A First Look at TCP Performance in Indoor IEEE 802.11ad WLANs

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

***Abstract*—Radio communication in the millimeter-wave (mmwave) frequency bands, recently standardized by the IEEE 802.11ad specification, is emerging as an alternative to the traditional 2.4/5GHz wireless systems, promising multi-Gigabit throughput. Most work in this area until now has focused on PHY/MAC layer design. In contrast, little attention has been paid towards transport layer performance. In this work, we experimentally characterize TCP's behavior and its interplay with the lower layers of the protocol stack over 60GHz links in a typical office environment. Our results differ significantly from the those reported in the past for 60GHz links in a stable datacenter environment.**

I. Introduction

IEEE 802.11ad WLANs show great promise of matching up the ever increasing demand for bandwidth, by allowing a multi-fold increase in the maximum throughputs offered by 802.11ac. However, before 802.11ad WLANs can be realized in practice, many challenges need to be worked out, at all layers of the network stack. In this work, we experimentally study the feasibility of using 802.11ad as the replacement/augmentation of existing 802.11n/ac links in *indoor* WLANs. Given that the majority of internet traffic is composed of TCP flows, we seek to understand how TCP performance is affected by changes in the underlying physical link.

Related work Recent work by Zhu et. al. [1] demonstrated the feasibility of mmwave technology for building *outdoor* picocells. The works in [2], [3], [4] study the feasibility of using 802.11ad in a datacenter environment which is very different from the WLAN environment we consider in this study. In [5], Tie et. al. used custom 60GHz radios in an indoor office setting, but they limit their work to characterizing link level performance.

II. Experimental Methodology

We conducted experiments in the lobby of an academic building. Our 802.11ad link setup (Figure 1) 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, even though the wireless link itself is capable of much higher speeds. Hence, we also use the average PHY data rate as an indicator of the maximum achievable throughput. The Wilocity radios do not allow us to control the PHY layer data rate and use their own rate adaptation algorithm[1] and an in-built mechanism to control the beamforming direction.

All the reported results are the average of 5 10-second iperf sessions. In order to experiment with different link conditions, we vary the distance between the transmitter and the receiver, which results in varying RSSI. Although the cards are capable of finding a NLOS path (if it exists) in case the LOS path is blocked, we use LOS paths so that our results at different distances can be compared. The cards report the PHY data rate chosen by the rate adaptation mechanism and an RSSI value between 0 and 100. We also obtain the application layer throughput every 0.5 seconds from iperf.

III. RSSI, PHY Data Rate vs. distance

We first examine how channel quality indicators, RSSI and PHY data rate, vary with the distance between the transmitting and receiving ends of the link. Figures 2(a) and 2(b) show the average values of PHY data rate and RSSI, respectively, at various distances, when the separation between the two ends of the link varies between 2 feet and 155 feet.

PHY data rate vs. distance Figure 2(a) shows an overall decreasing trend of the data rate with distance. However, there is significant variation in the intermediate distances, where data rates do not exhibit a monotonic behavior. We can divide the distance axis into three distinct regions. For distances up to 20 feet, the two highest possible rates of 3080 Mbps and 3850 Mbps can be supported. The zero standard deviations further indicate that those two rates were consistent across the 5 runs. The next region is between 20-135 feet, where the possible link rates vary from 2935 Mbps to 1317 Mbps. Also, the standard deviations observed in this region are rather large, indicating that even for a static client, the rate adaptation logic was not able to maintain a stable rate. Lastly, distances between 135-155 feet (after which the link breaks) are characterized by a sharp decrease in the data rate. However, standard deviations are low/zero here, indicating again that lower data rates were chosen consistently.

RSSI vs. distance To understand the data rate's variability with distance, we look at the corresponding values of RSSI, which serve as a direct measure of the quality of the underlying wireless medium. Figure 2(b) shows a trend very similar to what we observed for the data rate. We can in fact very closely match the three regions observed before. Previous works [2], [3], [4] observed a clear drop of RSSI with distance in the much more stable datacenter environment which justified the use of the free-space propagation model in simulation

[1] The following PHY data rates are supported by our cards (in Mbps): 385, 770, 1155, 1540, 1925, 2310, 3080, 3850.

Fig. 1. Experimental setup.

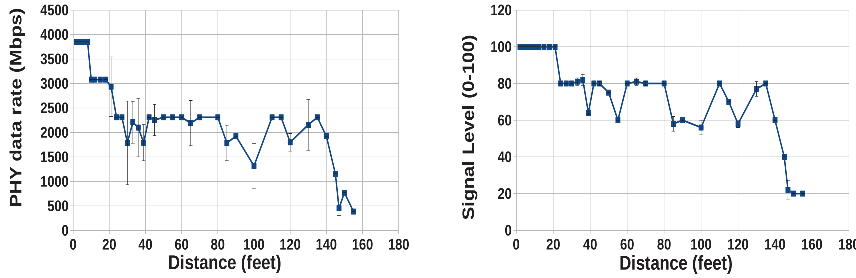

(a) PHY data rate vs. distance.
(b) RSSI vs. distance.

Fig. 2. Variation of PHY data rate and RSSI with distance.

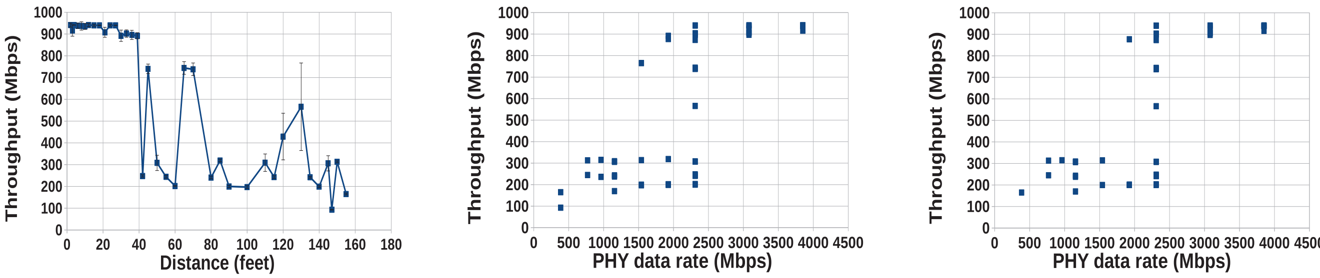

(a) Throughput vs. distance.
(b) Throughput vs. PHY data rate ($p = 75\%$).
(c) Throughput vs. PHY data rate ($p = 90\%$).

Fig. 3. TCP throughput vs. distance and PHY data rate.

studies [2], [3], [5], [4]. In contrast, *the large variability of RSSI with distance in Figure 2(b) indicates the presence of multipath fading in a typical WLAN environment, and the need for new propagation models in 802.11ad simulators.*

The validity of the free-space propagation model in datacenter environments has also led to (i) the use of simple rate adaptation algorithms in simulators based on the received SINR [2], [3], [5], [4] and the use of SINR as a direct indicator of the PHY data rate [1]. Figures 2(a), 2(b) show that, although PHY data rate's and RSSI's variability with distance exhibit a similar trend at a macroscopic level, standard deviations are much higher in the case of PHY data rate and the same RSSI value at different distances does not always result in the same data rate. *These observations show that RSSI can only serve as a coarse-grained indicator of PHY data rate and motivate the need for more intelligent rate adaptation algorithms.*

IV. TCP Throughput vs. Distance, PHY Data Rate

We now focus on TCP throughput and study how it varies as a function of distance and PHY data rate.

Throughput vs. Distance Figure 3(a) shows a picture very similar to the data rate and RSSI plots we observed before. Note that maximum possible TCP throughput is achieved until the distance of 35 feet, even through the corresponding data rates between 20 and 35 feet are lower than the maximum possible data rate. However, this could be attributed to the fact that our applied source data rate at the sender is only limited to 1 Gbps. The region of interest here is between the 35 feet and 135 feet, where throughput varies from as high as 750 Mbps to as low as 200 Mbps. Further, the average throughputs reported have large standard deviations. *We conclude that, contrary to the finding in previous studies [2], there is no correlation between TCP throughput and distance, except for short distances.*

Throughput vs. PHY data rate The reason for the lack of correlation between TCP throughput and distance becomes clear by looking at the corresponding section of the PHY rate plot (Figure 2(a)); the rate adaptation logic fails to provide a stable rate for a given distance and a monotonic behavior overall. However, it is possible that the PHY data rate itself could be used, instead of distance, as an indicator of TCP throughput. To answer this question, we consider only those sessions where a particular data rate was selected by the rate adaptation logic more than a threshold $p\%$ of time and calculate the average TCP throughput for each session. Figures 3(b) and 3(c) plot the distribution of throughputs for each data rate corresponding to the 75% and 90% thresholds, respectively. In both plots, PHY rates above 3000 Mbps yield the maximum TCP throughput of 900 Mbps, where rates below 1500 Mbps result in the lowest throughputs. Interestingly, for PHY rates between 1500 Mbps and 2500 Mbps, we observe multiple TCP throughput values, ranging from very high to very low. In fact, this particular set of data rates are the ones corresponding to the unstable region in the distance plots (Figures 2(a) and 2(b)). *We conclude that for a large number of intermediate PHY data rates, the corresponding TCP throughput is highly unpredictable.* Although TCP is already known to perform poorly over wireless links, the problem is exaggerated here because of significantly higher data rates.

Overall, our study shows that TCP performance over 60GHz links in an indoor WLAN is very different from the performance in a stable datacenter environment and cannot be accurately predicted from lower layer metrics. Our observations motivate future work in multiple directions including new rate adaptation algorithms and cross-layer design.

References

[1] Y. Zhu et al., "Demystifying 60ghz outdoor picocells," in *Proc of ACM MobiCom*, 2014.

[2] D. Halperinet al., "Augmenting data center networks with multi-gigabit wireless links," in *Proc. of ACM SIGCOMM*, 2011.

[3] X. Zhou et al., "Mirror Mirror on the Ceiling: Flexible Wireless Links for Data Centers," in *Proc. of ACM SIGCOMM*, 2012.

[4] Y. Zhu et al., "Cutting the Cord: a Robust Wireless Facilities Network for Data Centers," in *Proc. of ACM MobiCom*, 2014.

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