environments, potentially due to a combination of multipath and waveguide effects. Overall, we observe that *PHY data rate cannot be predicted from simple propagation models in indoor WLAN settings*. For comparison, Fig. 8c plots the measured dominant rates for 802.11ac in the

Corridor. Although there is a decreasing trend with distance indicating absence of waveguide effects, the dominant rate oscillates a lot in the range of 20–120 ft and shows no correlation with distance, as expected in indoor environments.

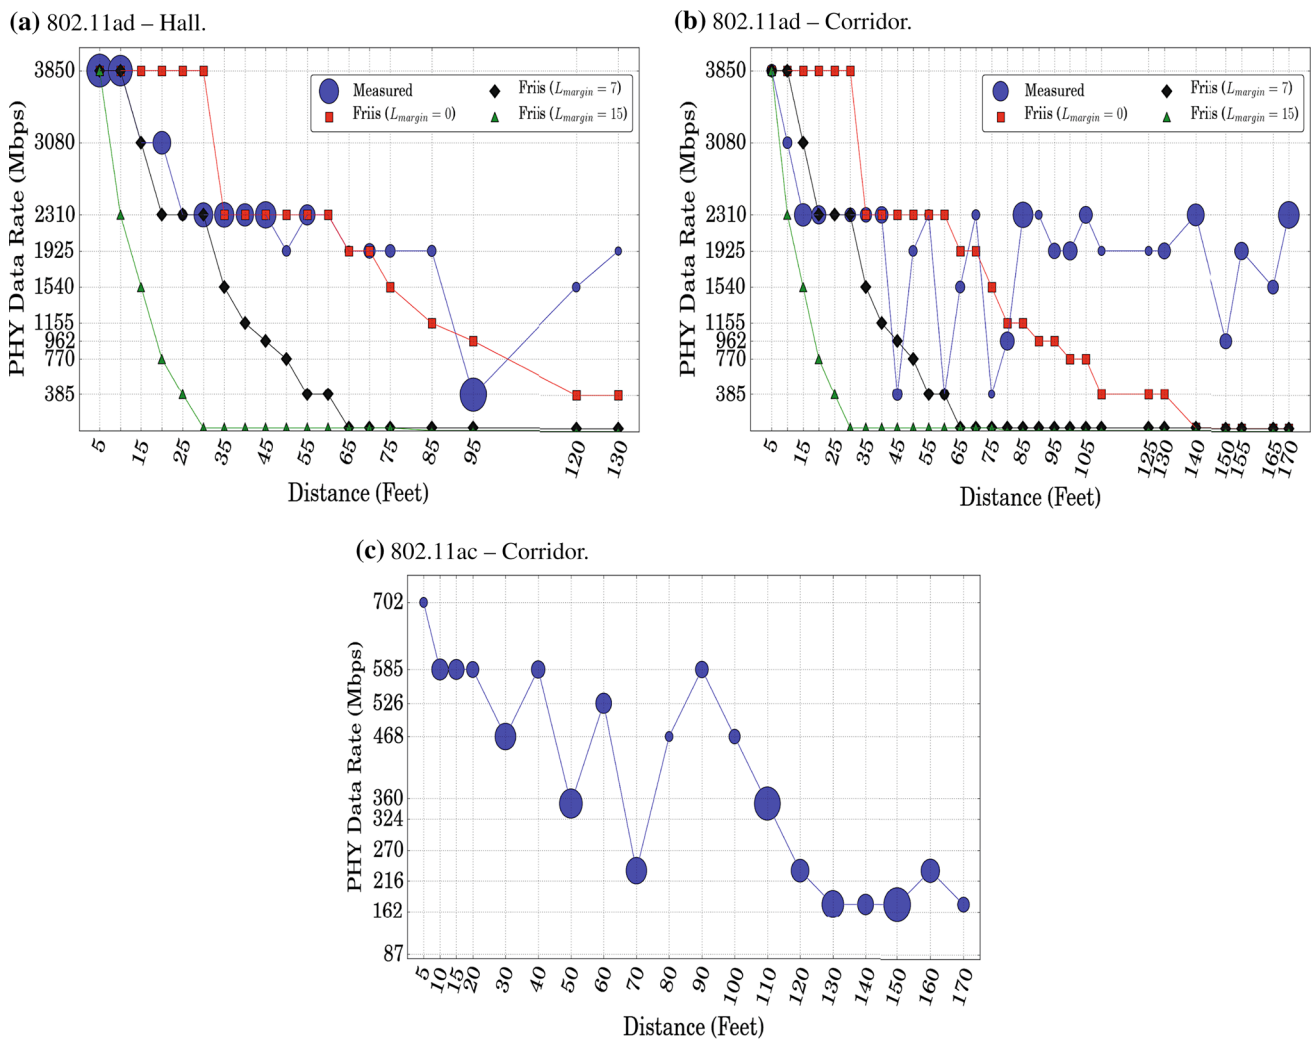
*Throughput vs. distance* Figs. 6e, f show again distinct regions although these regions do not always overlap with the RSSI regions. In the Hall (Fig. 6e), throughput sustains high values (above 800 Mbps) for distances up to 45 ft although RSSI starts dropping at 25 ft. It then exhibits a gradual drop up to a distance of 75 ft (boundary of the second RSSI region) and “dead zones” for longer distances. In the Corridor (Fig. 6f), we observe two small regions of high values (above 800 Mbps at 5–20 ft, around 800 Mbps at 25–40 ft), very large variations with distance (up to 400 Mbps within 5 ft) for distances up to 110 ft, and “dead zones” for longer distances. Overall, we observe *a weak correlation of throughput with distance for short/intermediate distances and no correlation for longer distances*.

Figure 7b plots the TCP throughput over distance for 802.11ac in the Corridor. Similar to RSSI, we observe again a stronger correlation with distance compared to 802.11ad and almost zero standard deviations for a given distance. However, note that TCP throughput is affected more in legacy WiFi than in 802.11ad—it never exceeds 200 Mbps in Fig. 7b although the dominant rates can be as high as 702 Mbps (Fig. 8c), probably due to contention from other WiFi networks.

**Remarks** The results in this section show that signal strength does not drop monotonically with distance in the case of wide-beam antennas in typical indoor WLAN environments, due to the presence of strong multipath and, in some cases, waveguide effects. In fact, in certain environments, the combined impact of these two phenomena is stronger than in the case of legacy WiFi. Hence, in contrast to observations made by previous works in indoor data center or outdoor picocell environments, propagation in typical indoor WLAN environments cannot be described by simple propagation models, and new models are needed for 802.11ad simulators. Similarly, the rate adaptation logic cannot converge to a single rate most of the time even in the case of high RSSI, indicating a weak (if any) correlation between the two metrics. Both PHY data rate and TCP throughput show a weak or no correlation with distance.

#### 4.4.2 HXI radios

Since the Wilocity radios do not allow us to fix the PHY data rate or the beam properties, we cannot pinpoint the exact causes of the observed behavior in Sect. 4.4.1 (for example, we do not know if we always have a direct LOS link, or in some cases low BER results in beam steering



**Fig. 8** Comparison of the measured dominant rate versus the theoretically calculated rate in the Hall (a) and corridor (b). The dominant rate for 802.11ac in the Corridor is included for comparison (c). A larger circle indicates larger dominance

and indirect communication via reflections). Additionally, since they report RSSI instead of RSS, as we mentioned in Sect. 4.4.1, we cannot directly compare the experimental results with theoretical propagation models (and instead, we resort to the indirect comparison, via the PHY data rate, in Fig. 8). In this section, we repeat the experiments from Sect. 4.4.1 using HXI radios equipped with horn antennas, in both symmetric and asymmetric configurations.

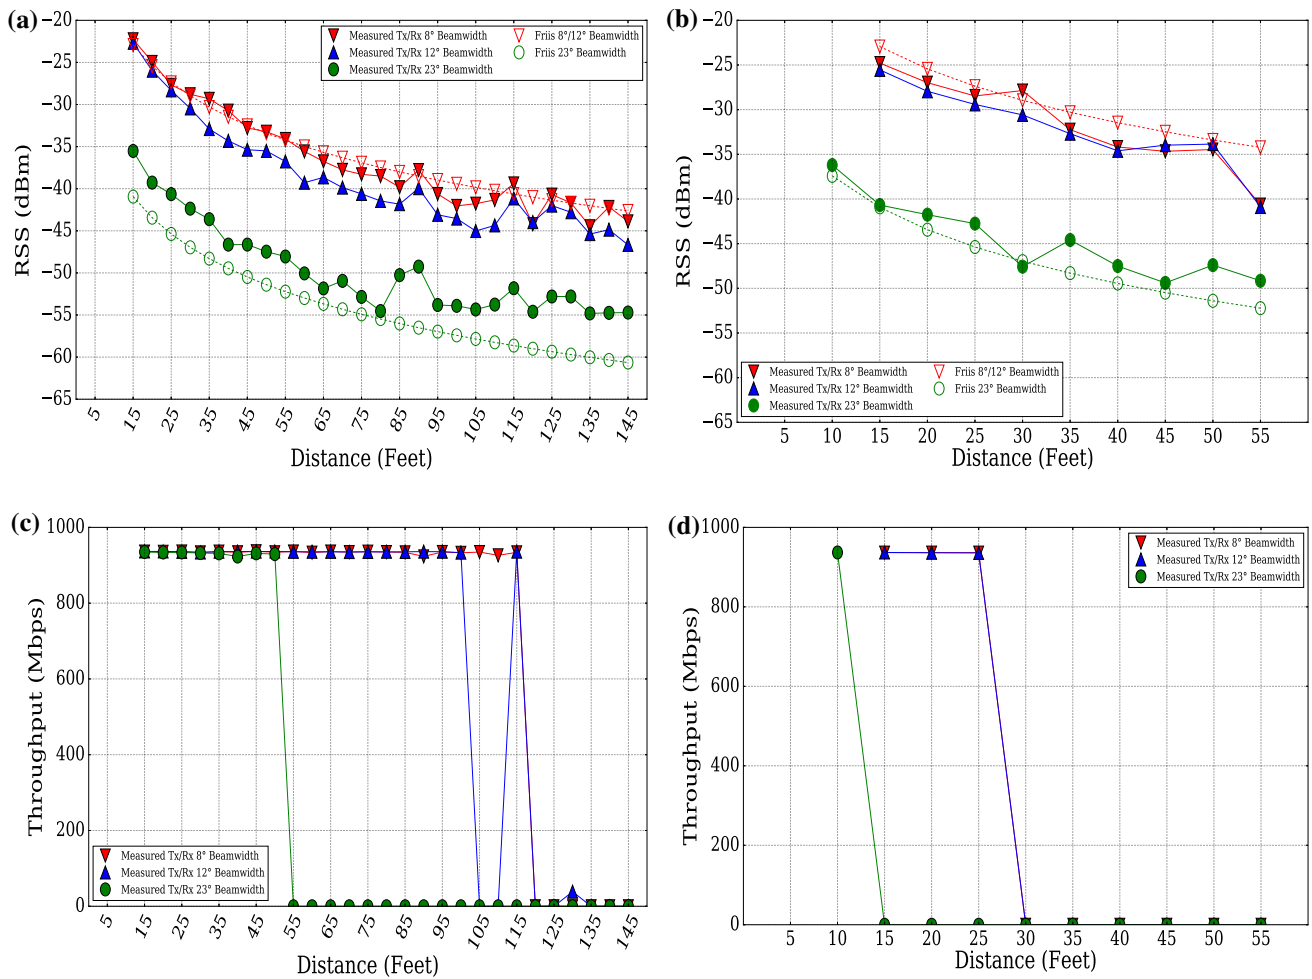
We begin with symmetric configurations. Figures 9a–d plot the RSS and TCP throughput in the two locations with each of the three horn antennas. Since HXI radios report RSS (in dBm), we also plot the theoretical received RSS using the Friis equation (equation (1) with no shadowing component and  $L_{margin} = 0$ ).

Figure 9a, b show a very different picture compared to the one we observed in Fig. 6a, b. First, RSS in Fig. 9a, b exhibits a gradual decrease with distance, and shows no “dead zones” and no large oscillations, unlike in Fig. 6a, b.

In particular, the measured RSS values with the narrowest beam ( $8^\circ$ ) antenna in the Hall follow very closely the theoretical values from the Friis model. As expected, the distance between measured and theoretical values increases with wider beams in both locations. Note though that we still observe the impact of multipath propagation (in the form of non-monotonic RSS drop with distance) for longer distances; again this impact is more pronounced for wider antennas. Note also that the non-monotonic behavior starts at longer distances in the open Hall than in the Corridor. Simple geometric calculations show that, in the narrow Corridor, the signal will “hit” the walls (potentially resulting in multipath fading), at a distance of 36, 24, and 12 ft, in the case of a  $8^\circ$ ,  $12^\circ$ , and  $23^\circ$  beam, respectively.

Similarly, Figs. 9c, d show a very different picture compared to the one we observed in Figs. 6e, f. We observe two striking differences: (i) throughput values are either 900 Mbps or 0; there are no intermediate values, (ii)





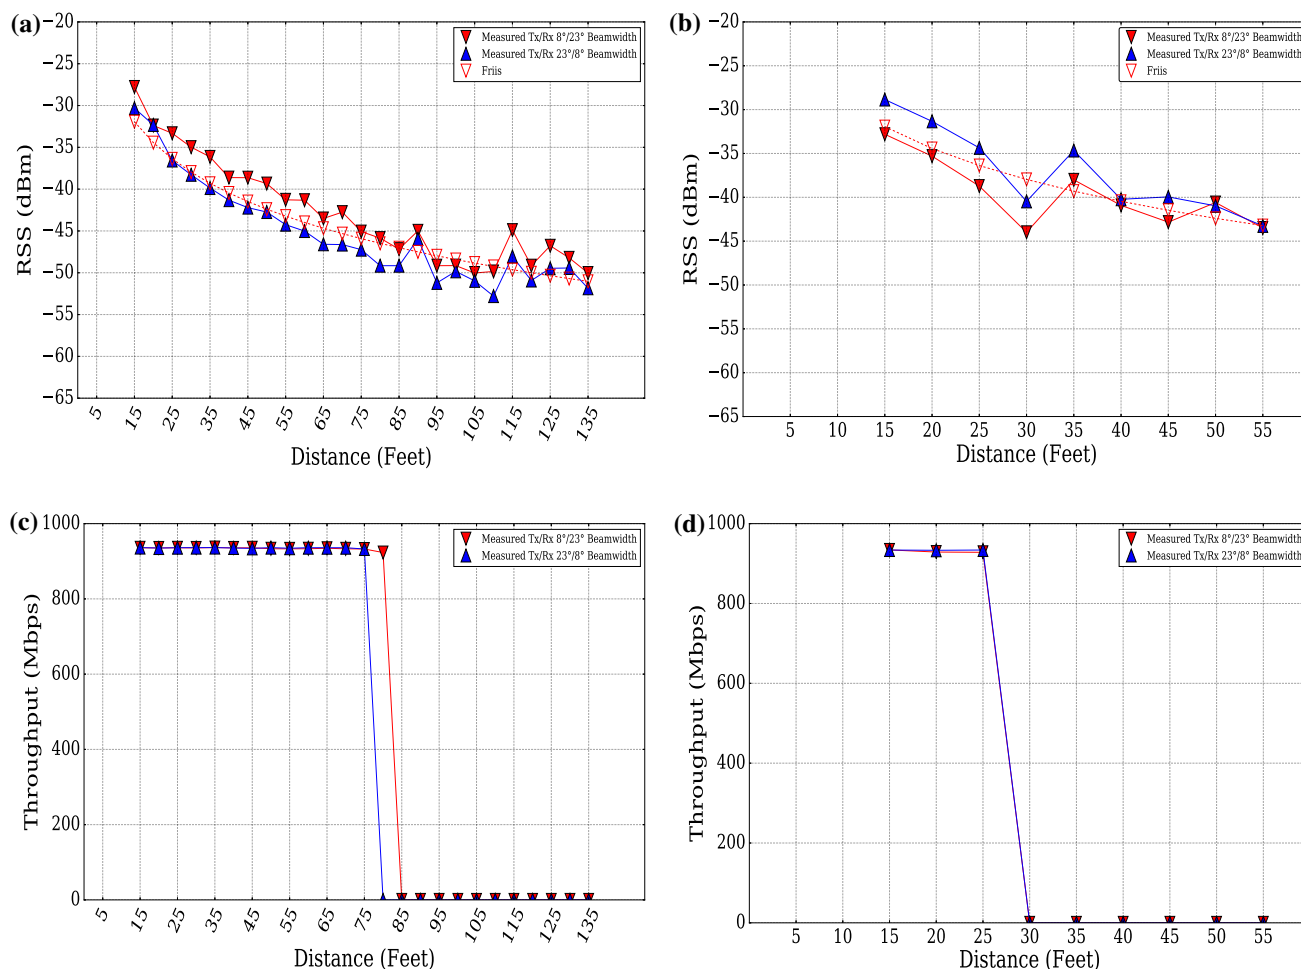
**Fig. 9** 60 GHz performance as a function of distance in two locations using the HXI radios—symmetric case. **a** RSS (Hall). **b** RSS (Corridor). **c** TCP throughput (Hall). **d** TCP throughput (Corridor)

the range is much shorter than in the case of Wilocity radios. Additionally, we observe no oscillations or dead zones (with one exception in Fig. 9c with the 12° beamwidth); throughput remains constant at 900 Mbps up to a certain range (different for different beamwidths) and sharply drops to zero at a distance 5 ft longer. In the Hall, the non-zero throughput range is 115 ft for the 8° and 12° beam antennas, and 50 ft for the 23° beam antenna. In the Corridor, the range is even shorter; only 25 ft for the for the 8° and 12° beam antennas, and only 10 ft for the 23° beam antenna. In our communication with HXI engineers, multipath propagation along with the simple PHY layer were identified as the causes of this behavior. Note again that these radios were initially designed for LOS communication between static nodes in outdoor open spaces and they were equipped with very narrow beam antennas (0.9°). However, such narrow beamwidths are not feasible in scenarios involving mobile clients.

We now consider asymmetric configurations, with one of the two nodes equipped with a 23° beam antenna and the

other one with an 8° beam antenna. In real scenarios, it is likely that APs may be able to accommodate a much larger number of antenna elements compared to pocket-size mobile devices, and be able to support much narrower beams. The results for the RSS and TCP throughput in the two locations for two configurations (narrow beam Tx/wide beam Rx, denoted as  $Tx/Rx$  8°/23°, and wide beam Tx/narrow beam Rx, denoted as  $Tx/Rx$  23°/8°) are shown in Fig. 10a–d. Again, we also plot the theoretical RSS values from the Friis model.

The results for both RSS and throughput are qualitatively similar to those in Fig. 9. The non-zero throughput range falls between the range achieved with the symmetric 8° and 23° beamwidth configurations in the Hall (Figs. 10c, 9c) and is the same as the 23° beamwidth configuration in the Corridor (Figs. 10d, 9d). Interestingly, it is not clear from the results in Figs. 10a–d whether one of the two configurations is better than the other. In the Hall,  $Tx/Rx$  8°/23° gives a consistently higher RSS for a given distance by roughly 5 dB compared to  $Tx/Rx$  23°/8°



**Fig. 10** 60 GHz performance as a function of distance in two locations using the HXI radios—asymmetric case. **a** RSS (Hall). **b** RSS (Corridor). **c** TCP throughput (Hall). **d** TCP throughput (Corridor)

(Fig. 10a), which is also translated in a longer non-zero throughput range by 5 ft (Fig. 10c). In the Corridor, RSS is higher by about 5 dB with  $T_x/R_x$  23°/8° for up to 35 ft and similar with both configurations for longer distances (Fig. 10b), while the non-zero throughput range is the same with both configurations (Fig. 10d).

**Remarks** Overall, the experiments with the HXI radios confirm our findings in Sect. 4.4.1; the assumption of a monotonic drop of signal strength with distance, based on the Friis model, is only valid in open spaces with nodes equipped with narrow-beam antennas. In contrast, modeling signal propagation becomes much more challenging in typical indoor WLAN scenarios, where nodes will potentially use wider beams to cope with client mobility, due to multipath.

#### 4.5 Relationship among metrics from different layers

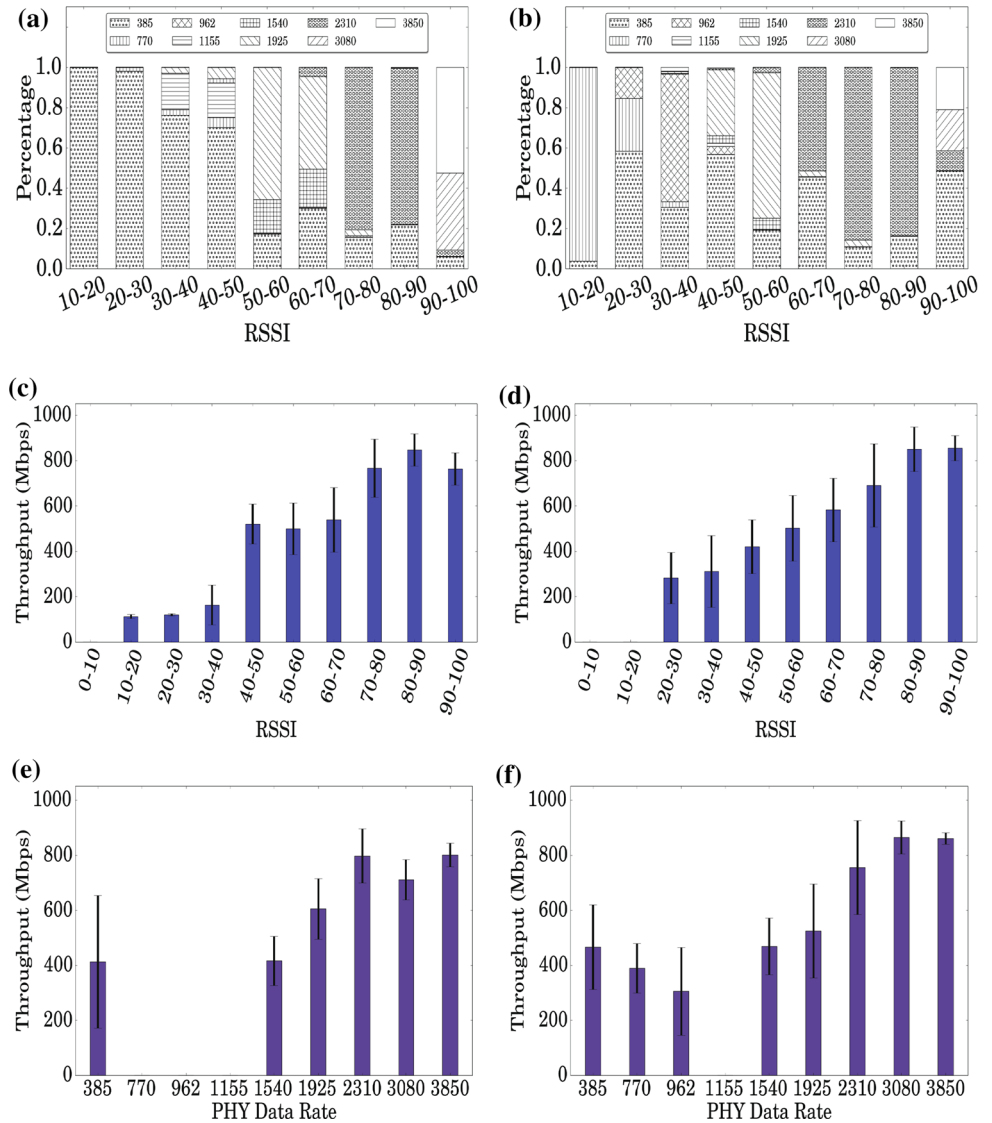
In this section, we take a closer look at the three metrics—RSSI, PHY data rate, and TCP throughput—and

investigate whether one of them can be used as a strong indicator of the other. In particular, we examine whether (i) RSSI can predict PHY data rate and/or TCP throughput and (ii) PHY data rate can predict TCP throughput.

*PHY data rate, TCP throughput vs. RSSI* Since RSSI varies during a 10 s iperf3 session, we had to consider a finer time granularity. We divided each session in 100 ms intervals and selected only those intervals where a particular RSSI value was observed at least 90% of the time. We then grouped the dominant RSSI values observed in the selected intervals in 10-unit bins. Figures 11a, c and b, d plot the PHY data rate distribution and the average TCP throughput over RSSI in the Hall and Corridor, respectively.

Figures 11a, b show that RSSI can serve as a weak indicator of PHY data rate at a given location; for most RSSI values, there is a dominant data rate appearing more than 60% of the time. In the Hall, we also observe a monotonic relationship between the two metrics—higher dominant data rates for higher RSSI values. However, for the same RSSI bin, the observed data rates can be very

**Fig. 11** Relationship among RSSI, PHY data rate, and throughput in two locations. **a** PHY data rate versus RSSI (Hall). **b** PHY data rate versus RSSI (Corridor). **c** TCP throughput versus RSSI (Hall). **d** TCP throughput versus RSSI (Corridor). **e** TCP throughput versus PHY data rate (Hall). **f** TCP throughput versus PHY data rate (Corridor)



different at the two locations. For example, for low RSSI (20–30), the data rate remains constant at 385 Mbps in the Hall but can take the values of 770 Mbps or 962 Mbps 40% of the time in the Corridor. As another example, for very strong RSSI (90–100), the data rate in the Corridor takes its lowest value (385 Mbps) 50% of the time.

Figure 11c shows that RSSI can serve as a reliable although coarse-grained indicator of throughput in the Hall. We clearly distinguish 3 regions—high throughput region (600–900 Mbps for RSSI higher than 70), medium throughput region (400–700 Mbps for RSSI between 40 and 70), and low throughput region (0–300 Mbps for RSSI lower than 40). The picture is very different in the Corridor (Fig. 11d). Instead of distinct regions, here we observe a monotonic increase of the average throughput with RSSI. However, the standard deviations are very large (100–200

Mbps) except in the case of very high RSSI values. We also observe that in the low RSSI region, for the same RSSI, TCP throughput can be very different in the two locations, making prediction difficult across locations.

*TCP throughput vs. PHY data rate* Similar to RSSI, the PHY data rate varies during a 10 s session. Hence, we used a similar methodology to investigate the relationship between PHY data rate and TCP throughput. We selected only the 100 ms intervals where a particular data rate was reported at least 90% of the time. Figures 11e, f plot the average TCP throughput over the PHY data rate in the Hall and Corridor, respectively. A first observation from these figures is that some data rates are never selected consistently over a 100 ms period. The two highest data rates result in high throughput values and low standard deviations in both locations. However, for the remaining data

rates, throughput varies significantly with standard deviations often higher than 200 Mbps.<sup>2</sup> Further, in the Corridor, several data rates have overlapping throughput ranges. Overall, the PHY data rate cannot serve always as a good indicator of TCP throughput.

**Remarks** Our results show that RSSI can serve as a weak indicator of PHY data rate and TCP throughput only at certain locations, but not across locations. Further, PHY data rate is not always a good indicator of TCP throughput. These observations have two immediate implications: (i) Translating signal strength to PHY data rate or PHY data rate to higher layer performance, a common practice in recent measurement studies [32, 45], can yield inaccurate results in typical indoor WLAN environments when nodes are equipped with wide-beam antennas. (ii) Simple RSS-based rate adaptation algorithms, which have been used in recent simulation studies [11, 35, 44, 46], may not work well in indoor WLANs.

#### 4.6 Impact of human blockage

The inability of 60 GHz links to pass through human body without suffering severe attenuation is often described as one of the major challenges for the mmWave technology. A human in the LOS between the transmitter and the receiver can attenuate the signal by 20–30 dB [28], resulting in link outage. [45] found that in outdoor picocell settings the impact of static pedestrians is limited in a very small area around the user due to the base station height (6 m). However, the impact of groups of moving pedestrians becomes heavier. Recent studies in indoor settings [32, 34, 35] showed that human blockage remains a major challenge, and becomes worse due to the long re-connection times of existing 802.11ad hardware. We investigate the severity of this problem through two sets of experiments.

##### 4.6.1 Controlled experiment

We performed experiments in the Hall and Corridor with controlled human placement and motion. We study the blockage caused by both permanent human obstruction of the LOS link and transient motion which disrupts the link temporarily. We are primarily concerned with measuring the *reconnection time* and *throughput degradation*.

At each location, we tried four different Tx-Rx distances (8'6", 16'6", 24'6", and 32'6") and considered two kinds of blockage: mobile and static. In the former, a person walks in random fashion along the LOS path between the

Tx and Rx. If the link breaks as a result of such motion, the human moves away from the link to allow it to recover. In the latter, a person stands permanently between the Tx and Rx, forcing the Tx to find an alternate NLOS path. In cases where the link does not break in spite of the blockage, we measured the throughput degradation caused due to human presence. Table 2 summarizes our findings, plotting the average values over 5 experiments. A zero reconnection time indicates that the link did not break; in that case, the value in brackets indicates the observed TCP throughput. An infinity ( $\infty$ ) value means that the Tx was unable to re-establish the link even several minutes after introducing the blockage.

In the Corridor, the link was resilient to transient human obstruction for distances up to 24'6". For a distance of 8'6", the link did not break even in the presence of permanent human blockage. Note though that throughput, which was always higher than 800 Mbps before the appearance of the human, dropped significantly in some cases. For larger distances, the Tx failed to find an alternative path in the presence of permanent human blockage. Note that in order to evaluate the worst case, the person in most cases stood very close to the Tx or Rx making it harder for the Tx to search for a new path. On the other hand, human blockage introduced mid-way between the Tx and Rx did not break the link. For even greater distance (32'6"), the link failed under human mobility and it took up to 16.6 seconds for it to be established again.

The results in the Hall are very different. The link broke under human mobility for all distances and reconnection times were in the range of 15.34 to 15.67 seconds, similar to the Corridor case, and much longer than the values reported in [35, 45]. The absence of any reflective surface in the immediate vicinity of the Tx made it harder for Tx to quickly find an alternative path when it was blocked momentarily. In case of static blockage, the Tx failed to find a NLOS path in all cases.

**Remarks** These results indicate that a static human near the Tx or Rx presents a much bigger challenge compared to transient blockages introduced by human motion. Nonetheless, the reconnection times are extremely high when considered in the context of multi-Gbps throughput.

**Table 2** Reconnection time and throughput in case of human blockage (M:mobile, S:Static)

	8'6"		16'6"		24'6"		32'6"	
	M	S	M	S	M	S	M	S
Corridor	0	0	0	$\infty$	0	$\infty$	16.61	$\infty$
	[573]	[596]	[728]		[441]			
Hall	15.34	$\infty$	16.07	$\infty$	15.42	$\infty$	16.67	$\infty$

<sup>2</sup> Note that the throughput corresponding to 385 Mbps data rate in Figs. 11e, f is higher (400/450 Mbps) since higher rates were used in the remaining 10% of the time.



One potential way to address this problem is beam dilation, although [32] showed that it only works under high SNR scenarios. Hence, there is a need for faster, more efficient rebeamforming algorithms, potentially combined with mechanisms that distinguish the cause of link outage (human blockage or client mobility), as different approaches work better in each scenario [32].

#### 4.6.2 In the wild experiment

The controlled experiments gave an insight into the challenges arising out of the presence of humans into the environment. However, two questions remain unanswered: (i) how often would such blockages occur in a typical WLAN and (ii) can serving a client with multiple APs (similar to the BS picocloud scenario in [45] for 60 GHz outdoor picocells) help mitigate this problem? To answer these questions, we used a methodology similar to that of [45] since Wilocity radios do not allow switching between APs on-the-fly. We deployed three links in our lab, with the three receivers very close to each other, emulating a single client which can potentially connect to any of the three docks/APs (Fig. 12a), for a period of 15 h. The 15 h experiment period included both night hours and day hours of the following day. We recorded per second TCP throughput for each of the three links.

Figures 12b, c, and d present the CDF of throughput in three cases when one, two, or three APs are considered to be deployed. In the first case, where we assume that only one of the 3 APs was available for connection to the client, we see

that each of the links was blocked/disconnected (zero throughput) for less than 5% of the time and two of the links maintained a throughput between 600 and 700 Mbps most (around 70%) of the time. However, throughput above 800 Mbps was achieved for less than 5% of time by each link.

When considering 2 APs, we have 3 possible combination of APs. Further, for each combination, we plot both the best throughput achieved out of the two APs and the worst one for comparison. Interestingly, all the three combinations gave a 0 % of disconnection time, when considering best throughput scenario, indicating that two APs would have been sufficient for maintaining 100% uptime.

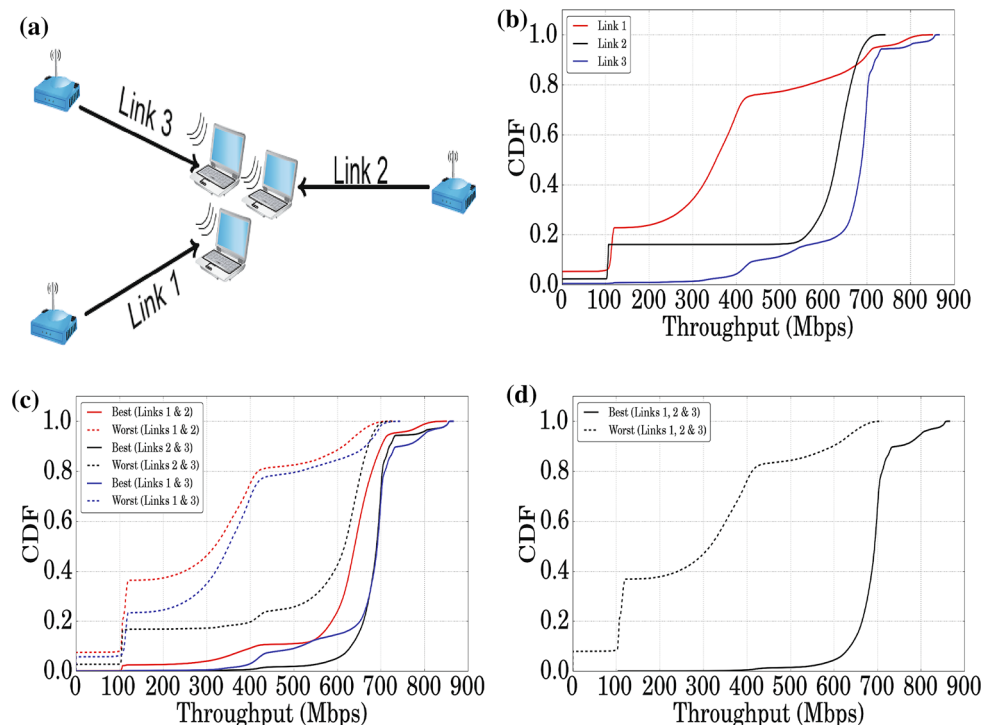
For the 3-AP case, we show the best and worst throughput CDFs. If a client were to connect to the best AP all the time, it would never experience disconnection and would maintain a median throughput of around 680 Mbps.

**Remarks** Our 15 h experiment showed that the presence of humans in a typical office environment does not have a significant impact on connectivity but can cause a significant throughput degradation. However, using 2 APs per client can provide 100% uptime and high throughput.

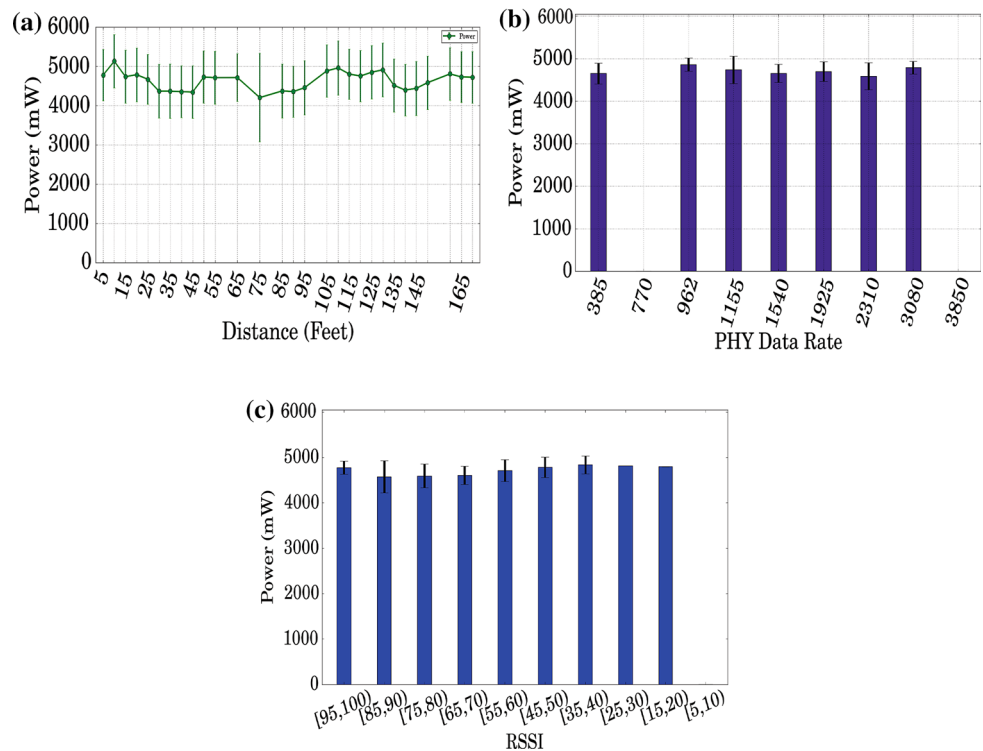
## 5 Power consumption measurements

We now study the power consumption of an 802.11ad NIC. All the measurements are performed in the Corridor. We begin with the the power in various non-communicating

**Fig. 12** In the wild experiment: Topology (a) and CDFs of TCP throughput over 15 h (b, c, d). **a** Topology. **b** 1 AP. **c** 2 APs. **d** 3 APs



**Fig. 13** 802.11ad Rx power consumption as a function of distance, PHY data rate, and RSSI. The error bars show the standard deviation. **a** Rx power versus distance. **b** Rx power versus PHY data rate. **c** Rx power versus RSSI



states in 5.1. We then examine the Rx power consumption (5.2), the impact of client motion (5.3), and the power consumption of beam steering (5.4).

### 5.1 Power in non-communicating states

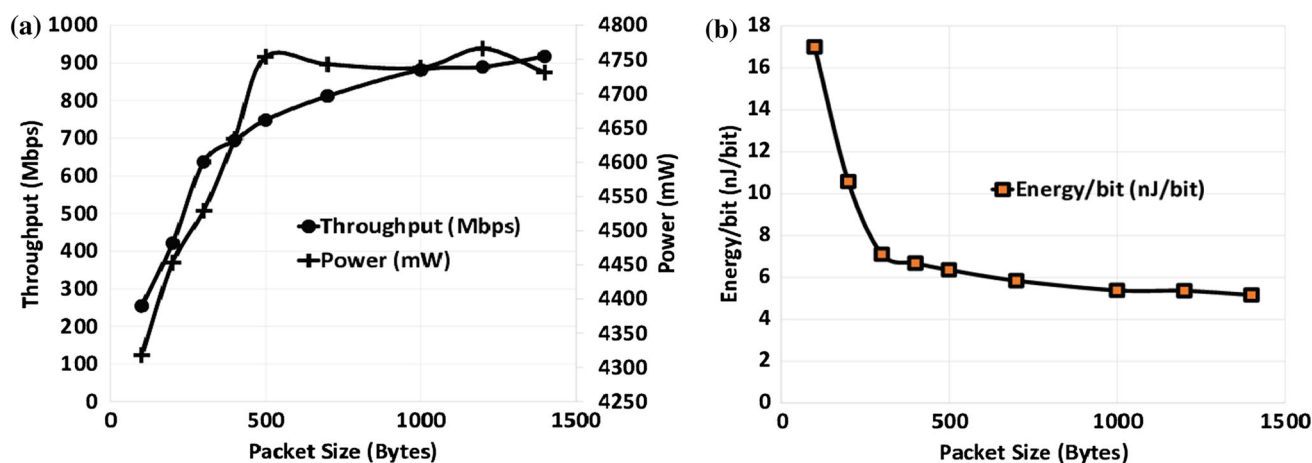
When the card is not connected, we distinguish two states: *Not connected/idle* and *Not connected/scan*; in the latter, the card is actively scanning for 802.11ad APs. The *Not connected/idle* state is the lowest power state (0.5 W). This is the minimum power that needs to be supplied to keep the card powered on. On the other hand, the scanning state consumes more than 2.5 W. Finally, in the *Connected/idle* state, the card is associated to the Dock but there is no Rx or Tx activity. Here the power is around 2.3 W.<sup>3</sup> For comparison, Halperin et al. [10] reported a power of only 820–1450 mW for an 802.11n WNIC in the *idle/connected* state, depending on the number of active antennas (1–3), and Zeng et al. [41] reported a power of 894–1196 mW for a  $3 \times 3$  802.11ac WNIC in the same state, depending on the channel width (20–80 MHz). As smartphones become the next target of 802.11ad, the high idle power consumption may become a major concern, calling for efficient power management schemes.

<sup>3</sup> We occasionally observed this value to vary between 3.5 and 4 W, e.g., just after re-connection. We believe that this is an energy bug in the chipset that leaves it in a high-power state after certain specific events.

### 5.2 Rx power consumption

*Distance, PHY data rate, and RSSI* Recall from Fig. 6f that throughput in the Corridor shows significant variations over distance. Given that WiFi power consumption is proportional to throughput [12, 31], one would expect a similar trend for power consumption over distance. Nonetheless, 802.11ad Rx power exhibits a very different trend over distance in Fig. 13a. The average value remains relatively stable over distance, in the range 4.5–5 W (2.2–2.7 W higher than in the *idle/connected* state), with standard deviations of around 0.5 W. In Figs. 13b, c, we plot the Rx power as a function of the PHY data rate and the RSSI. Similar to Sect. 4.5, we only considered 100 ms intervals where the PHY data rate or RSSI retained the same value at least 90% of the time. Again, although both these factors affect throughput significantly (Figs. 11d, f), their impact on power consumption is minimal.

*Packet size* To study the impact of packet size, we took measurements with different packet sizes while keeping the Tx close to the Rx to ensure that the card uses the highest PHY data rate. Figure 14a plots the power and throughput as a function of the packet size and Fig. 14b plots the energy per bit (in nJ/bit) as the average power consumption ( $W=J/a$ ) divided by the throughput (Mbps). We observe that packet size has a minimal impact on power consumption; as it increases from 100 to 1400 bytes, power consumption increases from 4300 to 4730 mW (10%). On the other hand, the large throughput improvement (from



**Fig. 14** Impact of packet size. **a** Throughput, power consumption. **b** Energy per bit

250 to 920 Mbps—3.68x) results in a significant reduction in the energy cost per bit, from 17.2 nJ/bit to 5.1 nJ/bit.

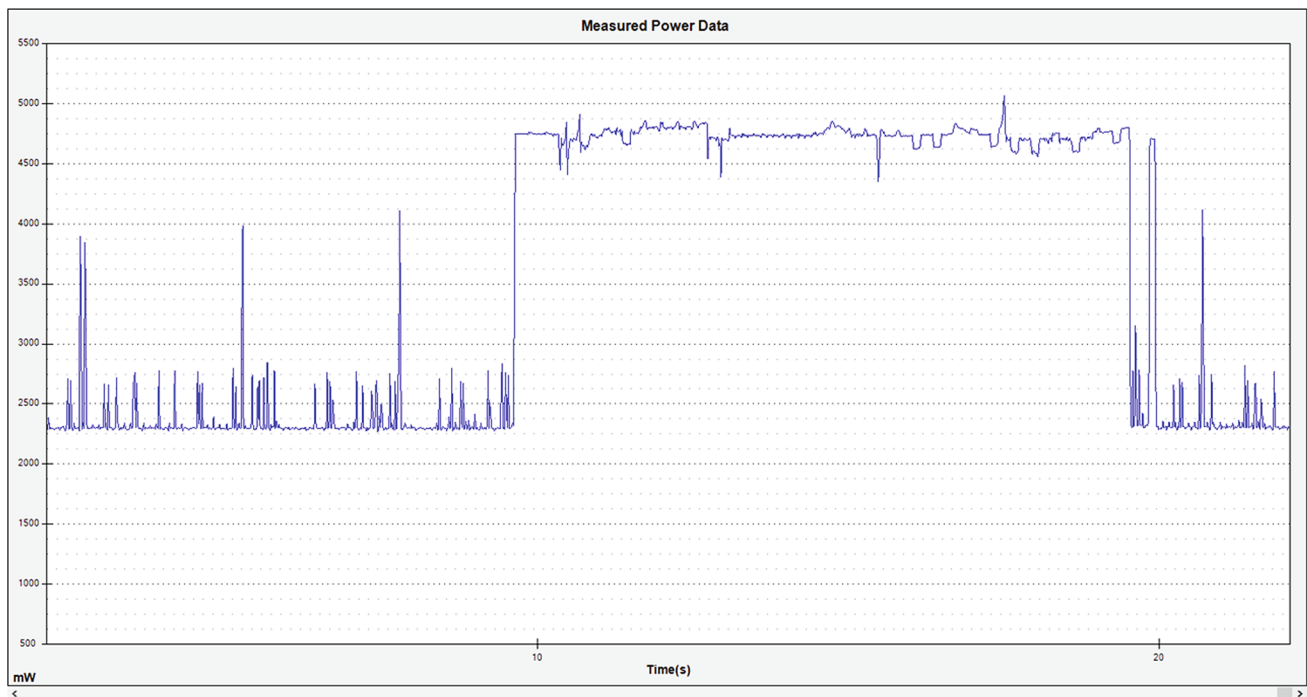
**Comparison with WiFi** Halperin et al. [10] reported an Rx power of 940–1600 mW and a Tx power of 1280–2100 mW for an 802.11n WNIC, depending on the number of active antennas (1–3). Zeng et al. [41] reported similar values (900–1500 mW) for the Tx power of a  $3 \times 3$  802.11ac WNIC, depending on the channel width (20–80 MHz) and the source data rate. Our measurements show that 802.11ad is much more power hungry; its average Rx power consumption (4700 mW) is 123–422% higher than the Rx and even the Tx power of legacy WiFi NICs. While this raises concerns for the viability of 802.11ad in power constrained mobile devices, the per bit energy cost shows a completely different picture. Specifically, [10] reports minimum energy costs (dividing power by the bitrate instead of the achieved throughput) from 4 to 200 nJ/bit for different bitrates and MIMO configurations, with higher values corresponding to lower bitrates. Using a similar methodology, the energy cost of 802.11ad varies from 1.22 nJ/bit (at 3850 Mbps) to 12.2 nJ/bit (at 385 Mbps). The benefit is more prominent at low data rates, where the per bit energy cost is an order of magnitude higher with 802.11n than with 802.11ad. Using the same methodology and the 802.11n Rx power values from [10] (940 mW for MCS0, 1 spatial stream and 1600 mW for MCS9, 3 spatial streams), we estimate an energy cost of 1.23–28.9 nJ/bit for 802.11ac with an 80 MHz channel width. Hence, in theory, 802.11ac can be as energy efficient as 802.11ad when both use their highest data rates. However, the highest data rates of 802.11ac in combination with large channel widths can only be used in very short distances ([41] observed that MCS 8 and 9 yield zero throughput with 80 MHz and 3 spatial streams for distances higher than 33 ft). In contrast,

802.11ad can achieve throughputs higher than 400 Mbps even at 100 ft (Fig. 6f).


Another benefit of 802.11ad becomes clear when we look at small packet sizes. Halperin et al. [10] found that the energy efficiency of 802.11n drops significantly for small packet sizes; they report energy costs of 40–100 nJ/bit for 100 byte packets. In contrast, the energy cost for 802.11ad with 100 byte packets is only 17.2 nJ/bit.

### 5.3 Client motion

We now evaluate the impact of client motion on power consumption. We considered three types of motion: moving towards the dock, moving away from the dock, and moving perpendicularly to the dock. In each case, the client moved for 10 s at walking speed while the dock was sending UDP traffic at full speed. We repeated each experiment 5 times. The results are shown in Fig. 15. With all 3 types of motion, we observed that, although power consumption exhibited higher oscillations than in Fig. 13a (especially in the case of the client away from or perpendicular to the dock), the average power consumption over the 10 s period remained the same as in the static case ( $\sim 4700$  mW). Since the Wilocity radios are equipped with wide-beam antenna arrays, it is possible that low speed motion does not result in large misalignment and either the beam steering process can quickly realign the antennas or beam steering is not triggered at all and rate adaptation deals with such small misalignments. Our conjecture is supported by the fact that throughput remained high (720–911 Mbps) in all experiments with all three types of motion. It is likely that the impact of motion will be higher in the case of narrow-beam antennas (802.11ad supports antenna beams as narrow as  $2.86^\circ$ ).



(a)



(b)


**Fig. 15** Impact of client motion. **a** Client moving towards the dock. **b** Client moving away from the dock. **c** Client moving perpendicular to the dock. **d** Static client

#### 5.4 Beam steering power consumption


Finally, we study the power consumption of the beam searching process triggered by a temporary link outage

(due to human blockage). Figure 16 shows that after a 2 s disconnection, the beamforming process starts and lasts for around 3.5 s. During this interval, power exhibits large variations from 1500–3600 mW. Interestingly, the





(c)




(d)

Fig. 15 continued

beamforming phase is followed by another power state (9–10 s) marked as “Re-Association” in Fig. 16, during which power remains almost constant at 3600 mW before it

drops again down to the idle level (2000 mW). We observed a similar behavior at all distances. The average power consumption of the beamforming phase varied from

**Fig. 16** Power consumption in the case of temporary link outage and reconnection



2942 to 3344 mW across distances and the combined power consumption of the Beamforming/Re-Association phase varied from 3406 to 3838 mW.

Previous studies [32, 35, 45] showed the significant impact of the re-beamforming process on performance ([32, 45] using the same hardware as ours), concluding that it can nullify the benefits of narrow beams. Our study reveals a similar negative impact on power consumption. Together, these results show the need for more efficient beam searching algorithms.

## 6 Conclusions

We evaluated 60 GHz performance *across layers* and power consumption in a typical indoor WLAN environment using 802.11ad compliant wide-beam COTS devices. Our results suggest that 60 GHz radios equipped with relatively wide-beam antennas can be a viable option for multi-gigabit WLANs as they are more robust to client mobility while they still provide sufficient communication ranges. We also found that an 802.11ad NIC consumes much higher power than legacy WiFi NICs but its much higher throughput makes it significantly more energy efficient. On the other hand, the large number of reflective surfaces in typical indoor WLAN environments combined with wider beams make performance highly unpredictable and invalidate several assumptions that hold true in static, LOS scenarios, calling for new propagation models, rate adaptation algorithms, and evaluation methodologies. Additionally, the 802.11ad idle power is much higher than the 802.11n/ac idle power and the beam

searching process after a link outage also incurs a significant amount of power consumption, calling for new power management schemes and beam searching algorithms.

**Acknowledgements** This work was supported in part by NSF grant CNS-1553447. Viral Vijay Vira and Anuj Garg completed this work during their MS studies at the University at Buffalo. The contents of this work are solely the responsibility of the authors and do not represent the opinions or views of Amazon LLC or FactSet Research Systems Inc.

## References

1. Abari, O., Hassanieh, H., Rodriguez, M., Katabi, D.: Millimeter Wave Communications: from Point-to-Point Links to Agile Network Connections. In: *Proceedings of the ACM HotNets (2016)*.
2. Smartphones will Account for Nearly Half of Both 802.11ac and 802.11ad Chipset Shipments in 2018. <https://www.abiresearch.com/press/smartphones-will-account-for-nearly-half-of-both-802-11ac-and-802-11ad-chipset-shipments-in-2018/>, (2013).
3. Abouelseoud, M., Charlton, G.: The effect of human blockage on the performance of millimeter-wave access link for outdoor coverage. In: *Proceedings of the IEEE Vehicular Technology Conference Spring (VTC-Spring) (2013)*
4. Anderson, C. R., & Rappaport, T. S. (2004). In-building wide-band partition loss measurements at 2.5 and 60 GHz. *IEEE Transactions on Wireless Communications*, 3(3), 922–928.
5. ath10k: mac80211 wireless driver for qualcom atheros qca988x family of chips. <http://wireless.kernel.org/en/users/Drivers/ath10k>.
6. PCI EXPRESS X1 to PCI Express Mini interface adapter. <http://www.adexelec.com/pciexp.htm>.
7. Collonge, S., Zaharia, G., & Zein, G. E. (2004). Influence of the human activity on wide-band characteristics of the 60 GHz indoor radio channel. *IEEE Transactions on Wireless Communications*, 3(6), 2396–2406.

8. Dong, K., Liao, X., & Zhu, S. (2012). Link blockage analysis for indoor 60GHz radio systems. *Electronic Letters*, 48(23), 1506–1508.
9. Haider, M.K., Knightly, E.W.: Mobility Resilience and Overhead Constrained Adaptation in Directional 60 GHz WLANs: Protocol Design and System Implementation. In: *Proceedings of the ACM MobiHoc (2016)*.
10. Halperin, D., Greenstein, B., Sheth, A., Wetherall, D.: Demystifying 802.11n power consumption. In: *Proceedings of the USENIX Workshop on Power Aware Computing and Systems (2010)*
11. Halperin, D., Kandula, S., Padhye, J., Bahl, P., Wetherall, D.: Augmenting data center networks with multi-gigabit wireless links. In: *Proceedings of the ACM SIGCOMM (2011)*
12. Huang, J., Qian, F., Gerber, A., Mao, Z.M., Sen, S., Spatscheck, O.: A Close Examination of Performance and Power Characteristics of 4G LTE Networks. In: *Proceedings of ACM Mobisys (2012)*
13. HXI Gigalink 6651. <http://www.hxi.com/doc/hximodel6651-10-30-2013>.
14. Langen, B., Lober, G., Herzig, W.: Reflection and transmission behavior of building materials at 60 GHz. In: *Proceedings of the IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC) (1994)*.
15. Maltsev, A., Maslennikov, R., Sevastyanov, A., Khoryaev, A., & Lomayev, A. (2009). Experimental investigations of 60 GHz WLAN systems in office environment. *IEEE Journal on Selected Areas in Communications (JSAC)*, 27(8), 1488–1499.
16. Rasekh, M.E., Marzi, Z., Zhu, Y., Madhow, U., Zheng, H.: Noncoherent mmWave path tracking. In: *Proceedings of the ACM HotMobile (2016)*.
17. Monsoon power monitor. <http://www.msoon.com/LabEquipment/PowerMonitor/>.
18. Mudumbai, R., Singh, S., Madhow, U.: Medium Access Control for 60 GHz Outdoor Mesh Networks with Highly Directional Links. In: *Proceedings of the IEEE INFOCOM 2009, Mini Conference (2009)*
19. Nitsche, T., Bielsa, G., Tejado, I., Loch, A., Widmer, J.: Boon and bane of 60 GHz networks: practical insights into beamforming, interference, and frame level operation. In: *Proceedings of the 11th ACM/SIGCOMM International Conference on Emerging Networking Experiments and Technologies (CoNEXT) (2015)*
20. Rappaport, T. S., Felix Gutierrez, J., Ben-Dor, E., Murdock, J. N., Qiao, Y., & Tamir, J. I. (2013). Broadband millimeter-wave propagation measurements and models using adaptive-beam antennas for outdoor urban cellular communications. *IEEE Transactions on Antennas and Propagation*, 4(61), 1850–1859.
21. Rappaport, T. S., Heath, R. W., Jr., Daniels, R. C., & Murdock, J. N. (2014). *Millimeter wave wireless communications*. Upper Saddle River: Prentice Hall.
22. Saha, S.K., Deshpande, P., Inamdar, P.P., Sheshadri, R.K., Koutsonikolas, D.: Power-throughput tradeoffs of 802.11n/ac in smartphones. In: *Proceedings of the IEEE INFOCOM (2015)*.
23. Saha, S.K., Sun, L., Koutsonikolas, D.: Improving connectivity, coverage, and capacity in 60 GHz indoor WLANs using relays. In: *ACM MobiCom S<sup>3</sup>Workshop (2015)*
24. Sato, K., Manabe, T.: Estimation of propagation-path visibility for indoor wireless lan systems under shadowing condition by human bodies. In: *Proceedings of the IEEE Vehicular Technology Conference (1998)*.
25. SiBEAM Captures World's First 60GHz Millimeter-Wave Smartphone Design Win in Letv's Flagship Smartphone, Le Max. <http://www.businesswire.com/news/home/20150519005350/en/SiBEAM-Captures-World's-60GHz-Millimeter-Wave-Smartphone-Design#.VZM9AkT9q9s>.
26. Singh, S., Mudumbai, R., & Madhow, U. (2011). Interference analysis for highly directional 60-GHz mesh networks: The case for rethinking medium access control. *IEEE/ACM Transactions on Networking*, 19(5), 1513–1527.
27. Singh, S., Ziliotto, F., Madhow, U., Belding, E.M., Rodwel, M.: Distributed coordination with deaf neighbors: efficient medium access for 60 GHz mesh networks. In: *Proceedings of the IEEE INFOCOM 2010 (2010)*.
28. Singh, S., Ziliotto, F., Madhow, U., Belding, E. M., & Rodwell, M. (2009). Blockage and directivity in 60 GHz wireless personal area networks: from cross-layer model to multihop MAC design. *IEEE Journal on Selected Areas in Communications (JSAC)*, 27(8), 1400–1413.
29. Smulders, P. F. M. (2009). Statistical characterization of 60-ghz indoor radio channels. *IEEE Transactions on Antennas and Propagation*, 57(10), 2820–2829.
30. Smulders, P. F. M., & Correia, L. M. (1997). Characterisation of propagation in 60 GHz radio channels. *Electronics & Communication Engineering Journal*, 9(2), 73–80.
31. Sun, L., Sheshadri, R.K., Zheng, W., Koutsonikolas, D.: Modeling WiFi active energy consumption in smartphones for app developers. In: *Proceedings of the IEEE ICDCS (2014)*.
32. Sur, S., Venkateswaran, V., Zhang, X., Ramanathan, P.: 60 GHz Indoor Networking through Flexible Beams: A Link-Level Profiling. In: *Proceedings of the ACM SIGMETRICS (2015)*.
33. Taori, R., & Sridharan, A. (2015). Point-to-multipoint in-band mmWave backhaul for 5G networks. *IEEE Communications Magazine*, 1(53), 195–201.
34. Sur, S., Zhang, X., Ramanathan, P., Chandra, R.: BeamSpy: Enabling robust 60 GHz links under blockage. In: *Proceedings of the USENIX NSDI (2016)*.
35. Tie, X., Ramachandran, K., Mahindra, R.: On 60 GHz wireless link performance in indoor environments. In: *Proceedings of the PAM (2012)*.
36. GHz development system. <https://www.pasternack.com/60-ghz-systems-and-modules-category.aspx>.
37. WARP: Wireless open access research platform. <http://warpproject.org/trac>.
38. Warty, N., Sheshadri, R.K., Zheng, W., Koutsonikolas, D.: A first look at 802.11n power consumption in smartphones. In: *Proceedings of the ACM Mobicom International Workshop on Practical Issues and Applications in Next Generation Wireless Networks (PINGEN) (2012)*.
39. WiMi: Wisconsin Millimeter-wave Software Radio. <http://xyz.hang.ece.wisc.edu/wimi/index.html>.
40. Xu, H., Kukshya, V., & Rappaport, T. S. (2002). Spatial and temporal characteristics of 60-ghz indoor channels. *IEEE Journal on Selected Areas in Communications (JSAC)*, 20(3), 620–630.
41. Yunze Zeng Parth H. Pathak, P.M.: A First Look at 802.11ac in Action: Energy Efficiency and Interference Characterization. In: *Proceedings of the IFIP Networking (2014)*.
42. Zhang, J., Zhang, X., Kulkarni, P., Ramanathan, P.: OpenMili: A 60 GHz software radio platform with a reconfigurable phased-array antenna. In: *Proceedings of the ACM MobiCom (2016)*.
43. Zhou, A., Zhang, X., Ma, H.: Beam-forecast: Facilitating mobile 60 GHz networks via model-driven beam steering. In: *Proceedings of the IEEE INFOCOM (2017)*.
44. Zhou, X., Zhang, Z., Zhu, Y., Li, Y., Kumar, S., Vahdat, A., Zhao, B.Y., Zheng, H.: Mirror Mirror on the Ceiling: Flexible Wireless Links for Data Centers. In: *Proceedings of the ACM SIGCOMM (2012)*.
45. Zhu, Y., Zhang, Z., Marzi, Z., Nelson, C., Madhow, U., Zhao, B.Y., Zheng, H.: Demystifying 60ghz outdoor picocells. In: *Proceedings of the ACM MobiCom (2014)*.

46. Zhu, Y., Zhou, X., Zhang, Z., Zhou, L., Vahdat, A., Zhao, B.Y., Zheng, H.: Cutting the cord: A robust wireless facilities network for data centers. In: *Proceedings of the ACM MobiCom* (2014).
47. IEEE 802.11 Task Group AD. [http://www.ieee802.org/11/Reports/tgad\\_update.htm](http://www.ieee802.org/11/Reports/tgad_update.htm).



**Swetank Kumar Saha** received the B.Tech. degree in Computer Science and Engineering (CSE) from IIT-Delhi, New Delhi, India in May, 2013. He is currently working towards a Ph.D. 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.



**Darshan Godabanahal Malle-shappa** received the M.S. degree in Electrical Engineering from the University at Buffalo, the State University of New York, in 2017.



**Avinash Palamanda** received the M.S. degree in Electrical Engineering from the University at Buffalo, the State University of New York, in 2017.



**Viral Vijay Vira** received his Master's degree in Computer Science from the State University of New York at Buffalo in 2016. He is currently a Software Development Engineer at Amazon. His interests are broadly in experimental wireless networking and distributed computing.



**Anuj Garg** received the M.S. degree in computer science and engineering from the University at Buffalo, the State University of New York, in 2015. He is currently a software engineer at FactSet Research Systems Inc.



**Dimitrios Koutsonikolas** is an associate professor in the Computer Science and Engineering Department at the University at Buffalo, The State University of New York. He received his Ph.D. in Electrical and Computer Engineering at Purdue University in August 2010 and worked as a post-doctoral researcher in the same department from September to December 2010. His research interests are broadly in experimental wireless networking and mobile computing, with a focus on high performance/power-aware protocol design and implementation, testbed prototyping, network measurements, and novel wireless applications. Dimitrios is serving as the General Co-Chair of the International Conference on Embedded Wireless Systems and Networks (EWSN 2018) and as an Associate Editor for IEEE Transactions on Mobile Computing. He received the NSF CAREER Award in 2016, the UB School of Engineering and Applied Sciences Early Career Researcher of the Year Award in 2015, and Best Paper awards in SENSORCOMM 2007 and WCNC 2017. He is a member of IEEE, ACM, and USENIX.