Multi-Gigabit Indoor WLANs: Looking Beyond 2.4/5 GHz

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

***Abstract*—This work explores the idea of building multi-Gigabit *indoor* enterprise WLANs using the millimeter-wave technology. Instead of the legacy 2.4/5 GHz band, we look at the unlicensed band around 60 GHz and the new 802.11ad standard, which supports PHY data rates of up to 6.76 Gbps. Through a set of measurements with commercial hardware, we show the feasibility of building multi-Gigabit WLANs. Using observations from our experiments and results reported in recent studies, we point out both the advantages and drawbacks of the 60 GHz technology, and identify the research challenges towards making multi-Gigabit 60 GHz indoor WLANs a reality.**

I. Introduction

We experience today an explosion in wireless network traffic; industry research predicts that the aggregate wireless bandwidth demand will increase by 1000x by 2020 [1]. One approach which has recently emerged as an alternative to the traditional 2.4/5 GHz wireless systems, promising multi-Gigabit throughput, is the use of millimeter-wave (mmWave) radios in the unlicensed 57-64 GHz spectrum (colloquially known as the 60 GHz band), supported by the new IEEE 802.11ad standard [2]. By leveraging very wide channels (2.16 GHz), 802.11ad provides bitrates between 385 Mbps and 6.76 Gbps. Moreover, the small form factor of 60 GHz hardware allows the use of large antenna arrays whose highly *directional* beams can significantly improve spatial reuse.

The caveat, however, is the high attenuation and vulnerability to blockage of 60 GHz links, owing to high frequency (small wavelength). E.g., free space propagation loss at 60 GHz, which scales up with square of the carrier frequency, is 21.6 dB worse than at 5 GHz. While high gain antenna arrays can compensate for the path loss, they introduce new challenges, due to human blockage and device mobility. Electronically steerable antenna arrays can theoretically overcome link outages, but the overhead of re-beamforming may counterbalance the potential gains [3], [4], [5].

These special characteristics of 60 GHz links have led to the assumption that mmWave radios are unsuitable for Non-Line-of-Sight (NLOS) communication. Consequently, the use of the 60 GHz technology had been limited (until very recently) to being a replacement for high-speed wired links in indoor scenarios, e.g., for HD video streaming in wireless personal area networks (WPANs) [6], or for augmenting data center networks with high capacity wireless links [7], [8], [4]. These scenarios are characterized by large open spaces providing direct LOS paths but largely devoid of reflective and/or obstructive surfaces and objects. The absence of phenomena such as reflection or multipath makes such environments easy to model as they exhibit near-free space propagation properties. However, the true potential of the mmWave technology cannot be realized if its use is limited to static, LOS scenarios. While a few recent works [3], [5] have considered the case of 60 GHz WLANs, the question of whether it is feasible to build general-purpose, indoor, multi-Gigabit WLANs out of 60 GHz radios remains largely unanswered.

In this paper, we explore the feasibility of building general-purpose, indoor, 60 GHz enterprise WLANs. Specifically, we are asking the question of whether multi-Gigabit WLANs can be built out of cheap, off-the-shelf, 802.11ad-compliant hardware. Note that, compared to WPAN or datacenter environments, typical indoor enterprise WLAN environment is highly complex, with many objects/surfaces that can attenuate, completely block, or reflect the signal, making it harder to predict the link behavior.

To answer this question, we conduct four sets of experiments in a typical academic office building. Our results suggest that even the currently available commercial 60 GHz radios, designed for short-range LOS scenarios, can provide Gigabit performance in a typical indoor WLAN setup and increase spatial reuse compared to legacy WiFi. However, it is often hard to predict or model performance in typical indoor WLAN environments. More importantly, our experiments reveal a number of serious challenges that have to be addressed for 60 GHz WLANs to become a reality. In the rest of the paper, based on our observations and findings of previous studies, we discuss the *good* (advantages), *bad* (drawbacks), and the *ugly* (unpredictable performance), of the 60 GHz technology and identify several research directions.

II. Related work

Our work is not the first to investigate the feasibility of 60 GHz technology in indoor WLANs. Initial experimental work focused on measuring and modeling channel propagation characteristics using dedicated channel sounding hardware (e.g., [9], [10], [11], [12]). Tie et al. [3] 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. More recently, Sur et al. [5] conducted a link-level profiling of indoor 60 GHz links, using a software-radio platform (WiMi). Their study offers many valuable insights, in particular about the potential capabilities and limitations of flexible beams. However, WiMi uses a small channel width of only 245 MHz and thus, it cannot achieve Gbps data rates. Hence, throughput values in this work are obtained from RSS and noise floor measurements in narrow channels using an 802.11ad specific rate table and they may not reflect the behavior of real 802.11ad links. In contrast to these works, in our study we measure both link and higher layer performance using off-the-shelf 802.11ad hardware and real data transfers over TCP.

Recent work also has argued for the use of 60 GHz technology to augment datacenters [7], [8], [4] and to build outdoor picocells [4], and demonstrated the feasibility of such approaches using both expensive proprietary devices [7], [8], [4] and the same cheap off-the-shelf hardware we use in this paper [4]. The datacenter environment, with static LOS links established on top of TOR switches, is very different from the complex indoor WLAN environment. Similarly, the outdoor picocell scenario differs greatly from the one we are concerned with, as also pointed out in [5], and several observations reported in that work do not hold for our use-case.

III. Experiments

In this section, we describe four sets of experiments we conducted along with the 802.11ad hardware and the experimental methodology.

Hardware and experimental methodology Our 802.11ad link setup consists of two commercially available 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 wireless interface and acts as an AP. Another laptop is connected to the dock through a Gigabit Ethernet interface to generate/receive TCP traffic. The use of the Ethernet interface limits the throughput in our experiments to 1 Gbps. The Wilocity radios support the following PHY data rates (in Mbps): 385, 770, 1155, 1540, 1925, 2310, 3080, 3850. They do not allow us to control the PHY layer data rate and use their own rate adaptation algorithm and an in-built beamforming mechanism. They report the current PHY data rate and an RSSI value between 0 and 100.

Although IEEE 802.11ad specifies antenna beams as narrow as 2.86^o, the Wilocity radios use 2x8 antenna arrays with a main beamwidth of $30^o - 40^o$ [13], [5]. Hence, our throughput and range measurements are probably lower bounds of the achievable performance of 802.11ad links. While radios equipped with narrow beam antennas can greatly extend range, recent work [5] has shown that they perform poorly with client mobility and human blockage. Hence, we believe that future WLANs may use wider beams.

We used iperf3 to generate TCP traffic. Each experiment consists of a 10-second TCP session. All the results are the average of 5 sessions. All experiments were performed late night to remove the possibility of human blockage. We leave the study of effects of human blockage on performance as future work. In our preliminary study here, we only deal with static objects present in the building environment.

Experiment 1 The goal of this experiment is to study the impact of location and antenna orientation on 802.11ad performance. We conducted measurements at 10 different locations inside an academic building following a methodology very similar to that in [3]. The chosen setups are diverse and represent the wide variety of scenarios that would typically occur in an office environment. We repeated the measurements at each of the locations with the same 16 different orientations of the Rx and Tx antennas as in [3] (see Table I and Figure 4(b) in [3]). We kept the transmitter and the receiver at a height of $5'6''$ and $2'6''$, respectively.

TABLE I
MEASUREMENT LOCATIONS AND ORIENTATIONS FOR EXPERIMENT 1

L#	Dist.	Desc.	Orientation					
0	8'6"	Open Space	**O#**	**Rx**	**Tx**	**O#**	**Rx**	**Tx**
1	16'	Open Space	0	→	←	8	←	←
2	8'6"	Corridor Sym.	1	→	↓	9	←	↓
3	8'6"	Corridor Asym.	2	→	→	10	←	→
4	16'	Corridor Asym.	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						

Figures 2(a), 2(b), 2(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. Figures 3(a), 3(b), 3(c) plot the same three metrics at each orientation, averaged across all locations. We also consider separately Location#0, which represents the best-case scenario (LOS 8ft).

Experiment 2 The goal of this experiment is to examine how channel quality indicators (RSSI, PHY data rate) and TCP throughput vary with the distance. We repeated the measurements at two different locations of the same building: a lobby on the first floor and a corridor on the third floor. Figures 4, 5 plot the RSSI, the PHY rate distribution, and the TCP throughput over distance at the two locations.

Experiment 3 This experiment aims to examine the spatial reuse due to the use of directional transmissions using three representative topologies (Figure 1). Each topology consisted of three links inside our lab, which houses 22 cubicles and is full of office-like furniture (chairs, desks, computer systems etc.). Each of the links was setup at a height of 4 ft and a LOS path existed between the Tx and Rx. In topology 1, the links were placed parallel with separation of 8 ft. between consecutive links. In topology 2, link 2 was perpendicular to links 1 and link 3. Topology 3 emulated a case where multiple users are co-located and are serviced by three different APs. All three links operated on the same channel. For each topology, we measured TCP throughput of each individual

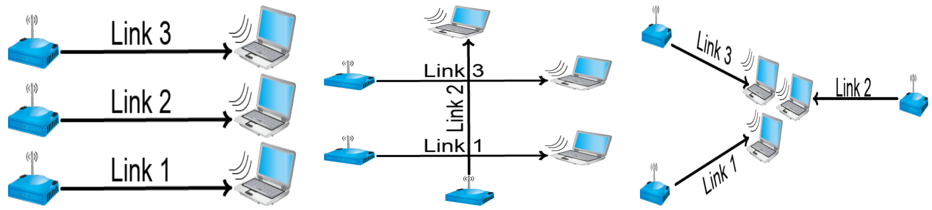

(a) Topology 1 (b) Topology 2 (c) Topology 3

Fig. 1. Spatial Reuse Topologies (Experiment 3).

link when it was operating alone and when all three links were active at the same time.

Experiment 4 In this experiment, we investigate the potential of using multiple APs as means to overcome blockage arising out of the presence of humans in the environment. In this context, we wanted to answer the following questions: (i) how often does such blockage occur in a typical WLAN and (ii) can serving a client with multiple APs (similar to the BS picocloud scenario in [4] for 60 GHz outdoor picocells) help mitigate this problem? To answer these questions, we used a methodology similar to that of [4] since Wilocity radios do not allow switching between APs on-the-fly. We deployed three links in our lab, very close to each other, emulating a single client which can potentially connect to any of the three docks/APs, in a topology very similar to that of Topology 3 (Figure 1(c)) for a period of 15 hours (which included both night and day time). We recorded per second TCP throughput for each of the three links. Figures 6(a), 6(b), and 6(c) present the CDF of throughput in three cases when one, two, or three APs are considered to be deployed.

IV. "The bad": Drawbacks of 60 GHz

Blockage mmWave frequencies are highly vulnerable to human blockage. A human in the LOS between the transmitter and the receiver can attenuate the signal by 20-30 dB [14], resulting in link outage. [4] 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 [3], [5] showed that human blockage remains a major challenge, and becomes worse due to the long reconnection times of existing 802.11ad hardware.

One potential way to address this problem is beam dilation, although [5] showed that it only works under high SNR scenarios. Hence, there is a need for faster, more efficient re-beamforming algorithms, potentially combined with mechanisms that distinguish the cause of link outage (human blockage or client mobility), as [5] showed that different approaches work better in each scenario. An alternative approach discussed in [4] for outdoor picocells is the use of a picocloud – multiple APs simultaneously serving a client; of course, such an approach would require a much denser AP deployment. We experiment with this approach to assess its effectiveness in Section VI.

Impact of antenna orientation Figure 2 shows that the performance averaged across all orientations is much lower than for Orientation #0 (when Tx and Rx are facing each other); RSSI (Figure 2(a)) and TCP throughput (Figure 2(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. This can be attributed to some orientations resulting 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 each.

Figure 3 shows that 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^o, 180^o, or 270^o, 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^o intervals are characterized by throughputs lower than 450 Mbps, and rather large standard deviations. 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 Figures 3(a), 3(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 equipped with steerable antenna arrays. However, practical 802.11ad antenna arrays cannot generate homogeneous beams across all directions [15]; this has been recently verified experimentally in [5].

Dead zones In Figures 4(a) and 4(c), we observe that RSSI and throughput drop to zero at 90 ft, raise to non-zero but very low values at 95 ft, drop again to zero for the range of 10-115 ft, and eventually raise back to high levels. Similar link outages can be observed in the corridor experiments, in Figures 5(a) and 5(c), at 115-120 ft, 135 ft, 145 ft, and 160 ft. 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 [4], [13] 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[5].

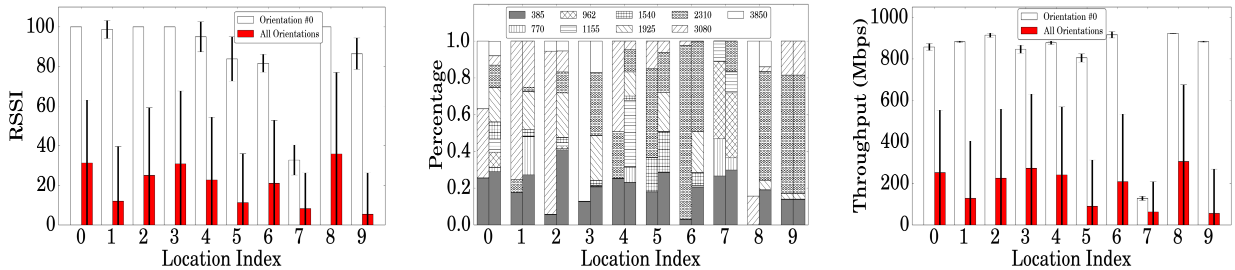

(a) RSSI. (b) PHY data rate distribution. (c) TCP throughput.

Fig. 2. Performance across different locations.

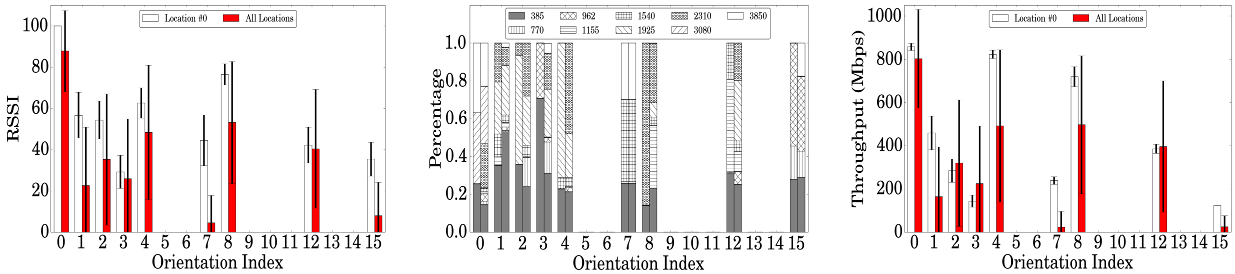

(a) RSSI. (b) PHY data rate distribution. (c) TCP throughput.

Fig. 3. Performance across different orientations

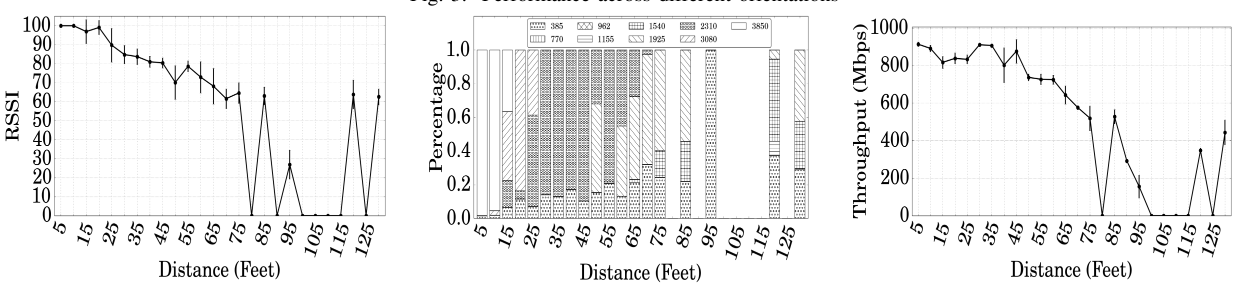

(a) RSSI. (b) PHY data rate distribution. (c) TCP throughput.

Fig. 4. Performance as a function of distance (lobby).

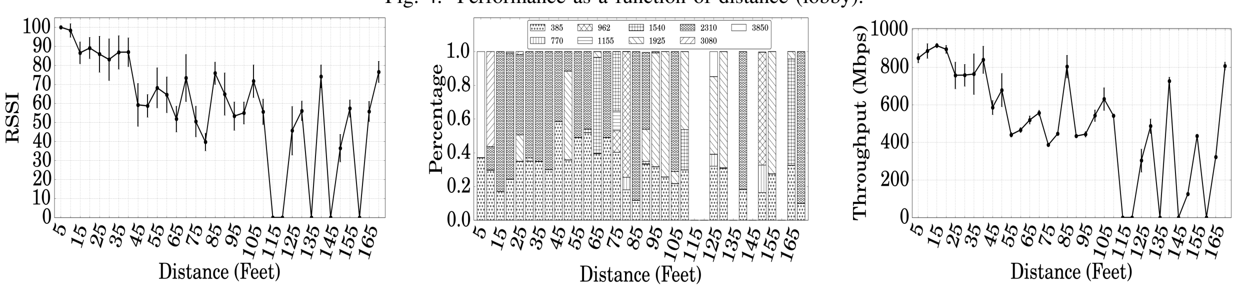

(a) RSSI. (b) PHY data rate distribution. (c) TCP throughput.

Fig. 5. Performance as a function of distance (corridor).

Investigating this tradeoff is part of our future work.

Impact of objects outside LOS Locations #8 and #9 are inside a research lab filled with "clutter" [10], i.e., objects that do not directly block the Tx-Rx LOS, such as office furniture, soft partitions that do not extend to the ceiling, and lab equipment. Figure 3 shows that the 24' link could sustain high data rates (1925 Mbps or higher for 85% of the time) and high throughput for Orientation #0, but no link was established for any of the remaining 15 orientations. Although [10] 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.

V. "The Ugly": Unpredictable performance

Impact of location Recent experimental work [7], [8], [4], [13] observed that 60 GHz signals attenuate with distance following closely the Friis model in LOS scenarios, both in stable datacenter and outdoor environments. The validity of the free-space propagation model has led to the use of simple RSS-based rate adaptation algorithms in simulators [7], [8], [3], [13] and the use of RSS as a direct indicator of the PHY data rate [4], [5]. Figures 4, 5 show that in indoor WLAN environments, these assumptions hold only partially in certain locations and are totally invalid in other locations.

Specifically, Figures 4(a) and 4(c) show that in the lobby both RSSI and throughput decrease with distance[1], with the exception of the range 85-120 ft, which we referred to in Section IV (dead zones). However, the measurements in the corridor (Figures 5(a) and 5(c)) show a very different picture. RSSI shows a decreasing trend with distance only for very short distances (5-15 ft), remains almost stable for distances 15-40 ft, and exhibits very large variations and non-monotonic behavior for distances longer than 40 ft. Throughput also exhibits non-monotonic behavior and large variations (up to 300 Mbps within 5 ft). More importantly, in certain cases, throughput variations do not follow the RSSI variations, e.g., between 45-55 ft or 75-80 ft. This large variability of RSSI and throughput with distance indicates the presence of strong multipath in certain typical WLAN environments, when nodes are equipped with relatively wide-beam antenna arrays, and the need for new propagation models in 802.11ad simulators.

As far as the PHY data rate is concerned, both Figure 4(b) and 5(b) show that for most distances there are 2 or 3 dominant data rates, and the lowest rate of 385 Mbps is used at least 10% and up to 60% of the time, even in the case of high RSSI. In the lobby experiments (Figure 4(b)) we still observe a monotonic decrease with distance and RSSI; lower data rates dominate at longer distances/lower RSSI values. In contrast, the is no such monotonicity in the corridor. Overall, we observe that RSSI cannot be used as an indicator of PHY data rate. This has two immediate implications. First, translating signal strength to PHY data rate can yield inaccurate results in typical indoor WLAN environments; the same observation is true for legacy 802.11 (e.g., [16]). Second, simple RSS-based rate adaptation algorithms may not be effective; more intelligent algorithms may be required in complex environments.

Orientation/link asymmetry Figures 3(a), 3(b), 3(c) show that, at Location #0 (lobby), orientations #4 and #12, which are symmetric w.r.t the Tx position, do not give similar throughput. 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 environmental asymmetry (there are still walls in the lobby 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 limitations). Further, [5] showed that 60 GHz links exhibit link asymmetry (downlink and uplink throughput are different when the Tx and Rx use different beamwidths). These asymmetries make it hard to predict and/or accurately model performance in indoor environments.

AP discovery and reconnection time Experiments in [5] show that AP discovery latency ranges from 5 ms to 1.8 s for static clients and up to 12.9 s for mobile clients. Further, experiments with Wilocity radios in [4] (which we confirmed with our hardware) and with non-802.11ad radios in [3] showed that re-beamforming delay after link directionality changes varies from 2-7 s. Such large delays can severely impact performance of 60 GHz links and large variations make it hard to model their impact.

[1]Since our hardware reports RSSI and not RSS, we cannot check the validity of the Friis model.

VI. "The Good": Advantages of 60 GHz

High throughput 802.11ad has the potential to cope with the predicted 1000x increase in mobile data traffic, by delivering multi-Gigabit throughput per client. By leveraging very wide channels (2.16 GHz), it supports data rates in the range of 385-6.76 Mbps. Our experimental results in Figures 2(b), 3(b), 4(b), and 5(b) confirm this showing data rates higher than 1 Gbps in several scenarios. Further, each AP can be equipped with three radios, each tuned on one of the three orthogonal channels, providing a maximum total downlink throughput of more than 20 Gbps per AP. In contrast, an 802.11ac AP equipped with 8 antennas, using 160 MHz channel, can provide a total of 6.9 Gbps using MU-MIMO technology and up to 13.8 Gbps if it is equipped with two radios (only 320 MHz of spectrum is available in the 5 GHz band). Note also that the antenna form factor in the 5 GHz band is much larger than in the 60 GHz band.

Spatial reuse Although the 802.11ac standard allows for 160 MHz wide channels, in practice it will be very hard to find such large chunks of free spectrum in the 5 GHz band, especially in dense WLAN deployments. In contrast, the use of directional mmWave links in 802.11ad allows for parallel transmissions over the same channel, further increasing channel capacity. To quantify the degree of spatial re-use, we use the Experiment 3 setup described Section III. We use the spatial reuse factor β introduced in [5], which is equal to the sum throughput of concurrent links divided by the average throughput of isolated links.[2]

Table II summarizes the results and compares them against legacy WiFi (802.11ac) using omni-directional antennas. Topology 1, where links are distant enough so that side lobes do not cause interference, provides for maximum spatial reuse. Topology 2 provides the least spatial reuse as Links 1 and 3 get around 60% of their isolated throughput. In topology 3, spatial reuse is better than in topology 2, even though the receivers of the three links are located very close to each other. On the other hand, spatial reuse for 802.11ac is close to 1 as it does not allow for concurrent transmissions.

TABLE II
Spatial Reuse Factor (β)

	Topology 1	Topology 2	Topology 3	802.11ac
β	2.95	2.16	2.78	0.92

Range A common belief is that communication range in 60 GHz is too short even in free space due to the very short wavelength. Figures 4 and 5 show that this is not true, even with radios using relatively wide beams and lower EIRP than the maximum allowed [4]. Our measurements in the corridor show that RSSI exhibits large oscillations[3] but does not drop with distance beyond 40 ft (Figure 5(a)) and a PHY data rate of 2310 Mbps can be supported at a distance of 170 ft

[2]β has a max. value of 3 (when there is no mutual interference) in the case of 3 links of similar quality. Higher value indicates better spatial reuse.

[3]Due to a phenomenon known as *waveguide effect* [11].

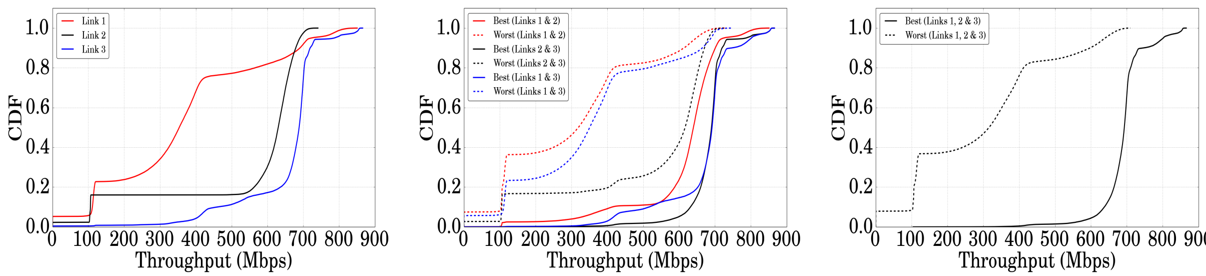

(a) 1 AP (b) 2 APs (c) 3 APs

Fig. 6. CDF of TCP throughput over 15 hours.

(Figure 5(b)). The measurements in the lobby show a different picture, closer to what one would expect, with RSSI dropping with distance (Figure 4(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 4(b)). These ranges are much longer than the values reported recently with the same hardware (770 Mbps at 72 ft in a datacenter [13], 385 Mbps at 72 ft and 2310 Mbps at only 33 ft in an outdoor environment [4]).

NLOS links Another major concern in the case of 60 GHz is the performance over NLOS links. Figure 2(c) shows that orientation #0 provides for near best possible performance (between 800-900 Mbps) at all locations, except one (Location #7). In fact, the standard deviations are negligible, indicating that the mean throughput was sustained across multiple runs. Hence, similar to the findings in [3], we observe that high-throughput NLOS 60 GHz links can be established through materials such as walls or glass. Although the signal does attenuate when it passes through such materials (Figure 2(a)), Figure 2(b) shows that, in the case of optimal antenna orientation, an NLOS link through a wall was able to sustain rates of 1540-3080 Mbps 80% of the time and an NLOS link through glass was able to sustain a rate of 2310 Mbps 95% of the time.

Overcoming Blockage In Figure 6(a), 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 (Figure 6(b)), we have 3 possible combination of APs. 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 percentage 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 in Figure 6(c). 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.

VII. Conclusion

Our preliminary results strongly suggest that there is indeed promise in pursuing the design of 802.11ad-based WLANs, to get the much needed multi-fold throughput increase. At the same time, we identified unique challenges in the use 60 GHz technology in a typical office environment, which differ significantly from findings reported in previous studies for LOS environments. We believe that our results will open multiple new research directions in MAC and PHY layer design towards multi-Gigabit mmWave indoor enterprise WLANs.

Acknowledgment

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

References

[1] "Mobile broadband usage is set to explode," http://www.techjournal.org/2011/09/mobile-broadband-useage-is-set-toexplode-infographic.

[2] "IEEE 802.11 Task Group AD," http://www.ieee802.org/11/Reports/tgad_update.htm.

[3] X. Tie, K. Ramachandran, and R. Mahindra, "On 60 GHz wireless link performance in indoor environments," in *Proc of PAM*, 2012.

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

[5] 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.

[6] Http://www.wirelesshd.org/.

[7] 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.

[8] 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.

[9] 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.

[10] 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.

[11] 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.

[12] 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.

[13] 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.

[14] S. Singh, F. Ziliotto, U. Madhow, E. M. Belding, and M. Rodwell, "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)*, vol. 27, no. 8, pp. 1400–1413, October 2009.

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

[16] S. H. Wong, H. Yang, S. Lu, and V. Barghavan, "Robust rate adaptation for 802.11 wireless networks," in *Proc. of ACM MobiCom*, 2006.