

Can Millimeter-Wave Support Interactive Extended Reality under Rapid Rotational Motion?

Jakob Struye*, Hany Assasa[†], Barend Van Liempd[†], Arnout Diels[‡], Jeroen Famaey*

*University of Antwerp - imec, Antwerp, Belgium. Email: {firstname}.{lastname}@uantwerpen.be

[†]Pharrowtech, Leuven, Belgium. Email: {hany,barend.vanliempd}@pharrowtech.com

[‡]Dekimo, Leuven, Belgium. Email: {firstname}.{lastname}@dekimo.com

Abstract—Using Millimeter-Wave (mmWave) wireless communications is often named as the prime enabler for mobile interactive Extended Reality (XR), as it offers multi-gigabit data rates at millisecond-range latency. To achieve this, mmWave nodes must focus their energy towards each other, which is especially challenging in XR scenarios, where the transceiver on the user’s XR device may rotate rapidly. To evaluate the feasibility of mmWave XR, we present the first throughput and latency evaluation of state-of-the-art mmWave hardware under rapid rotational motion, for different PHY and MAC-layer parameter configurations. We show that this parameter configuration has a significant impact on performance, and that specialized beamforming approaches for rapid rotational motion may be necessary to enable uninterrupted, high-quality mobile interactive XR experiences.

I. INTRODUCTION

Extended Reality (XR) has found a wide array of applications over the past decade, including in healthcare, education, entertainment and manufacturing [1]. The technical requirements of deployments depend heavily on the exact application. In the most challenging case, extremely-high-quality content is captured or generated in real-time, and transmitted wirelessly to XR-enabled devices, such as head-mounted devices or smartphones. In interactive applications, which require real-time content acquisition, data rates are usually extremely high. Heavy-duty compression cannot be deployed to reduce data rate requirements, as these algorithms are time-consuming, which increases latency. Furthermore, compression is less effective when there is no option of looking ahead to upcoming content. In several types of applications, incorporating a wireless link is essential. This includes applications where content acquisition inherently occurs off-site (e.g., remote conferencing) or when off-loading rendering to an edge cloud enables high-quality virtual experiences even with small, lightweight and silent devices (e.g., gaming). When a satisfactory Quality of Experience (QoE) requires multi-gigabit data rates with latencies in the order of milliseconds along with near-perfect reliability, sub-6 GHz wireless communications are no longer sufficient. Only higher frequencies offer the bandwidth necessary to fulfill these High Rate, High Reliability and Low Latency Communications (HR2LLC) requirements [2]. Of these, Millimeter-Wave (mmWave), spanning 24 to 300 GHz, is most likely to be deployed at a large scale in the near future, as mmWave devices aimed at consumers have been available for some years. Performance figures attainable with

mmWave are highly impressive, with the most recent mmWave amendment for Wi-Fi, IEEE 802.11ay, offering a maximum link rate of over 8.5 Gbps using a single 2.16 GHz channel. Through channel aggregation and Multiple-Input Multiple-Output (MIMO), an Access Point (AP) could theoretically reach over 275 Gbps, divisible among multiple users. Realizing these performance levels does come with a number of challenges not prominent in sub-6 GHz bands. As path and penetration losses increase along with frequency, establishing and maintaining links with sufficiently high signal strength for communications at high Modulation and Coding Schemes (MCSs) becomes highly challenging. The main mitigation technique for this is *beamforming*, a process in which transceivers focus their energy in specific, carefully selected directions. In its most common implementation, the singular antenna is replaced with an *antenna array* consisting of many carefully positioned antenna elements. By applying a different, intelligently selected, phase shift to each element’s signal, the different signals will be phase-aligned, and therefore interfere constructively, in certain directions, but interfere destructively in others. Essentially, this enables focusing a large fraction of the available energy budget in certain desirable directions, effectively increasing the resulting signal strength significantly. One configuration of phase shifts is colloquially called a beam, and beamforming is commonly implemented through *codebooks* containing a pre-defined list of beams. The optimal beam between two devices can be determined by exhaustively sweeping all beams in the codebook.

Despite the first version of the Wi-Fi mmWave amendment, IEEE 802.11ad, being over a decade old, hardware availability and, as a result, performance evaluations, have remained limited. Several works evaluate mmWave-capable Dell hardware, with evaluations on throughput given different distance, Angle of Arrival (AoA) and interference levels [3], [4], on reverse engineering the protocol configuration and beamforming effectiveness [5], and on power consumption [6]. Also frequently evaluated is the TP-Link Talon AD7200, with evaluations on Signal-to-Noise Ratio (SNR) with different beams [7], throughput under blockage and translational motion [8] and performance during mobility-induced blockage [9]. Evaluations with the mmWave-equipped Asus ROG phone include those on throughput under translational motion [10] and on performance in multi-AP scenarios [11]. Finally, several evaluations of experimental 60 GHz testbeds




Fig. 1. The Beacon Interval

investigate the impact of mobility on performance [12]–[14]. Overall, no existing work evaluates throughput and latency *during* rotational motion of an IEEE 802.11ad/ay system with tunable PHY and Medium Access Control (MAC)-layer parameters. In this work, we evaluate the performance of a state-of-the-art mmWave *Evaluation Kit (EVK)*, with a focus on its performance for XR. Specifically, we evaluate throughput and latency under mobility, and evaluate the impact of changing system parameters on performance.

II. IEEE 802.11AD/AY

This work considers hardware implementing the IEEE 802.11ad protocol along with its backwards-compatible successor IEEE 802.11ay, both incorporating mmWave functionality into Wi-Fi. In this section, we provide protocol details along with an overview of our hardware implementing it.

A. Protocol details

As shown in Fig. 1, IEEE 802.11ad/ay divides its transmission schedule into Beacon Intervals (BIs) of configurable length [15]. Each BI starts with a Beacon Header Interval (BHI) reserved for overheads. The BHI starts with a Beacon Transmission Interval (BTI) during which the AP transmits a beacon advertising its existence and capabilities using its different beams. Then, during the Association - Beamforming Training (A-BFT) phase, consisting of different slots, devices may attempt to associate to the AP. This is followed by the optional Announcement Transmission Interval (ATI) for additional signalling. After the BHI, there is a Data Transmission Interval (DTI), intended for actual data transmission. The AP orchestrates the DTI, subdividing it into periods. Each period is either a Contention-Based Access Period (CBAP), during which any node may transmit after a backoff period, or a Service Period (SP), which is reserved for communications between two specific nodes.

B. Hardware

These experiments use two EVKs supplied by Pharrowtech, configured as AP and client respectively. The EVK incorporates Pharrowtech’s SPIRIT PTM1060 module, containing the Radio-Frequency Integrated Circuit (RFIC) and antenna array, with the Renesas RWM6050 baseband processor. The EVK also contains an Intel NUC serving as Network Processor Unit (NPU), running Linux, through which users can configure the other components. The EVK complies with the IEEE 802.11ad/ay standard, and covers its full unlicensed spectrum range, from 57 to 71 GHz. It supports channel bonding of two channels, 64-QAM modulation, an MCS up to 9, and high-resolution phase-shifting.

III. EXPERIMENTS

In these experiments, we mimic a mmWave XR deployment, with one EVK serving as a static AP and another as a mobile user. During the evaluations, both EVKs were positioned on a table, with rotations being performed manually. The two devices were placed 3 m apart horizontally, which is within the range expected in XR deployments. Unless noted differently, we configure the devices to use MCS 9, enable full MAC-layer aggregation, set the transmit power to levels appropriate for indoor usage and a BI of 10 ms. Latency is measured using our open-source `bwdelaytester`¹ tool, configured to send 1000 B packets using UDP at 1.6 Gbps, which does not saturate the link using the default configuration. At the start of each experiment, the EVKs’ clocks are synchronized with sub-microsecond accuracy using Precise Time Protocol (PTP), and each packet contains a generation timestamp, meaning per-packet latency can be measured extremely accurately. We visualize the packet latency with a Cumulative Distribution Function (CDF), and include the packet loss in each plot. Note that, given the real-time nature of the application, we considered packets with a latency over 10 ms to also be lost. Each latency experiment is run for 4 minutes. In some scenarios, we also report the maximal throughput, which was measured in a separate experiment, using saturated UDP `iPerf`² traffic.

In this evaluation, we vary parameters one by one. We investigate the impact of different MCS settings, BI lengths, BHI configurations, channel access schemes, and beam tracking³ approaches. To isolate the impact of each parameter, we perform this first set of experiments in a static scenario, without rotating the user EVK. In a second step, we repeat a subset of well-performing configurations under mobility. We consider three different mobility levels: in the first two, the user performs either moderate or rapid rotational motion while the AP remains within its Field of View (FoV), while for the third level, the user EVK is rotated far enough for the AP to move outside of its FoV.

IV. RESULTS

In this section, we present the results of each experiment. We first evaluate a static deployment, before moving on to mobility experiments.

A. Static Experiments

First, we consider **MCS tuning**. By default, the MCS is at 9, the highest supported value. We also evaluate lower settings, along with automatic rate adaptation. For static indoor scenarios, MCS 9 is easily decodable, meaning we expect a reduction in performance for lower MCS and limited to no impact from enabling automatic rate adaptation. We first evaluate the throughput, being 1.85 Gbps for MCS 9 (and rate adaptation), 1.73 Gbps for MCS 8 and 1.45 Gbps for MCS 7.

¹<https://github.com/arnoutdekimo/bwdelaytester>

²<https://iperf.fr>

³We use “beam tracking” to refer to beamforming when it occurs for an already-associated mobile device




Fig. 2. Latency for different MCS settings




Fig. 4. Latency for different BI lengths




Fig. 3. Actual MCS selected by rate adaptation




Fig. 5. Latency for different BHI configurations (10 ms BIs)

As we evaluate latency at 1.6 Gbps, we do not lower MCS any further. Fig. 2 shows the results of the latency experiments. At MCS 7, the latency is inflated due to the data rate, and MCS 8 still leads to higher latency than MCS 9, due to increased transmission delay at lower throughput. Note that most experiments exhibit a loss between 0.5% and 1%. We investigate the cause of this loss throughout the section. For rate adaptation, Fig. 3 visualizes the actual MCS selected through time, showing that it mostly fluctuates between 7 and 9, with a few brief reductions to 6. We run the remaining experiments at MCS 9 to eliminate the adaptation algorithm as a variable.

Next, we evaluate the impact of the **BI length**, which crucially determines the time between BHIs. During each BHI, no data transmission can occur, meaning that increasing the BI length should reduce latency. On the other hand, fewer BHIs may lead to performance degradation, as this period may be used for beam tracking. However, the default configuration can beam track in the DTI, meaning we expect this impact to be limited. We experiment with 10 ms, 100 ms and 1000 ms BIs, with results in Fig. 4. Clearly, latency is slightly lower with fewer BHIs. This makes sense, as packets are briefly

blocked during the BHI, so more packets will experience this with more frequent BHIs. The three CDFs converge around the 1 ms mark, which is roughly the duration of the BHI. The throughput, instead, experiences a considerable increase when going from 10 ms to 100 ms, jumping from 1.85 Gbps to 1.98 Gbps. The further increase to 1000 ms only improves throughput by another 0.01 Gbps. This makes sense, as the absolute reduction in scheduled BHI time is only one tenth of what it was when going from 10 ms to 100 ms.

By default, the BHI is configured to contain 8 beacons and 1 Sector Sweep (SSW) slot of 16 frames in the A-BFT phase. The specification requires at least 1 beacon per BHI, 1 frame per SSW slot, and an A-BFT only once every 15 BHIs. Using this **shortened BHI configuration** is expected to improve performance in a static scenario [16]. At 10 ms, this results in the latencies shown in Fig. 5, measured with all optimizations enabled. Similar to the previous experiment, the CDF rises more quickly with the shorter BHIs, as blocked packets are released more rapidly with shorter BHI. At higher BI values (100 to 1000 ms) these optimizations had little to no impact, as the fraction of time reserved for BHIs is already significantly lower. Each optimization separately improves throughput by




Fig. 6. Latency for different channel access approaches (10 ms BI)




Fig. 8. Latency for beam tracking in different periods (10 ms BIs)




Fig. 7. Latency for short vs long SP at different BI lengths.




Fig. 9. Repeats of well-performing configurations, with A-BFT beam tracking

0.04 to 0.07 Gbps. All three enabled together results in only a 0.11 Gbps increase, as reducing the length of the A-BFT is less impactful when also reducing its frequency.

Channel access can be configured with either CBAPs or SPs. As the EVKs implement CBAPs on top of SPs, and as most transmissions are from the AP (which does not need to request SPs) we assume the configuration will make little difference. However, with SPs there is also a parameter which controls the maximal SP length. As there are buffer periods between subsequent SPs, increasing this may improve performance. By default, the SP length is 2 ms, with a maximum allowed value of 65 ms. Fig. 6 shows that, as assumed, the difference between CBAP and default SP is negligible, indicating the CBAP is implemented using a 2 ms SP. The impact of the longer SPs is noticeable but limited. However, this was run at a BI of 10 ms, meaning the full 65 ms could not be reached. We therefore repeat the SP measurements at 100 ms and 1000 ms BIs. We compare this to the 10 ms long SP configuration in Fig. 7. While long SPs perform consistently better latency-wise at 10 ms BI, this is no longer the case for higher BI values: the short SP CDF eventually overtakes the long SP's. There may be an optimal

configuration somewhere between 2 ms and 65 ms. We leave optimizing this for future work, as the potential gains appear limited. In addition, we note that these parameters affect throughput significantly: lengthening either the SPs to 65 ms or BI to 100 ms gains 0.12 Gbps and 0.13 Gbps respectively, with the combination gaining 0.27 Gbps, as this allows for the full 65 ms SP to be scheduled. Increasing the BI further to 1000 ms did not markedly improve latency, but did lead to significant fluctuations in throughput.

As a final parameter, we vary where in the BI **beam tracking** is scheduled. By default, beam tracking (for an already-associated station) is scheduled during the DTI, normally intended for data transmission. By beam tracking here, instead of during the A-BFT, the station does not have to wait for the next BHI to perform urgent beam tracking, for example during rapid motion. Fig. 8 shows that latency improves somewhat by disabling DTI beam tracking. More importantly, packet loss is reduced by two orders of magnitude. Especially in an interactive XR scenario, this is a massive improvement. We suspect that the firmware attempts to transmit packets during this DTI beam tracking and does not reschedule these properly. While this is an apparent improvement, note that the experiments




Fig. 10. Latency under moderate motion




Fig. 12. Latency under rapid motion




Fig. 11. Latency for different BHI configurations, under moderate motion

presented so far were for fully static environments, meaning the true impact of disabling DTI beam tracking remains to be evaluated.

Before proceeding to the mobility experiments, we repeat a selection of the previous experiments with DTI beam tracking disabled, with Fig. 9 showing the results. Differences in performance for previously evaluated parameters are now less pronounced. 100 ms BIs perform best, with long SPs having slightly lower packet loss. The slightly higher loss with these parameters is negligible. The lowest-loss configuration, with regular SPs, saw 99.61 % of all packets arrive within 1 ms and 99.92 % within 2 ms. Less than 0.0028 % of packets arrived with higher latency, and 0.075 % were lost.

B. Mobility Experiments

The mobility patterns in these experiments were chosen to mimic mobile XR, meaning the transmitter is always static. Firstly, the receiver rotates at moderate speed. The receiving EVK was rotated continuously (alternating between clockwise and counterclockwise), completing a 90° turn every 2 s. The AoA and broadside were aligned at the midpoint of the rotation, meaning the AP stays within the XR device’s FoV. Fig. 10 shows latencies for possibly viable configurations.

Notably, a short BI with A-BFT beam tracking performs best, with a packet loss of only 0.41 %. DTI beam tracking is meant to speed up beam tracking, but clearly, at moderate rotational speeds, the BHI alone can offer sufficiently responsive beam tracking. The DTI beam tracking only hinders data transmission even in this case. At larger BI, non-negligible loss occurs even with DTI beam tracking (though it is difficult to tell if this occurs due to collisions or link degradation). Looking at latencies, the 10 ms configuration without DTI beam tracking, despite its CDF rising more slowly than others’, has the highest percentage of packets arrive rapidly, with 97.32 % within 1 ms and 99.55 % within 2 ms. 0.037 % arrived with a higher latency, and 0.41 % were lost. For the next experiments, we use a 10 ms BI and A-BFT beam tracking.

Then, we repeat the BHI length optimization experiments, which improved latency using a 10 ms BI in the static case. This gives the latencies in Fig. 11. Decreasing the number of beacons has a slight positive effect on latency but increases loss somewhat, while reducing the A-BFT frequency has a serious impact on reliability. This likely occurs because the EVK needs to wait significantly longer before it can beam track, causing queues to fill up and latencies to increase to over half a second.

As a final experiment with rotations within the FoV, we consider fast rotations. We repeat the best-performing configuration, now rotating the EVK as rapidly as safely possible every 2 s, leaving the EVK in the final position until the next rotation. Fig. 12 shows that, while the overall latency distribution is similar, the loss increases five-fold with faster rotations, indicating the beam tracking algorithm does not always adapt in time.

Finally, we evaluate what happens when the AP leaves the XR device’s FoV. In this extreme scenario, the device is turned 180° every 2.5 s, with the rotation itself taking 1.5 s, meaning the antenna arrays are perpendicular to each other around 40 % of the time, making successful reception extremely challenging. As Fig. 13 shows, every considered configuration performs too poorly to be usable. While there are significant differences between the configurations, re-




Fig. 13. Latency under extreme motion

establishing a broken link sometimes took multiple seconds, meaning the only valid conclusion from this experiment is that more than 1 AP is needed to reliably cover a highly mobile user.

V. CONCLUSIONS

In this work, we evaluated the performance of state-of-the-art IEEE 802.11ad/ay hardware. We measured latency distribution and throughput in both static and mobile scenarios, with a focus on rapid rotational motion expected in XR scenarios. In addition, we measured the impact of modifying several PHY and MAC-layer parameters. By tuning parameters related to the structure of the BI, we were able to increase throughput from 1.85 Gbps to 2.12 Gbps, a nearly 15% increase. This configuration, with a 100 ms BI, A-BFT beam tracking and long SPs, exhibits a loss around 0.4% and near-consistent sub-2 ms packet latency for arriving packets, with shorter SPs reducing packet loss to under 0.1%, but reducing throughput by 0.14 Gbps. Under moderate motion, we maintained a loss of around 0.4% by reducing the BI length to 10 ms. This allowed for more frequent beam tracking, but reduced throughput back to 1.85 Gbps. Fast motion increased this loss to 2%, while extreme motion, with the AP leaving the XR device's FoV, increased loss to 35%. Overall, these experiments show that, while mmWave is a promising enabler for mobile interactive XR, faster or even proactive beam tracking solutions for mobile nodes are needed. Furthermore, solutions such as multi-AP deployments, Reconfigurable Intelligent Surface (RIS) [17] and distributed antenna systems [18] can increase coverage. Even then, the streaming protocol will need to be robust to some packet loss.

ACKNOWLEDGMENT

This research was partially funded by the ICON project INTERACT and Research Foundation - Flanders (FWO) project WaveVR (Grant number G034322N). INTERACT was realized in collaboration with imec, with project support from VLAIO (Flanders Innovation and Entrepreneurship). Project partners are imec, Rhinox, Pharrowtech, Dekimo and TEO.

REFERENCES

- [1] L. Adriana Cárdenas-Robledo, Óscar Hernández-Urbe, C. Reta, and J. Antonio Cantoral-Ceballos, "Extended reality applications in industry 4.0. – a systematic literature review," *Telematics and Informatics*, vol. 73, p. 101863, 2022.
- [2] C. Chaccour, M. N. Soorki, W. Saad, M. Bennis, and P. Popovski, "Can terahertz provide high-rate reliable low latency communications for wireless vr?" *IEEE Internet of Things Journal*, pp. 1–1, 2022.
- [3] T. Yamada, T. Nishio, M. Morikura, and K. Yamamoto, "Experimental evaluation of ieee 802.11ad millimeter-wave wlan devices," in *2015 21st Asia-Pacific Conference on Communications*, 2015, pp. 278–282.
- [4] A. Loch, G. Bielsa, and J. Widmer, "Practical lower layer 60 ghz measurements using commercial off-the-shelf hardware," in *Proceedings of the Tenth ACM International Workshop on Wireless Network Testbeds, Experimental Evaluation, and Characterization*, 2016, p. 9–16.
- [5] 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 *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies*, 2015.
- [6] S. K. Saha, T. Siddiqui, D. Koutsonikolas, A. Loch, J. Widmer, and R. Sridhar, "A detailed look into power consumption of commodity 60 ghz devices," in *2017 IEEE 18th International Symposium on A World of Wireless, Mobile and Multimedia Networks*, 2017, pp. 1–10.
- [7] D. Steinmetzer, D. Wegemer, M. Schulz, J. Widmer, and M. Hollick, "Compressive millimeter-wave sector selection in off-the-shelf ieee 802.11ad devices," in *Proceedings of the 13th International Conference on Emerging Networking Experiments and Technologies*, 2017, p. 414–425.
- [8] S. K. Saha, H. Assasa, A. Loch, N. M. Prakash, R. Shyamsunder, S. Aggarwal, D. Steinmetzer, D. Koutsonikolas, J. Widmer, and M. Hollick, "Fast and infuriating: Performance and pitfalls of 60 ghz wlns based on consumer-grade hardware," in *2018 15th Annual IEEE International Conference on Sensing, Communication, and Networking*, 2018, pp. 1–9.
- [9] J. Struye, H. K. Ravuri, H. Assasa, C. Fiandrino, F. Lemic, J. Widmer, J. Famaey, and M. T. Vega, "Opportunities and challenges for virtual reality streaming over millimeter-wave: An experimental analysis," in *2022 13th International Conference on Network of the Future (NoF)*, 2022, pp. 1–5.
- [10] S. Aggarwal, A. Thirumurugan, and D. Koutsonikolas, "A first look at 802.11ad performance on a smartphone," in *Proceedings of the 3rd ACM Workshop on Millimeter-Wave Networks and Sensing Systems*, 2019, p. 13–18.
- [11] S. Aggarwal, M. Ghoshal, P. Banerjee, D. Koutsonikolas, and J. Widmer, "802.11ad in smartphones: Energy efficiency, spatial reuse, and impact on applications," in *IEEE INFOCOM 2021 - IEEE Conference on Computer Communications*, 2021, pp. 1–10.
- [12] S. Aggarwal, U. S. Sardesai, V. Sinha, and D. Koutsonikolas, "An experimental study of rate and beam adaptation in 60 ghz wlns," in *Proceedings of the 23rd International ACM Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems*, 2020, p. 171–180.
- [13] J. O. Lacruz, D. Garcia, P. J. Mateo, J. Palacios, and J. Widmer, "mm-flex: an open platform for millimeter-wave mobile full-bandwidth experimentation," in *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services*, 2020, p. 1–13.
- [14] J. O. Lacruz, R. R. Ortiz, and J. Widmer, "A real-time experimentation platform for sub-6 ghz and millimeter-wave mimo systems," in *Proceedings of the 19th Annual International Conference on Mobile Systems, Applications, and Services*, 2021, p. 427–439.
- [15] T. Nitsche, C. Cordeiro, A. B. Flores, E. W. Knightly, E. Perahia, and J. C. Widmer, "Ieee 802.11ad: directional 60 ghz communication for multi-gigabit-per-second wi-fi [invited paper]," *IEEE Communications Magazine*, vol. 52, no. 12, pp. 132–141, 2014.
- [16] J. Struye, F. Lemic, and J. Famaey, "Towards ultra-low-latency mmwave wi-fi for multi-user interactive virtual reality," in *GLOBECOM 2020 - 2020 IEEE Global Communications Conference*, 2020, pp. 1–6.
- [17] E. Basar, M. Di Renzo, J. De Rosny, M. Debbah, M.-S. Alouini, and R. Zhang, "Wireless communications through reconfigurable intelligent surfaces," *IEEE Access*, vol. 7, pp. 116 753–116 773, 2019.
- [18] A. Moerman, J. Van Kerrebrouck, O. Caytan, I. L. de Paula, L. Bogaert, G. Torfs, P. Demeester, H. Rogier, and S. Lemey, "Beyond 5g without obstacles: mmwave-over-fiber distributed antenna systems," *IEEE Communications Magazine*, vol. 60, no. 1, pp. 27–33, 2022.