

Potholes Ahead: Impact of Transient Link Blockage on Beam Steering in Practical mm-Wave Systems

Adrian Loch
IMDEA Networks Institute
Madrid, Spain
adrian.loch@imdea.org

Irene Tejado
IMDEA Networks Institute and
Universidad Carlos III, Madrid, Spain
irene.tejado@imdea.org

Joerg Widmer
IMDEA Networks Institute
Madrid, Spain
joerg.widmer@imdea.org

Abstract—The practical realization of beam steering mechanisms in millimeter wave communications has a large impact on performance. The key challenge is to find a pragmatic trade-off between throughput performance and the overhead of periodic beam sweeping required to improve link quality in case of transient link blockage. This is particularly critical in commercial off-the-shelf devices, which require simple yet efficient solutions. First, we analyze the operation of such a commercial device to understand the impact of link blockage in practice. To this end, we measure TCP throughput for different traffic loads while blocking the link at regular intervals. Second, we derive a Markov model based on our practical insights to compute throughput for the case of transient blockage. We use this model to evaluate the trade-off between throughput and periodic beam sweeping. Finally, we validate our results using throughput traces collected using the aforementioned commercial device. Both our model and our practical measurements show that transient blockage causes significant signal fluctuation due to suboptimal beam realignment. In particular, fluctuations increase with traffic load, limiting the achievable throughput. We show that choosing lower traffic loads allows us to reduce fluctuations by 41% while achieving the same net throughput than with higher traffic loads.

I. INTRODUCTION

Distinguishing link degradation due to mobility from that due to blockage in millimeter wave (mm-Wave) networks is challenging. Existing consumer-grade transceivers operating in this band often cannot tell both apart. Due to the high path loss at mm-Wave frequencies, transceivers use beamforming to overcome attenuation [1], [2]. Both the transmitter and the receiver must steer their beams towards each other to achieve a high Signal-to-Noise Ratio (SNR). If one of them moves, they must adjust their beams accordingly. However, in case of transient link blockage due to, e.g., a person crossing the link, both transceivers should maintain their original beam steering. If a device classifies such blockage as mobility, adjusting the beam while the link is blocked may lead to a misalignment.

Figure 1 shows an example of the impact of transient blockage on the Transmission Control Protocol (TCP) throughput for a commercial off-the-shelf (COTS) 60 GHz device. While related work also studies the impact of blockage on 60 GHz communications [3], [4], it does not consider the beam steering misalignment that may occur as a result. In Figure 1, we observe that after the first blockage at second 25, the throughput stabilizes at more than 200 mbps less than that prior to the blockage, even though the link is unobstructed. In other words, our COTS device has interpreted the first blockage as mobility, thus trying to adapt its beam steering. While the new

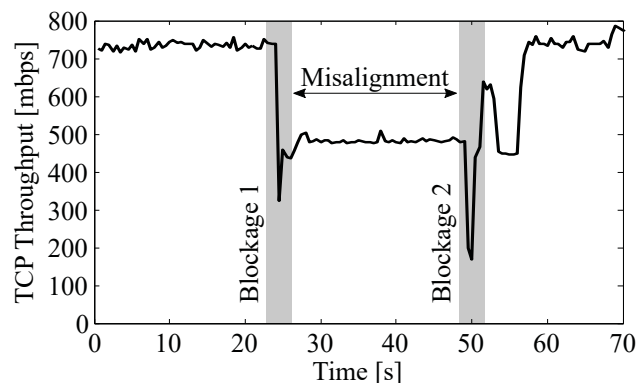


Fig. 1. Throughput “pothole” during subsequent transient blockage events.

steering allows for communication, the SNR is lower, resulting in lower TCP throughput. At second 50 in Figure 1, a second blockage causes another beam realignment, which in this case fortunately results in the original beam steering and allows TCP to achieve again roughly 700 mbps. Figure 1 shows that the impact of transient blockage can be highly detrimental to throughput if the transceiver classifies blockage as mobility, causing a throughput “pothole” while the link is misaligned. First, the link operates significantly below its potential. Second, the throughput exhibits high fluctuations. Regarding the latter, we find that throughput can stabilize at a number of different levels after a blockage. This is critical for traffic which requires stable, high throughput links, such as uncompressed video.

Related work proposes a number of beam steering solutions that partially solve the above problem. The 802.11ad standard suggests beam tracking to address mobility. That is, devices continuously track SNR variations and adapt their beam steering accordingly. This avoids costly beam sweeping, i.e., exhaustively probing all directional beam patterns of a transmitter to find the receiver after the transmitter and/or the receiver move. Still, beam tracking itself cannot distinguish mobility from blockage, resulting in the behavior in Figure 1. Similarly, other approaches that aim at reducing the complexity of beam sweeping [5] cannot distinguish them neither. In contrast, Blind Beam Steering (BBS) [6] estimates the direction in which the receiver is located. Essentially, the transmitter uses an antenna array operating at 2.4 GHz to estimate the angle-of-arrival, and then uses this information to perform beam steering at 60 GHz. Hence, transceivers can distinguish mobility from blockage. However, BBS considers a 2.4 GHz antenna array with at least four antennas, which at a minimum antenna separation of

$\lambda/2$ might be impractical for certain devices. Other approaches exploit information about transmit directions which worked prior to an SNR drop to narrow down the exhaustive beam steering search [7]. Alternatively, geometric analysis can also reveal which alternative beams are available at any location within a room [8]. However, such approaches are meaningful for long-lived blockages which require finding alternative non-line-of-sight paths, while we consider a *transient* blockage scenario. Recent work [9] also suggests analyzing the initial samples of an SNR drop in order to determine whether it is due to mobility or blockage. Such an approach must operate at the physical layer since it requires a timely reaction, i.e., the device must identify the cause of the SNR drop immediately when it starts. This strict time constraint may hinder its implementation in COTS devices. Although the propagation characteristics in the 60 GHz band allow for precise tracking [10], distinguishing mobility from blockage remains challenging.

As a result of the above challenge, COTS devices typically implement simple heuristics that result in the fluctuating behavior shown in Figure 1. In this paper, we analyze such throughput fluctuations and show their relation to frame-level aggregation [11]. Specifically, our contributions are as follows:

- We measure TCP throughput, Modulation and Coding Scheme (MCS) adaptation, frame-level aggregation, and beam steering for a COTS device to understand the impact of transient link blockage in practice.
- We derive an analytical Markov chain model for transient link blockage based on our practical insights. We use it to analyze the trade-off between throughput and the overhead caused by periodic beam sweeps.
- We show that frame aggregation can compensate for MCS drops due to steering misalignments, thus reducing throughput fluctuations by up to 41% while maintaining high throughput on COTS 60 GHz devices.

The remainder of this paper is structured as follows. In Section II we analyze the practical observations that allow us to understand the impact of transient blockages. After that, Section III presents our Markov chain model, and Section IV discusses the results that we obtain both from the model and from a COTS device. Finally, Section V concludes the paper.

II. PRACTICAL OBSERVATIONS

In this section, we use a 60 GHz COTS device to understand the impact of transient link blockage. To this end, we set up a 60 GHz link and walk through it at indoor pace.

A. Experiment setup

COTS networking devices that operate at mm-Wave frequencies are not yet widely available. However, some notebook manufacturers have implemented the WiGig [12] standard—which is very similar to 802.11ad—to design wireless docking stations. Basically, docking station and notebook establish a 60 GHz link to replace the physical connection between both devices. We use such a setup for our practical experiments, namely, a Dell D5000 docking station and a Dell Latitude E7440 notebook. More details about these devices can be found in [13]. We place the devices on two tables and separate

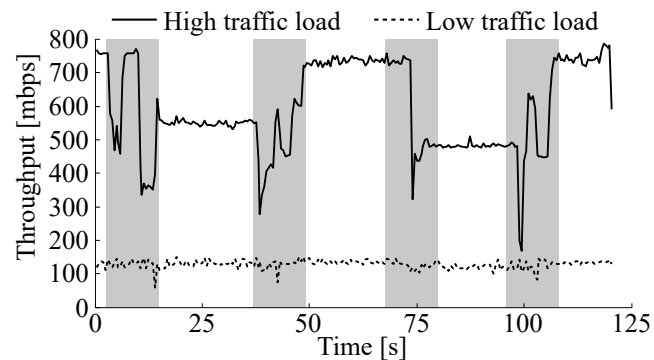


Fig. 2. Suboptimal throughput levels after transient blockages (gray shaded).

them about two meters. We use iperf to transmit data on the 60 GHz link connecting both devices. To control the traffic load, we adjust the TCP window size (WS). For each experiment, we record the signal level, the MCS, the TCP throughput, and the frame aggregation size. To record the signal level, we overhear the communication using a SiversIMA FC1005V/00 V-band converter connected to a Universal Software Radio Peripheral (USRP) X310. We use it to obtain energy level traces that allow us to infer the beam alignment and the frame aggregation size.

B. Effects at the application layer

Figure 2 shows the application layer TCP throughput for both a high and a low link load. In this case, we measure the throughput every 500 ms, and induce transient link blockages roughly every 20 to 30 seconds. Each time we cross the link, we observe a significant throughput drop. Moreover, for increasing link loads our measurements show that the throughput drops at each blockage become larger and the number of suboptimal levels increases. As sketched in Section I, the highest level corresponds to the best beam alignment for a certain load. The other levels correspond to suboptimal beam alignments resulting in a range of MCS levels. While Figure 2 clearly shows that transient link blockages have a larger impact at higher traffic loads, the underlying reason is not evident.

C. Analysis at the lower layers

To understand the reason for the behavior in Figure 2, we measure lower layer metrics during a transient link blockage. Figure 3 and Figure 4 show these metrics for the case of low (WS 10 kB) and high (WS 100 kB) link load, respectively.

1) *Low link load:* At low load, the link achieves a rate of roughly 400 Mbps before the blockage, as shown in Figure 3(c). As soon as we walk across the link, the signal level shown in Figure 3(a) drops significantly. However, after the blockage, it only rises again to about half of the amplitude. This indicates that the docking station has switched to a different sector due to the blockage. Since the docking station is steering in a different direction while the transmit power remains unchanged, we observe a different signal level at the SiversIMA converter. In other words, the signal level drop in Figure 3(a) has triggered a sector level sweep which has led to a suboptimal sector since the person was still crossing the link when it took place. As a result, the SNR at the receiver dropped and thus the link switched to a significantly lower MCS, as shown in Figure 3(b). Still, while the physical layer rate has

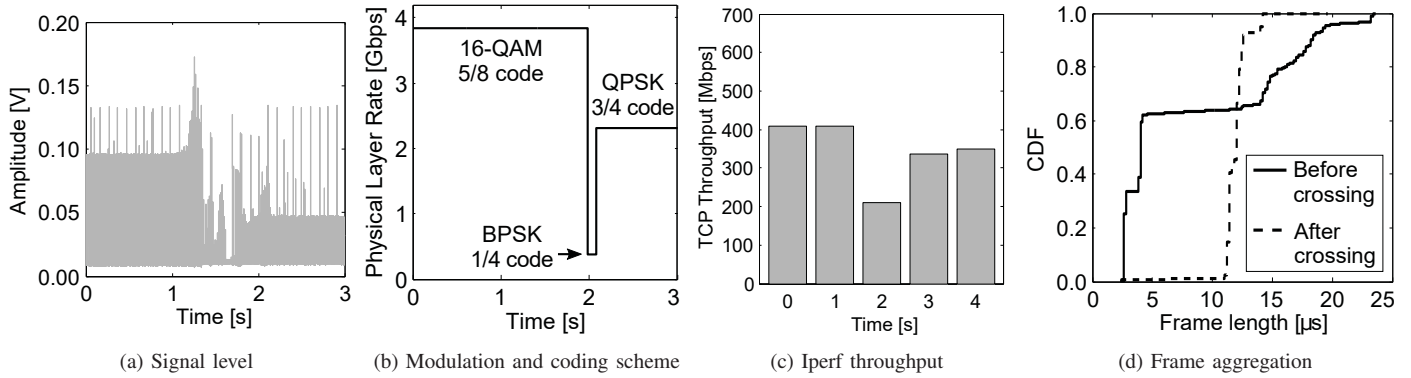


Fig. 3. Blockage with low link load in practice. The long blockage in (a) is likely due to the automatic gain control adjustment of the SiiversIMA converter.

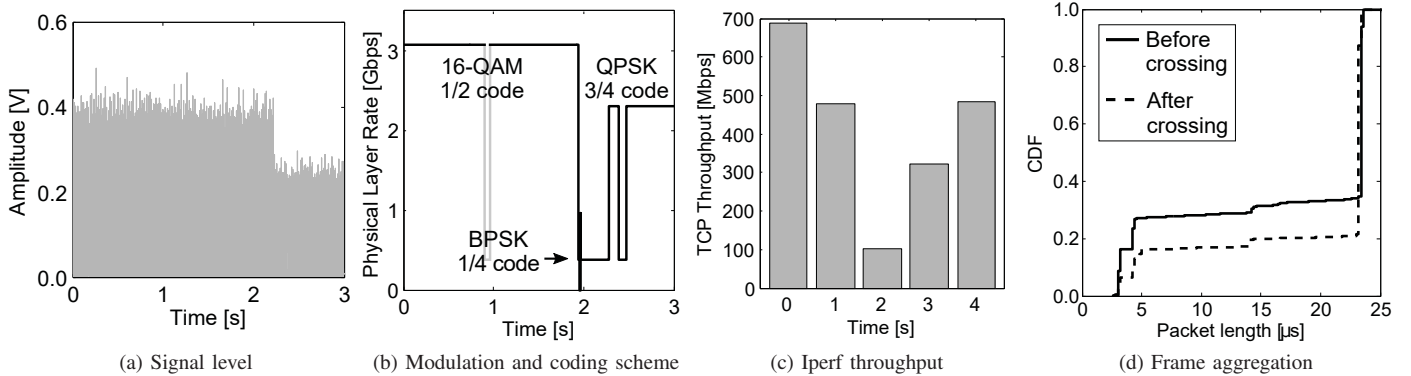


Fig. 4. Blockage with high link load in practice. The brief MCS drop grayed out in (b) is not due to the blockage but occurred by chance in this measurement.

nearly halved, the iperf throughput in Figure 3 (c) stabilizes at only 14.5% less throughput than prior to the blockage.

Figure 3 (d) explains this somewhat surprising behavior. Before the blockage occurs, the duration of more than 60% of all data frames is less than $5 \mu\text{s}$. However, after the blockage most frames are longer than $10 \mu\text{s}$. That is, frame aggregation increases. The Dell D5000 is known to aggregate all packets in the transmit queue whenever it gets access to the channel [13]. Due to the drop in MCS after blockage, the transmission time increases. During the transmission, packets continue arriving at the queue at the same pace. Hence, the average queue length after the blockage is larger than before the blockage, resulting in larger aggregated frames. This means that sector mismatches due to transient link blockages have a small impact on application layer throughput for low link loads because the frame aggregation in 802.11ad compensates for it.

2) *High link load*: For high link load, in Figure 4 (a) we observe again a sector mismatch after the blockage. As expected, also the MCS in Figure 4 (b) drops. Still, in Figure 4 (c) we observe that the impact on application layer throughput is much larger than for the low link load case—the throughput drops by roughly 30%. The underlying reason is that, at high link load, the docking station is already using a high level of aggregation. As shown in Figure 4 (d), approximately 70% of all data frames are longer than $20 \mu\text{s}$ before the blockage. Thus, the D5000 can only increase aggregation by about 10%, which is not enough to compensate for the lower MCS. This suggests that increasing the traffic load beyond a certain level might not pay off for a given frequency of transient link blockages.

III. SYSTEM MODEL

From Section II we conclude that transient blockage may cause sector misalignment which in turn affects MCS and ultimately frame aggregation. Moreover, we show that aggregation can mitigate blockage at moderate link loads. In the following, we derive a Markov model based on these insights.

A. Model overview

Due to suboptimal sector alignments, a link may be in $N+1$ different states: N states due to all feasible combinations of transmit and receive sectors, and one state when the link is blocked. Figure 5 shows an example of such a Markov model for a link with $N = 3$ throughput levels, that is, one high level L_H , and two suboptimal levels SL_1 and SL_2 . The blockage state is named B. Whenever a transient link blockage occurs, the model transitions to blockage state B. From there, it moves to one of the N states until the link is blocked again. If the state is one of the suboptimal ones, the model directly transitions to L_H whenever the device performs a periodic sector level sweep. We choose a time-slot size of one millisecond for our Markov model since this matches the timescale at which sector level sweeps are expected to occur. However, it is straightforward to adjust the model parameters for other time-slot sizes.

We model the occurrence of transient link blockage as a Gaussian distribution $\mathcal{N}(\mu, \sigma^2)$. The probability p_S of staying in state L_H for one time-slot is directly related to μ . To derive it, we first formulate the probability $P(k)$ of staying $k > 1$ time-slots at that state as in Equation 1. Then, we compute the average time T that we stay at the state as in Equation 2.




Fig. 5. Example Markov model for a link experiencing transient blockages.

$$P(k) = p_S^{k-1}(1 - p_S) \quad (1)$$

$$T = \sum_{k=1}^{\infty} kP(k) = \sum_{k=1}^{\infty} k \cdot p_S^{k-1} \cdot (1 - p_S) = \frac{1}{1 - p_S} \quad (2)$$

Since blockage occurs on average every μ seconds, $T = \mu$ and thus the probability of staying at L_H is $p_S = 1 - 1/\mu$. Note that the units of μ are time-slots, i.e., milliseconds in our case. From this we directly obtain that the transition probability from L_H to B is $1 - p_S = 1/\mu$. We follow a similar reasoning to compute the probabilities of staying at state B . If blockages last on average t_b , the probability of staying at state B is $1 - 1/t_b$. However, after a blockage, the link may recover to any of the other N states. Hence, we weight the probability of leaving state B , i.e., $1/t_b$, with the probabilities of transitioning to each of the states. In the example in Figure 5, p_H is the probability to transition to L_H while p_{SL1} and p_{SL2} are the probabilities of transitioning to each of the suboptimal states, respectively. In general, $p_H + \sum_{\forall i} p_{SLi} = 1$ must hold.

Finally, the probability of staying at the suboptimal states is related to μ in the same manner as for L_H . Still, in this case we must take into account that the device may perform a periodic sector level sweep that triggers a transition to L_H . We define p_C as the probability that a transient link blockage occurs before the next periodic sector level sweep takes place. Hence, the transition probability from SL_i to L_H is \bar{p}_C . Since the probability of a sector sweep and a blockage are independent, we weight the result of Equation 2 with p_C to obtain the probabilities of leaving and staying at a suboptimal state.

B. Periodic sector level sweeps

In the following, we derive the analytic expression of p_C based on the statistical characteristics of blockages and the frequency of sector level sweeps. We consider that devices perform such sector level sweeps at regular intervals of length S . For 802.11ad, S translates directly into the Beacon Interval (BI), that is, the interval at which a station may transmit beacons to improve beam steering. Without loss of generality, we define the time origin $t = 0$ as the point in time when the last sector level sweep took place. Given that a blockage occurs at time $t = t_b$, the time that remains until the next

sector level sweep is $S - t_b$. Hence, the probability of another blockage C occurring prior to the next sector level sweep for time t_b is as in Equation 3, where erf is the error function.

$$\begin{aligned} P(C, t = t_b) &= \int_0^{S-t_b} \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(t-\mu)^2}{2\sigma^2}} dt \\ &= \frac{1}{2} \left(\operatorname{erf} \left(\frac{\mu}{\sqrt{2}\sigma} \right) - \operatorname{erf} \left(\frac{\mu - S + t_b}{\sqrt{2}\sigma} \right) \right) \end{aligned} \quad (3)$$

Next, we obtain the probability in Equation 3 for any value of $t_b \in \{0, S\}$. Using the definition of conditional probability, we obtain $P(C, t_b) = P(C, t = t_b)P(t_b)$, where $P(t)$ is evenly distributed because sector level sweeps occur at fixed intervals. Hence, p_C is the marginal probability of $P(C, t_b)$ for all values of t_b . Based on this, we compute p_C as in Equation 4.

$$\begin{aligned} p_C &= \int_0^S P(C, t_b) dt_b = \int_0^S P(C, t = t_b) \cdot P(t_b) dt_b \\ &= \int_0^S \frac{1}{2} \left(\operatorname{erf} \left(\frac{\mu}{\sqrt{2}\sigma} \right) - \operatorname{erf} \left(\frac{\mu - S + t_b}{\sqrt{2}\sigma} \right) \right) \cdot \frac{1}{S} dt_b \\ &= \left(\frac{S - \mu}{2S} \right) \left(\operatorname{erf} \left(\frac{\mu}{\sqrt{2}\sigma} \right) - \operatorname{erf} \left(\frac{\mu - S}{\sqrt{2}\sigma} \right) \right) - \\ &\quad \sqrt{\frac{2}{\pi}} \frac{\sigma}{2S} \left(e^{-\frac{\mu^2}{2\sigma^2}} - e^{-\frac{(\mu-S)^2}{2\sigma^2}} \right) \end{aligned} \quad (4)$$

C. Throughput

Next, we derive the throughput for each of the $N + 1$ states of our Markov model. We consider a 60 GHz transmitter that operates as observed in Section II. That is, the transmitter has a queue of size q . Each time it gets access to the channel, the device aggregates up to a_{\max} data from the queue and transmits it. While it is transmitting, more packets arrive at the queue. Hence, the size of the aggregated frames directly depends on the transmission time t_{TX} and the channel access time t_{acc} . The latter includes the Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) overhead as in 802.11ad. Further, control messages increase each BI by a factor f .

The application generates a link load l . The maximum link load that the device supports occurs when it aggregates a_{\max} data in each transmission. We compute l for a certain aggregation factor $a_{\text{factor}} \in \{0, 1\}$ such that the device aggregates a_{\max} for $a_{\text{factor}} = 1$. In this case, blockage has a significant impact since the device cannot aggregate more to mitigate the impact of lower MCS values (c.f. Section II). Based on l , we compute the queue level q_L for each of the N states. Each state is related to a certain MCS. For transmission i , the queue level is the transmission time t_{TX} of the previous transmission multiplied by the link load l , as shown in Equation 5.

$$q_{L,i} = t_{TX,i} \cdot l = \left(\frac{q_{L,i-1}}{\text{MCS}} + t_{\text{acc}} \right) \cdot f \cdot l \quad (5)$$

While $l \leq \text{MCS}$, the above equation converges to a stable queue level q_L for $i \rightarrow \infty$. Based on this, we derive the throughput THP for each state. If the queue level is less than or equal to a_{\max} , the throughput is directly the link load l , as shown in Equation 6. However, if $q_L > a_{\max}$, the application

is generating more data than the link can transmit. Eventually, the queue level reaches its maximum length q . In this case, we limit the amount of data in each transmission to a_{\max} .

$$\text{THP}_N = \frac{qL}{t_{\text{TX}}} = \begin{cases} \frac{t_{\text{TX}} \cdot l}{t_{\text{TX}}} = l & \text{if } q_L \leq a_{\max} \\ \frac{a_{\max}}{\left(\frac{a_{\max}}{\text{MCS}} + t_{\text{acc}}\right) f} & \text{if } q_L > a_{\max} \end{cases} \quad (6)$$

Finally, we compute the throughput in state B. To this end, we define the blockage duration t_B as the time during which the blockage affects the throughput of the link. In general, t_B is the sum of (a) the time t that the link is physically blocked and (b) the time to transmit the data that accumulates at the application while the link is blocked. The latter includes the new data d_{new} that arrives while serving the accumulated data. Further, the accumulated data is limited to the queue size q . If the new state to which the Markov model transitions after a blockage already operates at the maximum aggregation size, (b) is infinite since the link cannot transmit more data. In that case, the duration of the blockage effect is just (a) and the related throughput is zero. Otherwise, we observe a higher throughput while (b) is ongoing. In this case, we compute the throughput in state B as in Equation 7. Since the duration of (b) depends on the MCS of the state to which the Markov model transitions after the blockage, we compute t_B as the average of all t_{B_i} values for $j \in [1, N]$ weighted with the probabilities $p_j \in \{p_H, p_{SL_i}\}$ of transitioning to each of the states.

$$\text{THP}_B = \sum_{\forall j} p_j \frac{\min(t \cdot l, q) + d_{\text{new}}}{t_{B_j}} \quad (7)$$

IV. EVALUATION

Using our Markov model in Section III, we study the impact of transient link blockage both in theory and based on practical traces from our 60 GHz wireless docking station.

A. Analytical evaluation

For our analytical evaluation, we set the parameters of our model as in Table I. We consider a 60 GHz link with $N = 3$ states. Whenever a blockage occurs, the probability of transitioning to each of the states is the same. We evaluate the impact of four parameters, namely, the Sector Sweep Interval (SSI), μ , a_{factor} , and v . Parameter v is the speed at which a person causing the blockage walks across the link. We use it to compute t , i.e., the amount of time that the link is physically blocked. To this end, we assume a beamwidth of α and that the person crosses the link at a distance d of the transmitter.

1) *Link load*: Figure 6 depicts our results for increasing link loads and average blockage intervals μ . As a starting point, we set the SSI to infinity, that is, the device only switches to a different sector when a blockage occurs. While the average throughput increases with the link load, the improvement becomes very small as the load approaches one. As discussed in Section II, this occurs because aggregation cannot mitigate the impact of lower MCS values. Further, Figure 6 also shows that the throughput fluctuation increases significantly with the link load. For instance, while increasing the link load from 0.4 to 0.6 improves throughput by just 8% for $\mu = 20$, it worsens throughput fluctuations by 25%. Hence, increasing the

link load would hardly pay off in this case. For comparison, we include the result for an SSI of 100 ms. As expected, the average throughput is much higher since the link quickly switches back to the best sector after a blockage. Also, the variance decreases much faster with μ because the transmitter is using the same sector most of the time.

2) *Sector Sweep Interval*: In our second analysis, we study the impact of the SSI. Figure 7 depicts a trade-off regarding how often a device performs sector level sweeps. If the SSI is very short, the resulting overhead limits the throughput significantly. Conversely, if it is too long, the blockages also reduce the average throughput. We achieve the best performance when the SSI matches the average blockage interval μ . Hence, a meaningful strategy for 60 GHz devices would be to estimate how often significant SNR drops occur in a given scenario, and set the SSI to that value. Regarding throughput fluctuations, very short SSI values are best since this ensures that the link is most of the time in the best state. As the SSI increases, the variance becomes larger because the link is more often in suboptimal states. Beyond $\text{SSI} = \mu$, the variance increases again for large μ since link state changes are mostly due to blockages. While sector sweeps ensure that the link recovers to the best state, blockages may lead to suboptimal states.

3) *Speed*: Next, we study how the speed of the person causing the blockage affects throughput. In Figure 8, we observe that the higher the speed, the higher the throughput is. This is expected, since the person blocks the link for a shorter amount of time t . Regarding variance, fluctuations generally decrease with speed because the system stays less time in state B, which often only achieves very low or even zero throughput. Initially, the opposite occurs for low values of v . In this case, the system is most of the time in state B since blockages last long. Hence, increasing the time that the link operates in any of the other N states also increases fluctuations.

TABLE I. EVALUATION PARAMETERS (DEFAULTS UNDERLINED)

Parameter	Value	Parameter	Value
t_{DIFS}	$13\mu\text{s}$	α	20°
t_{SIFS}	$3\mu\text{s}$	d	1.5
t_{ACK}	$6\mu\text{s}$	v	[0.1m/s, 4m/s] Default: <u>1m/s</u>
t_{SLOT}	$5\mu\text{s}$	MCS	3.85Gbps, 1.925Gbps, 1.155Gbps
CW_{\min}	15	P_H	1/3
t_{sweep}	4ms	P_{SL1}, P_{SL2}	1/3, 1/3
SSI	[10ms, ∞]	q	793.5kB
μ	[2s, 20s]	a_{factor}	0.2, 0.4, 0.6, 0.8, <u>1.0</u>
σ	0.1s	a_{\max}	79.35kB

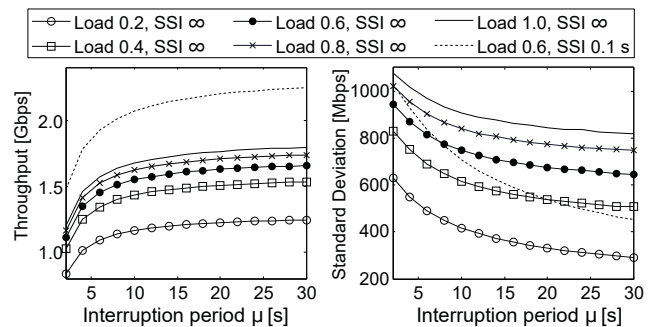
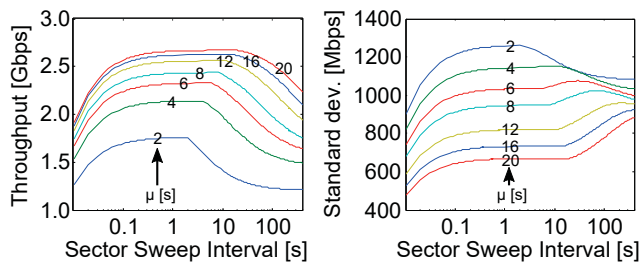



Fig. 6. Throughput and throughput fluctuation for different link traffic loads.

Fig. 7. Throughput and throughput fluctuation for different SSI and μ values.Fig. 8. Throughput and throughput fluctuation for different v and μ values.

B. Practical evaluation

We obtain practical results using our COTS 60 GHz system described in Section II-B. We use the traces to compute empirically the number of states N as well as the probabilities p_H and p_{SLi} . Further, $p_C = 1$ since we do not observe periodic sector level sweeps in the traces. To control the link load, we adjust the TCP WS. We configure our Markov model with the above practical parameters to obtain the throughput and the fluctuations for any value of μ . Figure 9 depicts our result, which is the practical counterpart of Figure 6. We obtain equivalent results in both, which validates our theoretical approach. However, in Figure 6 fluctuations decrease with μ while in Figure 9 they are roughly stable. The reason is that in Figure 6 larger μ values imply that throughput is zero less often, thus decreasing fluctuations. However, in Figure 9, iperf measures throughput at most every 0.5 s. Since blockage often lasts less, we do not observe zero throughput but always the average of the blockage and the following milliseconds. Hence, larger μ values have a much smaller impact on the dispersion of the throughput values. Further, in Figure 9 we observe that TCP WSs beyond 100 kB barely provide additional throughput because aggregation cannot mitigate MCS degradations. Still, the respective variance continues to increase significantly. As an example, for $\mu = 5$ s, the throughput for WS 100 kB is the same as for WS 150 kB but the variance is 41% larger. This confirms our link load trade-off hypothesis in Section II-C2.

V. CONCLUSION

We study the practical impact of transient link blockages on 60 GHz links. We observe that sector level sweeps during a blockage may cause a device to use a suboptimal sector. We analyze two methods to mitigate this. The first one introduces periodic sector level sweeps. To optimize throughput, we find that such sweeps should take place on average as frequently as blockages. However, to minimize throughput variance, sweeps should take place as often as possible. The second method optimizes the link load at the application layer. We find that

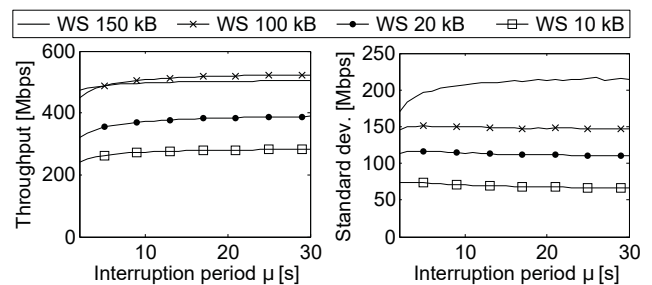


Fig. 9. Average throughput for different link loads in our practical setup.

moderate loads allow frame aggregation to compensate for low MCSs due to suboptimal sectors. In particular, we show that choosing a suitable link load can reduce throughput variance by 41% compared to a higher load with the same throughput.

ACKNOWLEDGMENTS

This paper is partially supported by the European Research Council grant ERC CoG 617721, the Ramon y Cajal grant of the Spanish Ministry of Economy and Competitiveness RYC-2012-10788, and the Madrid Regional Government through the TIGRE5-CM program (S2013/ICE-2919).

REFERENCES

- [1] D. R.C. and R. Heath Jr., "60 GHz Wireless Communications: Emerging Requirements and Design Recommendations," *IEEE Vehicular Technology Magazine*, vol. 2, no. 3, 2007.
- [2] D. De Donno, J. Palacios, D. Giustiniano, and J. Widmer, "Hybrid Analog-Digital Beam Training for mmWave Systems with Low-Resolution RF Phase Shifters," in *Proc. ICC'16*, 2016.
- [3] Y. Zhu, Z. Zhang, Z. Marzi, C. Nelson, U. Madhow, B. Y. Zhao, and H. Zheng, "Demystifying 60GHz Outdoor Picocells," in *Proc. ACM Mobicom*, 2014.
- [4] N. Moraitis and P. Constantinou, "Indoor Channel Measurements and Characterization at 60 GHz for Wireless Local Area Network Applications," *IEEE Trans. Antennas and Propagation*, vol. 52, no. 12, 2004.
- [5] B. Li, Z. Zhou, H. Zhang, and A. Nallanathan, "Efficient Beamforming Training for 60-GHz Millimeter-Wave Communications: A Novel Numerical Optimization Framework," *IEEE Trans. on Vehicular Technology*, vol. 63, no. 2, 2014.
- [6] T. Nitsche, A. Flores, E. Knightly, and J. Widmer, "Steering with eyes closed: Mm-wave beam steering without in-band measurement," in *Proc. IEEE INFOCOM 2015*, 2015.
- [7] A. Patra, L. Simić, and P. Mähönen, "Smart mm-Wave Beam Steering Algorithm for Fast Link Re-Establishment Under Node Mobility in 60 GHz Indoor WLANs," in *Proc. of the 13th ACM International Symposium on Mobility Management and Wireless Access*, 2015.
- [8] S. Sur and X. Zhang, "BeamScope: Scoping Environment for Robust 60 GHz Link Deployment," in *Proc. of Information Theory and Application Workshop*, 2016.
- [9] S. Sur, V. Venkateswarana, X. Zhang, and P. Ramanathan, "60 GHz Indoor Networking through Flexible Beams: A Link-Level Profiling," in *Proc. of ACM SIGMETRICS'15*, 2015.
- [10] T. Wei and X. Zhang, "mTrack: High-Precision Passive Tracking Using Millimeter Wave Radios," in *Proc. of ACM Mobicom*, 2015.
- [11] D. Skordoulis, Q. Ni, H. h. Chen, A. P. Stephens, C. Liu, and A. Jamalipour, "IEEE 802.11n MAC Frame Aggregation Mechanisms for Next-Generation High-Throughput WLANs," *IEEE Wireless Communications*, vol. 15, no. 1, 2008.
- [12] C. Hansen, "WiGiG: Multi-Gigabit Wireless Communications in the 60 GHz Band," *IEEE Wireless Communications*, vol. 18, no. 6, 2011.
- [13] 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. CoNEXT'15*, 2015.