

A Feasibility Study of 60 GHz Indoor WLANs

Swetank Kumar Saha, Viral Vijay Vira, Anuj Garg, Dimitrios Koutsonikolas
University at Buffalo, The State University of New York
Buffalo, NY 14260-2500
Email: {swetankk, viralvij, anujgarg, dimitrio}@buffalo.edu

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 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 (LOS) 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.

I. 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 [1], 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 [2], [3], [4]. Signal propagation is easy to model in these scenarios as it exhibits near-free space propagation properties [5], [2], [3]. However, the true potential of the mmWave technology cannot be realized if its use is limited to such scenarios. Recent work [6] 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* [7], [8]. The 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 [9]. A study from ABI Research predicts that smartphones will account for nearly half of 802.11ad chipset shipments in 2018 [10]. However, improved communication speeds generally come at the cost of higher power consumption. Studies in 802.11n/ac chipsets and smartphones have shown that power increases with PHY data rate [11], [12], [13] and channel width [14], [13], as well as with application layer throughput [15], [16]. 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.

This paper presents a feasibility study of 60 GHz indoor WLANs by evaluating 60 GHz client performance and power consumption in a typical academic office building. Recent work [8] has shown that, although narrow-beam antennas can greatly extend range, they yield poor performance in scenarios involving client mobility and human blockage. Our study assumes that 60 GHz WLAN APs and clients will use relatively wide-beam antennas to cope with mobility; this may invalidate some of the assumptions that hold true in static, LOS scenarios.

With respect to performance, *in contrast to previous works [7], [8], [17] which measured performance at a single layer (typically PHY) often using custom, non-standard compliant hardware and narrow-beam antennas [7], [8], we investigate performance across multiple layers using 802.11ad-compliant wide-beam COTS devices.* We add to the findings of recent experimental studies of 60 GHz link performance in indoor WLAN environments by answering three questions.

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

(2) How does link distance affect performance at different layers? We find 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 [2], [3], [7], [4].

(3) Is there any correlation among performance metrics at different layers? Can metrics at a lower layer serve as good

indicators of performance at higher layers? We find 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 [6], [8] due to limitations of available 60 GHz hardware,¹ can yield inaccurate results in typical indoor WLAN environments. Similarly, simple signal-strength based rate adaptation algorithms which have been used in recent simulation studies [2], [3], [7], [4] may not perform well in such environments.

Additionally, *we present what we believe to be the first measurements of power consumption of an 802.11ad NIC.* We answer the following questions:

(4) 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 find that 802.11ad NIC consumes much higher power than legacy WiFi (802.11n/ac) NICs but its much higher throughput 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.

(5) 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 find 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 antenna array orientation [7], human blockage [7], [6], [8], mobility [6], [8], interference patterns and spatial reuse [6], [8], [17], communication through reflections [17], and the benefits from using multiple-APs or relays [6], [19].

II. RELATED WORK

Initial experimental studies of 60 GHz in indoor environments focused on measuring and modeling channel propagation characteristics using dedicated channel sounding hardware (e.g., [20], [21], [22], [5], [23]). Tie et al. [7] studied link level performance of 60 GHz links with respect to blockage and antenna orientation. However, they used custom designed non-802.11ad hardware and measured performance of IP-over-wireless-HDMI. Sur et al. [8] conducted a link-level profiling of indoor 60 GHz links, using a custom software-radio platform (WiMi) [18] focusing on the capabilities and limitations of flexible beams. WiMi uses a small channel width of only

¹Commercial hardware often does not report PHY data rate and custom hardware [18] does not provide throughput commensurate with 802.11ad rates.

245 MHz and thus, it cannot achieve Gbps data rates. Hence, findings in [8] are extrapolated from RSS and noise floor measurements in narrow channels and they may not reflect the behavior of real 802.11ad links. In contrast to these works, we are using *COTS 802.11ad-compliant hardware* equipped with phased arrays and measure performance *across different layers* of the protocol stack via TCP data transfers. More recently, Nitsche et al. [17] conducted a measurement study of 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 are 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. Additionally, to our best knowledge, this is the first work that investigates 802.11ad power consumption.

Recent work also has argued for the use of 60 GHz technology to augment datacenters [2], [3], [6] and demonstrated the feasibility of this approach using both expensive proprietary devices [2], [3], [6] and the same COTS hardware we use in this paper [4]. The datacenter 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.

Channel sounding measurements have also been conducted in outdoor environments [24], [20], [25]. More recently, Zhu et al. experimentally demonstrated the feasibility of 60 GHz-based outdoor picocells [6] 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 [8], and several observations reported in that work do not hold for our use-case.

III. MEASUREMENT METHODOLOGY

Hardware Our 802.11ad link setup 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 2x8 phased array antennas with relatively wide main beams (30° – 40°) [4], [8] 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.

We repeated a 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 [26] driver. The card supports 3x3 MIMO operation, channel widths of 20 MHz, 40 MHz, and 80 MHz, and PHY data rates in the range 6.5-1170 Mbps.

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 an 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 ft wide) offering ample opportunities for reflection/multipath 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 (see Table I).

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 [27] 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 laptop's Mini PCIE slot. Finally, we used a Monsoon Power Monitor [28] to supply power to the setup and record the power consumed. 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.

Transmitter Height To decide upon the optimal height for the dock/transmitter (Tx), we fixed the height of the receiver (Rx) to 2'6", varied the Tx height from 2'6" to 6'6" at the interval of 1', and measured TCP throughput at each height. We picked two different distances (8 ft and 16 ft) at each of the two locations. We further repeated the measurements for four different orientations of the Rx antenna array, since in a real deployed WLAN, the AP's location and orientation will be fixed but mobile clients can have any possible orientation with respect to the AP. A general observation from our experiments was that *performance over 60 GHz indoor links is very sensitive to a number of factors – location, orientation, Tx height, Tx-Rx distance*. For example, changing the Tx height by just 1' can change TCP 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.

Our experiments suggested that 5'6" performs better in general, although in some cases 3'6" showed better performance. We picked 5'6" to reduce the impact of multipath caused by reflections from the floor. Also, a larger height for the AP is desirable in practice for a larger coverage distance/area.

IV. PERFORMANCE ACROSS LOCATIONS

In this section, we study the impact of location on the performance of 60 GHz links. We performed multiple experiments at each of the locations shown in Table I with same 16 different Tx/Rx orientations as in [7]. Figures 1(a), 1(b), 1(c) plot the average RSSI, the selected PHY rates, and the average TCP throughput at each of the 10 locations. 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.

TABLE I
MEASUREMENT LOCATIONS

Location#	Distance	Description
0	8'6"	Hall
1	16'	Hall
2	8'6"	Corridor/Sym.
3	8'6"	Corridor/Asym.
4	16'	Corridor/Asym.
5	8'6"	Drywall
6	8'6"	Glass
7	8'6"	Corner
8	8'6"	Lab
9	24'	Lab

Figures 1(a) and 1(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, Figure 1(b) 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 (Figure 1(a)) and TCP throughput (Figure 1(c)) 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" [22], 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%

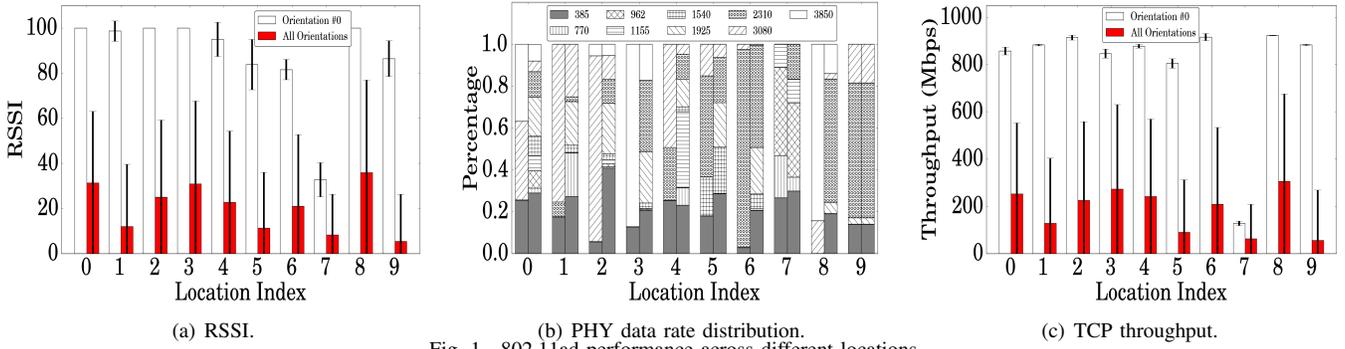


Fig. 1. 802.11ad performance across different locations.

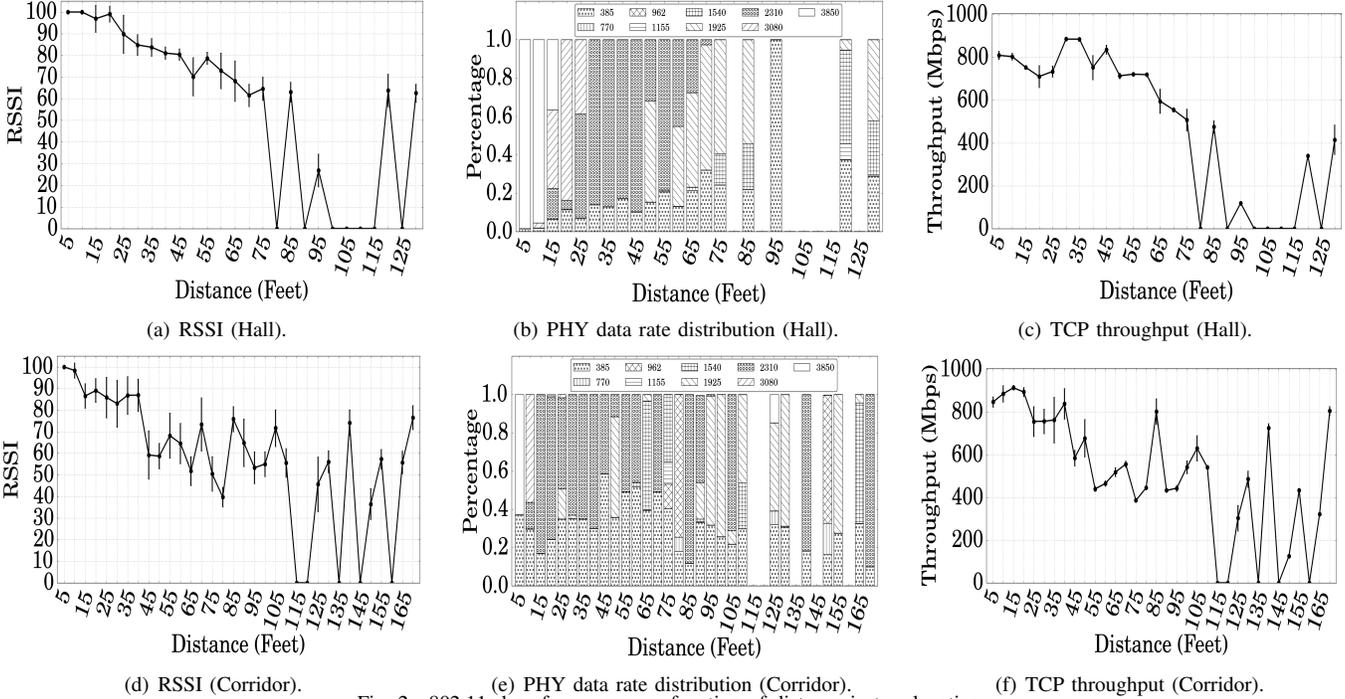


Fig. 2. 802.11ad performance as a function of distance in two locations.

of the time) and high throughput, but no link was established for any of the remaining 15 orientations. Although [22] found that attenuation due to clutter decreases as we move from 2.5 GHz 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 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.

V. IMPACT OF DISTANCE

We now focus on the impact of distance in LOS scenarios. Figure 2 plots the RSSI, the PHY rate distribution, and the TCP throughput over distance at the two main locations.

Range Figure 2 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 (23 dBm [6]) than the maximum allowed by FCC (40 dBm). The Corridor measurements show that RSSI exhibits large oscillations (due to a phenomenon known as *waveguide effect* [5]) but does not drop with distance beyond 40 ft (Figure 2(d)) and a PHY data rate of 2310 Mbps can be supported at a distance of 170 ft (Figure 2(e)). The Hall measurements show a different picture, closer to what one would expect, with RSSI gradually dropping with distance up to 75 ft (Figure 2(a)) 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 (Figure 2(b)). These ranges are much longer than the values reported recently with the same hardware (770 Mbps at 72 ft in a datacenter [4], 385 Mbps at 72 ft and 2310 Mbps at only 33 ft in an outdoor environment [6]).

RSSI vs. distance Recent experimental work [2], [3], [6], [4] observed that the attenuation of 60 GHz signals with distance follows closely the Friis model in LOS scenarios, both in stable

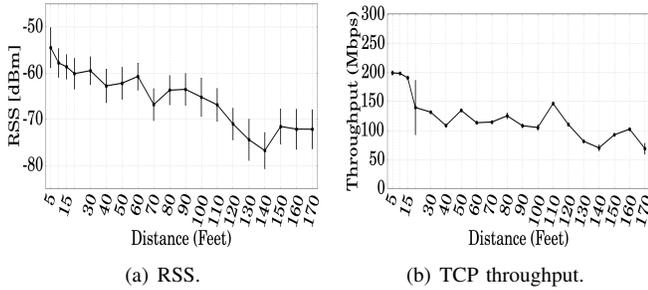


Fig. 3. 802.11ac performance as a function of distance in the Corridor.

datacenter and outdoor picocell environments. Figure 2 shows that this assumption does not hold true in the case of wide beams in indoor environments.

In Figure 2(a) (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 ft and 75 ft where RSSI decreases with distance. Lastly, distances between 80 ft 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 [6], [4] 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 [8].

In Figure 2(d) (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, Figure 3(a) 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 to Figure 2(d), despite the fact that the 802.11ac cards are equipped with omni-directional antennas.

PHY data rate vs. distance Figure 2(b) and 2(e) 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 (Figure 2(b)), 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 (Figure 2(e)).

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 [2], [3],

[7], [4] and the use of RSS as a direct indicator of the PHY data rate [6], [8]. Our results in Figures 2(a) and 2(d) 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, we calculate a theoretical RSS value based on the commonly used log-distance path loss model adjusted to account for shadowing fading [5] and potential losses due to reflections in case of NLOS links [6].

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

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

where $EIRP$ is equal to $23dBm$ for Wilocity Radios [6], $G_{RX}(dBi) = 10 \log_{10} N_{RX}$ is the receiver antenna gain as a function of the antenna elements $N_{RX} = 16$ [6], n is the path loss exponent (we use different values for Corridor and Hall based on [5]) and X_{Ω} represents a shadowing component (zero mean Gaussian random variable with standard deviation values also obtained from [5] for different environments). In [6], 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 [8]) to convert P_{RX} to a PHY data rate.

Figures 4(a) and 4(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, Figure 4(c) 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** Figures 2(c) and 2(f) show again distinct regions although these regions do not always overlap with the RSSI regions. In the Hall (Figure 2(c)), 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

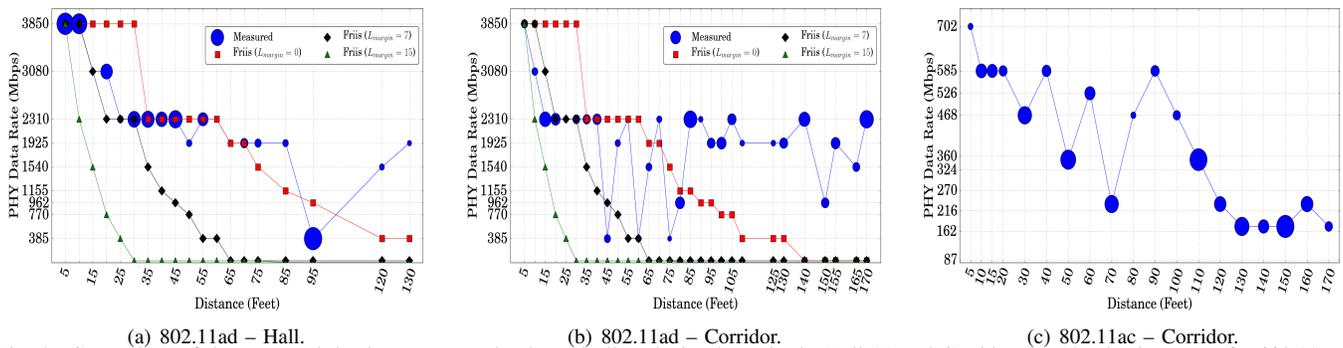


Fig. 4. Comparison of the measured dominant rate vs. the theoretically calculated rate in the Hall (a) and Corridor (b). The dominant rate for 802.11ac in the Corridor is included for comparison (c). A larger circle indicates larger dominance.

Corridor (Figure 2(f)), 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 3(b) 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 consistent performance (almost zero standard deviations) for a given distance. Note, however, that TCP throughput is affected more in legacy WiFi than in 802.11ad – it never exceeds 200 Mbps in Figure 3(b) although the dominant rates can be as high as 702 Mbps (Figure 4(c)), probably due to contention from other 802.11 networks in the 5 GHz band.

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 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 datacenter 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.

VI. 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 sec 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 5(a), 5(b) and 5(d), 5(e) plot the PHY data rate distribution and the average TCP throughput over RSSI in the Hall and Corridor, respectively.

Figures 5(a) and 5(d) 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, the picture changes when we compare the two locations. For the same RSSI bin, the observed data rates can be very 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 5(b) 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 (Figure 5(e)). 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 sec 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 5(c) and 5(f) plot the average TCP throughput over the PHY data rate in the Hall and Corridor, respectively. Note that the throughput corresponding to 385 Mbps data rate in Figure 5(f) is higher (450 Mbps) since higher rates were used in the remaining 10% of the time.

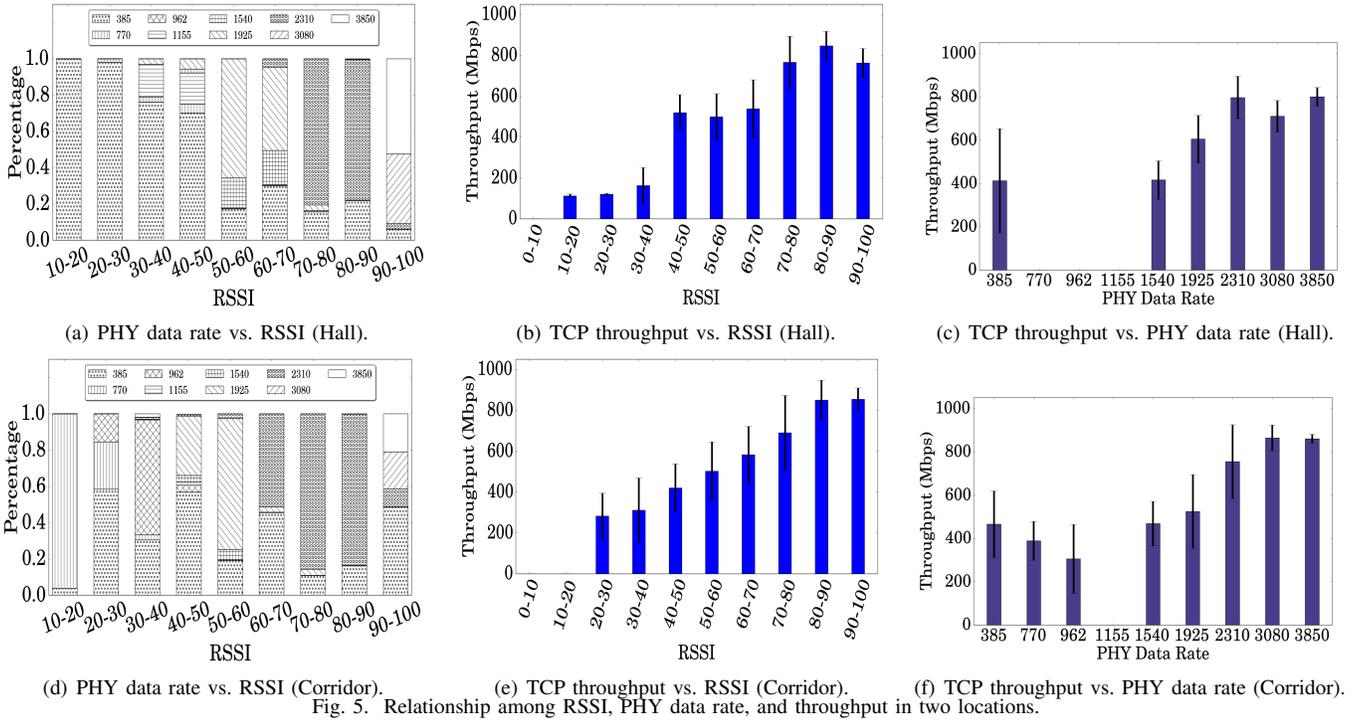


Fig. 5. Relationship among RSSI, PHY data rate, and throughput in two locations.

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. 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 [6], [8], 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 in datacenters [2], [3], [4] but also in indoor WLAN environments [7], may not work well in indoor WLANs.

VII. POWER CONSUMPTION

We now study the power consumption of an 802.11ad NIC. All the measurements are performed in the Corridor.

A. 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

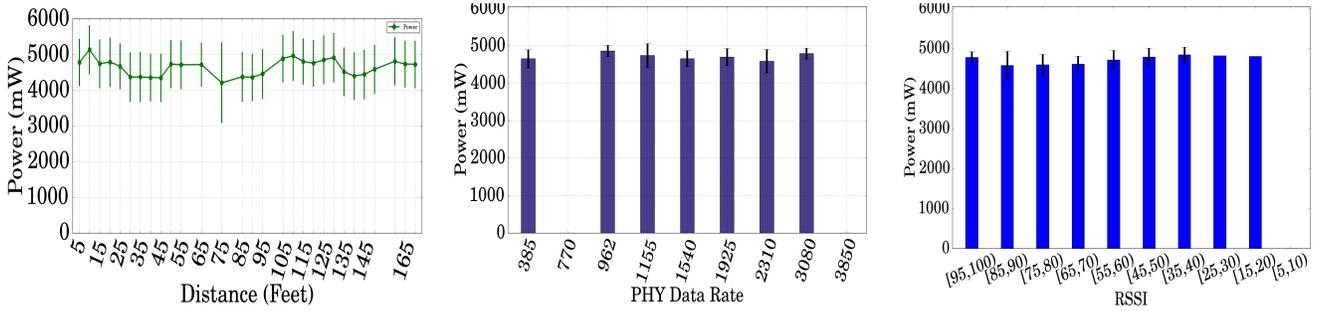
more than 2.5 W. Finally, in the *Connected/idle* state, the card is associated to the Dock but there is no Rx or Tx activity. Here the power is around 2.3 W.² For comparison, Halperin et al. [11] reported a power of only 820-1450 mW for an 802.11n WNIC in the *idle/connected* state, depending on the number of active antennas (1-3), and Zeng et al. [14] reported a power of 894-1196 mW for a 3x3 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.

B. Rx power consumption

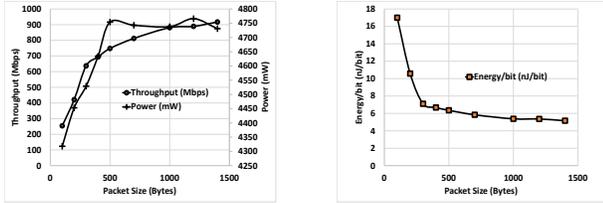
Distance, PHY data rate, and RSSI Recall from Figure 2(f) that throughput in the Corridor shows significant variations over distance. Given that WiFi power consumption is proportional to throughput [15], [16], one would expect a similar trend for power consumption over distance. Nonetheless, 802.11ad Rx power exhibits a very different trend over distance in Figure 6(a). 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 Figures 6(b), 6(c), we plot the Rx power as a function of the PHY data rate and the RSSI. Again, although both these factors affect throughput significantly (Figures 5(e), 5(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

²We occasionally observed this value to between 3.5 - 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.



(a) Rx power vs. distance. (b) Rx power vs. PHY data rate. (c) Rx power vs. RSSI.
 Fig. 6. 802.11ad Rx power consumption as a function of distance, PHY data rate, and RSSI. The error bars show the standard deviation.



(a) Throughput, power consumption. (b) Energy per bit.
 Fig. 7. Impact of packet size.

PHY data rate. Figure 7(a) plots the power and throughput as a function of the packet size and Figure 7(b) 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 bytes to 1400 bytes, power consumption increases from 4300 mW to 4730 mW (10%). On the other hand, the large throughput improvement (from 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. [11] 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. [14] reported similar values (900-1500 mW) for the Tx power of a 3x3 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, [11] reports minimum energy costs (dividing power by the bitrate instead of the achieved throughput) from 4-200 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 [11] (940 mW for MCS0, 1 spatial stream and 1600 mW for MCS9, 3 spatial streams),

we estimate an energy cost of 1.23-28.9nJ/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 ([14] 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 (Figure 2(f)).

Another benefit of 802.11ad becomes clear when we look at small packet sizes. Halperin et al. [11] 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 (Figure 7(b)).

C. 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 sec at walking speed while the dock was sending UDP traffic at full speed. We repeated each experiment 5 times. With all 3 types of motion, we observed that, although power consumption exhibited higher oscillations than in Figure 6(a) (especially in the case of the client away from or perpendicular to the dock), the average power consumption over the 10 sec period remained the same as in the static case (~ 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°).

D. 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 8 shows that after a 2 sec disconnection, the beamforming process starts and lasts for around

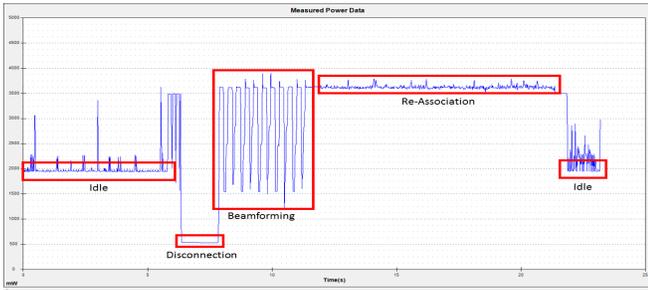


Fig. 8. Power consumption in the case of temporary link outage and reconnection.

3.5 sec. During this interval, power exhibits large variations from 1500-3600 mW. Interestingly, the beamforming phase is followed by another power state (9-10 sec) marked as “Re-Association” in Figure 8, 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 2942-3344 mW across distances and the combined power consumption of the Beamforming/Re-Association phase varied from 3406-3838 mW.

Previous studies [7], [6], [8] showed the significant impact of the re-beamforming process on performance ([6], [8] 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.

VIII. 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 (802.11n/ac) 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 algorithm after a link outage also incurs a significant amount of power consumption, in addition to the performance penalty which was observed by previous studies. Together, these two results call for new power management schemes and beam searching algorithms.

ACKNOWLEDGMENTS

This work was supported in part by NSF grant CNS-1553447.

REFERENCES

- [1] “IEEE 802.11 Task Group AD,” http://www.ieee802.org/11/Reports/Agad_update.htm.
- [2] 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.
- [3] 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.
- [4] 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.
- [5] 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.
- [6] Y. Zhu, Z. Zhang, Z. Marzi, C. Nelson, U. Madhoo, B. Y. Zhao, and H. Zheng, “Demystifying 60ghz outdoor picocells,” in *Proc of ACM MobiCom*, 2014.
- [7] X. Tie, K. Ramachandran, and R. Mahindra, “On 60 GHz wireless link performance in indoor environments,” in *Proc of PAM*, 2012.
- [8] 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.
- [9] “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>.
- [10] “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-8, June 2013>.
- [11] D. Halperin, B. Greenstein, A. Sheth, and D. Wetherall, “Demystifying 802.11n power consumption,” in *Proc. of USENIX Workshop on Power Aware Computing and Systems*, October 2010.
- [12] N. Warty, R. K. Sheshadri, W. Zheng, and D. Koutsonikolas, “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.
- [13] S. K. Saha, P. Deshpande, P. P. Inamdar, R. K. Sheshadri, and D. Koutsonikolas, “Power-throughput tradeoffs of 802.11n/ac in smartphones,” in *Proc. of IEEE INFOCOM*, 2015.
- [14] P. M. Yunze Zeng, Parth H. Pathak, “A First Look at 802.11ac in Action: Energy Efficiency and Interference Characterization,” in *Proc. of IFIP Networking*, 2014.
- [15] J. Huang, F. Qian, A. Gerber, Z. M. Mao, S. Sen, and O. Spatscheck, “A Close Examination of Performance and Power Characteristics of 4G LTE Networks,” in *Proc. of ACM Mobisys*, 2012.
- [16] L. Sun, R. K. Sheshadri, W. Zheng, and D. Koutsonikolas, “Modeling wifi active energy consumption in smartphones for app developers,” in *Proc. of IEEE ICDCS*, 2014.
- [17] 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/SIGCOMM International Conference on emerging Networking Experiments and Technologies (CoNEXT)*, December 2015.
- [18] “WiMi: Wisconsin Millimeter-wave Software Radio,” <http://xyzhang.ece.wisc.edu/wimi/index.html>.
- [19] S. K. Saha, L. Sun, and D. Koutsonikolas, “Improving Connectivity, Coverage, and Capacity in 60 GHz Indoor WLANs Using Relays,” in *ACM MobiCom S³ Workshop*, 2015.
- [20] 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.
- [21] 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.
- [22] 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.
- [23] 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.
- [24] B. Langen, G. Lober, and W. Herzig, “Reflection and transmission behavior of building materials at 60 GHz,” in *Proc. of IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, 1994.
- [25] T. S. Rappaport, J. Felix Gutierrez, E. Ben-Dor, J. N. Murdock, Y. Qiao, and J. I. Tamir, “Broadband Millimeter-Wave Propagation Measurements and Models Using Adaptive-Beam Antennas for Outdoor Urban Cellular Communications,” *IEEE Transactions on Antennas and Propagation*, vol. 4, no. 61, 2013.
- [26] “ath10k: mac80211 wireless driver for qualcomm atheros qca988x family of chips,” <http://wireless.kernel.org/en/users/Drivers/ath10k>.
- [27] “PCI EXPRESS X1 to PCI Express Mini interface adapter,” <http://www.adexelec.com/pciexp.htm>.
- [28] “Monsoon power monitor,” <http://www.monsoon.com/LabEquipment/PowerMonitor/>.