Received XX Month, XXXX; revised XX Month, XXXX; accepted XX Month, XXXX; Date of publication XX Month, XXXX; date of current version 11 January, 2024.

Digital Object Identifier 10.1109/OJCOMS.2024.011100

XR Streaming Performance with Wi-Fi 7 Multi-Link Operation

Marc Carrascosa-Zamacois[◇⋆♭], Lorenzo Galati-Giordano[♭], Francesc Wilhelmi[⋆], Gianluca Fontanesi[♭], Anders Jonsson[⋆], Giovanni Geraci[♯⋆], and Boris Bellalta[⋆]

[◇] ***Centre Tecnològic de Telecomunicacions de Catalunya, Castelldefels, Spain***
[♭] ***Nokia Bell Labs, Stuttgart, Germany***
[⋆] ***Universitat Pompeu Fabra, Barcelona, Spain***
[♯] ***Nokia Standards, Spain***

ABSTRACT Extended Reality (XR) has stringent throughput and delay requirements that are hard to meet with current wireless technologies. Missing these requirements can lead to worsened picture quality, perceived lag between user input and corresponding output, and even dizziness for the end user. In this paper, we study the capability of Wi-Fi 7, and its novel support for Multi-Link Operation (MLO), to cope with these tight requirements. Our study is based on simulation results extracted from an MLO-compliant simulator that realistically reproduces VR traffic. Results show that MLO can sustain VR applications. By jointly using multiple links with independent channel access procedures, MLO can reduce the overall delay, which is especially useful in the uplink, as it has more stringent requirements than the downlink, and is instrumental in delivering the expected performance. We show that using MLO can support more users per Access Point (AP) than an equivalent number of links using Single Link Operation (SLO). We also show that, while maintaining the same overall bandwidth, a higher number of MLO links with narrower channels leads to lower delays than a lower number of links with wider channels. We also study the impact of Overlapping Basic Service Sets (OBSS) on performance of XR applications, showing how SLO struggles to maintain a low delay when OBSS activity increases, and how MLO can cope with this interference by using MLO-aware channel assignment strategies. Finally, we consider random positioning for the users, showcasing that MLO can support 5× more users than SLO.

INDEX TERMS Multi-Link Operation, XR streaming, IEEE 802.11be, Wi-Fi 7

I. Introduction

Extended Reality (XR) applications, which include Virtual Reality (VR) and Augmented Reality (AR), are growing in popularity as they unlock novel use cases across many domains, such as healthcare, industry, education and gaming. Most use cases are planned for indoor use, and thus Wi-Fi is expected to become the main technology to support them [1]–[4], with most headsets including high-grade Wi-Fi capabilities[1], and services like *Steam Link*[2] allowing to stream games wirelessly from a computer to a headset. To deliver a good performance to the end user, XR traffic has stringent requirements in both application-level throughput, which can go over 100 Mbps, and delay, which needs to be well below 10 ms [5].

Wi-Fi currently struggles to provide delay guarantees. Wi-Fi's operation at the MAC layer is based on distributed channel access due to its operation in the unlicensed spectrum and the inherent requirement of using Listen Before Talk (LBT). More specifically, the contention among devices that share the same frequency channels has a direct impact on the delay experienced by the users, which deteriorates as the number of contenders increases. Further, XR applications have stricter requirements for the uplink (UL) delay, which is harder to control by the Access Point (AP) due to the spontaneous nature of this type of traffic.

In this paper, we focus on realistic VR streaming applications, and study Wi-Fi 7's ability to effectively meet their stringent requirements. We study the performance of VR traf-

[1] https://www.meta.com/help/quest/articles/headsets-and-accessories/oculus-link/connect-with-air-link/

[2] https://store.steampowered.com/app/353380/Steam_Link/

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

fic on Wi-Fi 7 networks, which we test for both Single Link Operation (SLO) and Multi-Link Operation (MLO). Using simulation results, we showcase the relationship between different transmission parameters, namely the Modulation Coding Scheme (MCS) and channel bandwidth, and provide insights on their required configuration for achieving the desired performance for VR applications. From that basis, we then study the limits of Wi-Fi 7 in terms of the number of supported VR streams. This paper builds upon the work presented in [6], which we expand by providing a broader set of validation insights—we consider randomly placed VR users instead of static positions, we study the implications of using narrower versus wider channels, and we consider Overlapping BSS (OBSS) competition to study how neighboring Wi-Fi MLO activity impacts network latency, potentially preventing the use of VR.

Our main contributions are as follows:

- We provide an overview of VR gaming traffic based on real traces, and analyze the associated requirements defined by the Wi-Fi Alliance (WFA).
- We conduct an extensive performance evaluation based on several simulations and discuss the feasibility of Wi-Fi 7 for supporting VR traffic. We show that the number of VR users with MLO and N links is higher than N independent SLO APs.
- We show that for MLO, multiple narrow channel links provide a better support for VR applications than a few but wider channel links, as the channel access delay—the most limiting factor—is significantly scaled down. However, when users are placed randomly, this shows diminishing returns as the MCS worsens and transmission times increase, thus additional complexity and hardware costs associated to handle a large number of independent links may outweigh the performance benefits.
- We study the impact of OBSS activity on VR streams, showing that SLO devices cannot maintain the necessary latency for VR while contending with other APs, which is something that MLO can address thanks to a better management of links.

The rest of the paper is organized as follows: Section II provides an overview of the related work. Section III introduces the VR setup considered, as well as a description of VR traffic obtained from traces and the throughput and delay requirements considered throughout this paper. Section IV shows the simulation system model and scenarios considered. The performance analysis of VR over a single MLO BSS is presented in V, while the results obtained with multiple BSS, as well as coexistence with SLO BSSs an be found in Section VI. Finally, Section VII closes the paper.

II. Related Work

A. Wi-Fi 7

IEEE 802.11be (Wi-Fi 7) [7] is soon to be amended into the main IEEE 802.11 standard, and devices including its features are already appearing in the market. Wi-Fi 7, also named Extremely High Throughput (EHT), brings several improvements over Wi-Fi 6, such as wider channels going up to 320 MHz, an additional modulation scheme of 4096-QAM, and the ability to use multiple OFDMA Resource Units per user [8], [9]. The most disruptive feature of Wi-Fi 7 however is a new addition to Wi-Fi: Multi-Link Operation (MLO), which allows devices to connect to multiple channels through a single association and to transmit packets simultaneously over them, thus multiplying the available bandwidth by the number of radios on a device. Medium access is also independent for each radio, leading to more transmission opportunities and reduced contention as well [10]. MLO has received a lot of attention in the last years: An overview of MLO and its different implementations can be found in [11], where the potential throughput gains is studied. In terms of latency, the work in [12] showed that adding a second link to traditional Wi-Fi can lead to order of magnitude gains in the 90th percentile delay for real-time applications. In [13], a dataset using real-world channel occupancy traces was used to test MLO performance, also showing a similar order of magnitude improvement over SLO in the 95th percentile delay. A survey focusing on the extensive research done on MLO can be found in [14].

B. Wi-Fi and XR applications

Wi-Fi's capability to support XR applications has been tested in [15], comparing the performance of wired and wireless deployments, and showing that, while Wi-Fi can achieve similar results than a wired connection, this only happens for devices that are close to the AP and with direct line of sight. In [16], user experience was studied for wired and wireless VR setups, also highlighting that a direct line of sight is necessary to ensure comparable Quality of Experience (QoE) between wired and wireless setups. In [17], [18], a setup with multiple VR users over Wi-Fi was studied, showing that scheduling the uplink transmissions (i.e., using OFDMA to improve multi-user contention) leads to worsened performance overall than just using DCF. They also propose several configuration options for split-streaming multi-user setups, including encoding bitrate, frame rater, channel bandwitdth, and RSSI, showing that a high RSSI is necessary to ensure good performance, as a single user with a low datarate worsens the performance for all users. Different types of XR applications were studied and classified in [19], [20], as well as their requirements and the possibility to cover them based on the current efforts done in 5G and Wi-Fi standardization. In [21], a performance evaluation model was proposed for edge-assisted XR applications, considering battery usage, end-to-end delay and handoff delays. In [22], the authors

analyzed a multi-user VR setting deployed over Wi-Fi, and concluded that Quality of Service (QoS) enforced by the standard Enhanced Distributed Channel Access (EDCA) is insufficient to support the stringent delay and packet loss requirements of such a setting. They then proposed a new architecture to improve performance by separating the downlink and uplink using the 802.11ad/ay 60 GHz band and the 802.11ax 5 GHz band, respectively.

C. Multi-Link Operation for XR applications

An early discussion on the use of Virtual Reality over 802.11ac Wi-Fi showed the inadequacy of Wi-Fi to meet the necessary requirements, and proposed the use of multiple Wi-Fi radios using Multi-Path TCP to reduce overall latency [23]. The performance of Wi-Fi 7 for AR applications was studied in [24], concluding that MLO can serve more users than SLO with equivalent bandwidth. In [25], the authors the use of MLO as an enabler of extended reality applications, using the secondary links for redundancy and seamless handovers to ensure uninterrupted service. To the best of our knowledge, this is the first study to focus on how MLO can enable VR traffic with OBSSs interference.

III. VR Gaming in Wi-Fi 7

A. VR Streaming Setup

We describe here the VR streaming setup used to obtain VR traffic traces, which we reproduce in our simulations to carry out the performance analysis over Wi-Fi 7. The setup includes a server connected to an AP directly via a 1 Gbps Ethernet link, and a Head Mounted Display (HMD) is connected to the AP through Wi-Fi. Fig. 1 shows the main components of VR streaming for gaming. A detailed description of the setup and applications used can be found in [17].

Streaming VR gaming applications offload the main processes to a server, taking the computational load of rendering the video, audio, and any other necessary data away from the HMD. The rendered video and audio are then transmitted through the Internet to the client HMD, which reproduces the incoming video for the user and captures the user inputs, which are then sent back to the server. This approach is also known as split-rendering VR.

B. VR Traffic Distribution

To analyze and replicate VR traffic, we use the traces from [17], which can be found in Zenodo as a dataset [26]. They were obtained using Air Light VR (ALVR), which was installed in both the server and the HMD. ALVR allows streaming VR games over Wi-Fi, as well as gives the user control over several stream settings, such as resolution, refresh rate, codec used, and transport protocol (UDP or TCP). ALVR creates a bridge between the server and the HMD, transmitting audio, video, and tracking. Video is compressed on the server using either the H.264 or H.265

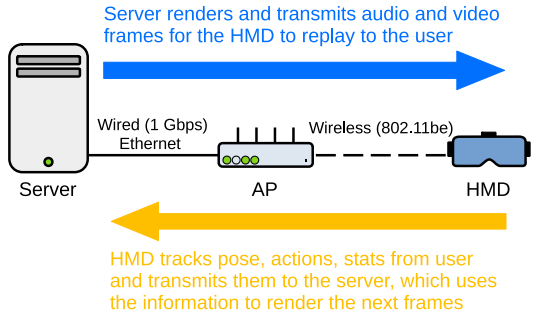

FIGURE 1: VR streaming components.

TABLE 1: Reliability and delay requirements defined by the WFA for VR gaming.

Type of traffic stream	Required reliability (percentile)	Maximum delay
DL: Video frames	75th	5 ms
	95th	10 ms
	99.9th	50 ms
UL: Pose, IMU, controller inputs	90th	2 ms
	99.9th	10 ms

codecs. Audio is sent raw using Pulse-Code Modulation (PCM).

Tests were performed at different resolutions and refresh rates, using H.264 coding and UDP for the transport protocol. Wireshark was used to capture the traffic on the server. These captures, which have been replicated in our simulations, reflect the generation patterns for VR gaming. We describe the trace patterns here, and the simulation traffic patterns are fully described in Section IV A.

- **Downlink traffic:** There are two types of packets for the downlink (DL): video and audio. Video is transmitted in batches of packets, separated as a function of the frame rate. In this case, it is 90 frames per second (FPS), which leads to an interval of 11.11 ms between batches. At a 100 Mbps application rate, each batch contains an average of 96 video packets of 1448 bytes. The audio is transmitted at a different rate (25 ms) and in batches of 4 packets.
- **Uplink traffic:** In the uplink (UL), we have information about tracking, pose, and stream statistics. They also follow the video frame rate, with 3 packets of size 106 bytes per video frame, and a single one of 212 bytes, which we believe accounts for the pose and stats, respectively.

C. Throughput and Delay Requirements

VR content has strict throughput and delay requirements. It requires real-time rendering, meaning that a large amount of data needs to be transmitted constantly and consistently, and

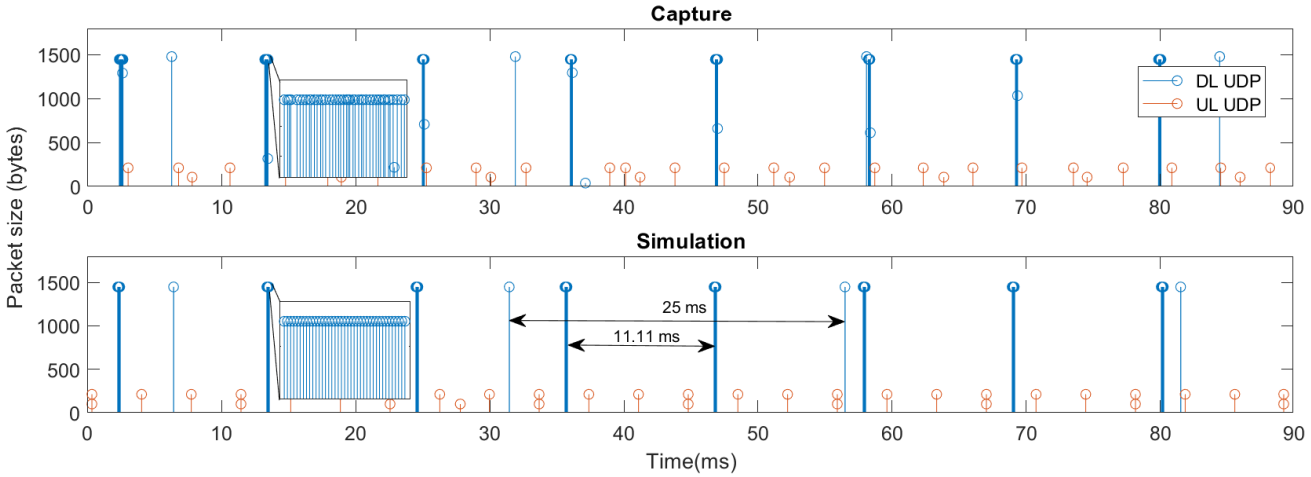

FIGURE 2: VR traffic distribution obtained from the capture from [17] (top figure) and the simulation (bottom figure).

buffering is not possible. VR frame rates are high, ranging from 72 FPS to 120 FPS. This frame rate sets the pace at which traffic is generated, and so the higher the quality of the stream, the more frequent the transmissions. Video quality is also affected by its bitrate, which can range from 40 to 200 Mbps. The delay is particularly important, not only to deliver a good video experience, but to avoid that the user suffers dizziness. Finally, the rendered video transmitted in the downlink changes based on the inputs of the user, which are delivered in the uplink. In this work, we will look at the delay requirements set by the Wi-Fi Alliance (WFA) for VR gaming [27], which we summarize in Table 1.

IV. System Model

A. VR Traffic Characterization

VR traffic is mostly periodic, defined mainly by its total traffic load and the frame rate used by the VR application. As feedback from the HMD is continuously transmitted to the server, a constant video bitrate is generated without exception. We match these main characteristics in our simulation: the frame rate ϕ sets the inter-arrival time Δ of the video data, with downlink packet batches separated by $\Delta = \frac{1}{\phi}$ secs. The video bitrate ω sets the size of the downlink video batches (in packets per batch), which corresponds to $N_{\text{batch}} = \Delta\frac{\omega}{\text{L}}$, where L is the video packet size. A comparison between the traces and our simulator output can be found in Fig. 2, showing an accurate representation of the real traffic.

B. Channel Access

We consider two main modes of operation (represented in Fig. 3):

- **Single Link Operation (SLO):** Legacy Wi-Fi operation, used as our baseline. APs and STAs connect through a single link and then perform backoff to access the medium. Packets are transmitted sequentially.

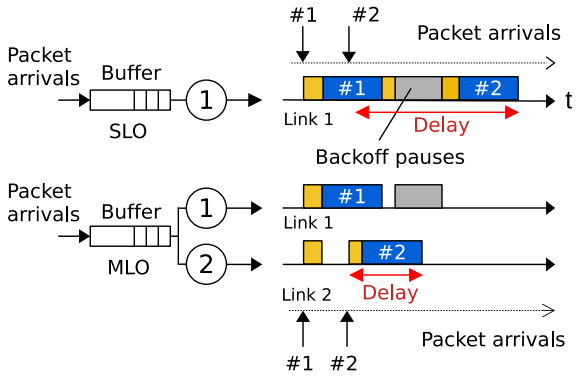

FIGURE 3: Channel access modes: SLO (top) and MLO (bottom). MLO allows simultaneous transmission, reducing the total delay for packet #2.

- **Multi-Link Operation STR (MLO):** MLO Simultaneous Transmit and Receive (STR)[3] allows APs and STAs to use multiple channels at the same time. Each link uses an independent backoff timer, thus packets can be transmitted opportunistically through both links. Fig. 3 shows packet #2 arrives at the buffer once packet #1 is in the middle of being transmitted through link 1. Backoff is then performed on the second link, and the packet is transmitted at the same time as packet #1, reducing the delay in comparison to the sequential transmission in SLO.

C. Scenarios

We first consider a single BSS scenario, analysing VR performance in isolation and studying configurations that enable the highest number of users. We then move on to coexistence scenarios, first studying the impact of MLO in

[3]In the remainder of the paper, we use MLO to refer to MLO STR.

TABLE 2: Simulation parameters.

Name	Value
Channel bandwidth	20, 40, 80, 160, 320 MHz
Transmission power	23 dBm
Clear channel assessment	−82 dBm
Spatial streams	2
Packet Error Rate (PER)	0.1
Max. A-MPDU size	1024 packets
Buffer size	5000 packets
Number of random drops	100
Simulation time	10 seconds

VR performance with external SLO influence, and finally analysing a coexistence scenario where every BSS is using MLO.

We consider Wi-Fi 7 modulation and coding schemes (up to 4096-QAM), and path loss in the 5 GHz band that is modeled considering the 802.11ax residential scenarios [28]. Simulation parameters can be found in Table 2. The main scenarios are as follows:

- **Single BSS:** a single Wi-Fi network composed of an AP and K VR stations. The VR server has a direct wired connection to the AP.
- **Single MLO and two SLO BSSs:** Three inline APs as depicted in Fig. 4, with a central MLO VR BSS surrounded by two SLO BackGround (BG) BSSs. Each of the SLO BSSs shares a channel with the MLO BSS. The VR AP has a VR STA that can be placed randomly in a 10 meters radius, while the BG APs have an STA at a fixed position to ensure the same data rate.
- **Two MLO BSSs:** Two MLO VR BSSs with the same coverage area that can either share both channels (Fig. 5a) or share only one channel while having an independent orthogonal channel each (Fig. 5b).

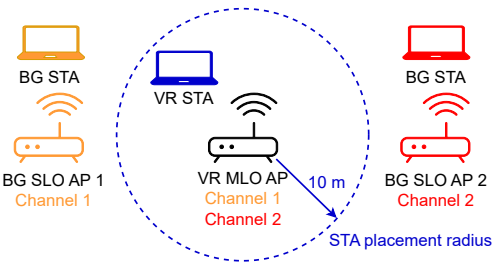

FIGURE 4: A single VR AP sharing a link with two other BG SLO BSSs.

D. Performance Metrics

The main metrics considered in the following sections are the latency constraints given by the WFA (see Section III C and Table 1), so we consider the 75th, 95th and 99.9th percentiles of the delay for the downlink, and the 90th and 99.9th percentile of the delay for the uplink. When random user placement is considered, we focus on the simulation results where the buffers remain below full capacity, avoiding a mix of unsaturated and saturated network conditions. Nevertheless, we also include the fraction of unsuccessful deployments to show how often the simulated load causes packet losses. In any case, the buffer occupancy ρ is useful to understand the contention of the OBSS scenarios (a higher ρ indicates higher contention).

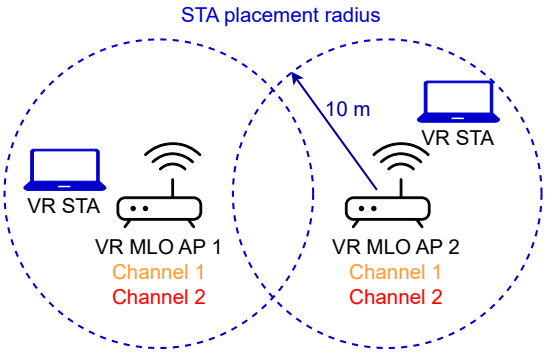

(a) Both links are shared

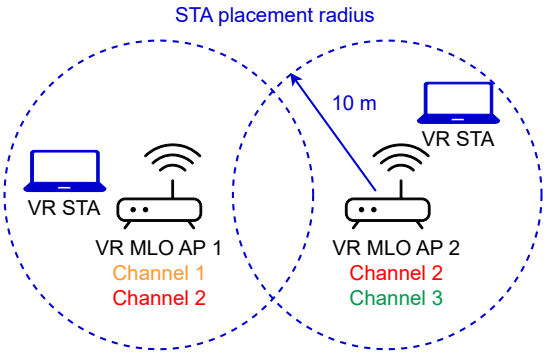

(b) Only one link is shared

FIGURE 5: Multiple MLO BSS VR scenarios.

V. Results for a Single BSS

This section studies the performance of VR applications, their requirements in terms of delay and focuses on how to enable multiple VR users in a network. A single BSS is considered here to minimize the amount of random variables affecting performance, while OBSS interference and coexistence will be studied in the next section.

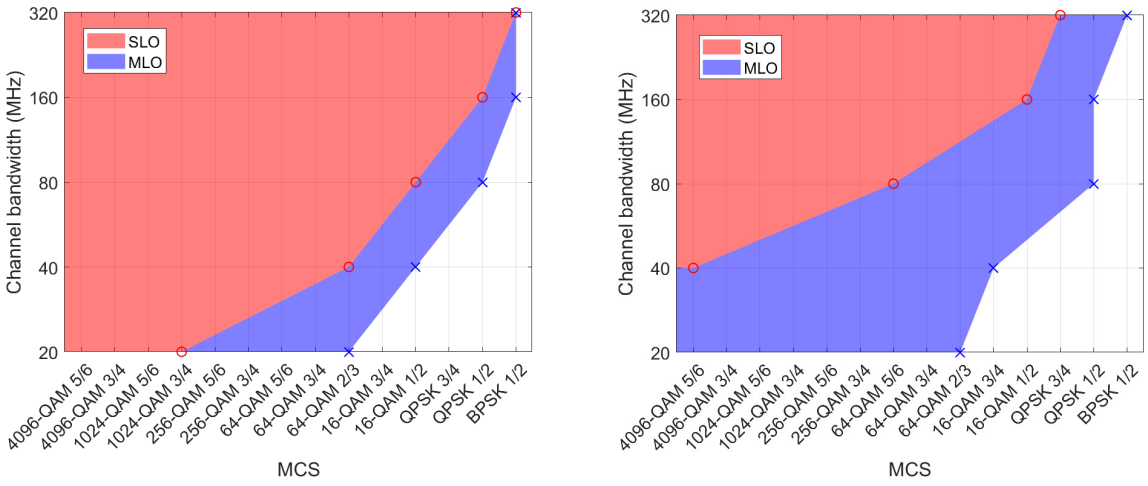

(a) Downlink (75%-tile of 5 ms)

(b) Uplink (90%-tile of 2 ms)

FIGURE 6: Minimum MCS to meet WFA requirements for different channel widths.

A. Required MCS and Channel Bandwidth

We start our study by verifying the minimum combination of MCS and channel bandwidth required to ensure a good experience to the end user according to the WFA specification for XR gaming [27] (see also Table 1). In our analysis, we compare Wi-Fi 7 MLO with respect to SLO. We consider a single AP and STA operating at different channel bandwidths, from 20 MHz to 320 MHz. For each bandwidth, different MCS values are evaluated to find the combinations that meet a good user experience.

Fig. 6 compares the results obtained by Wi-Fi 7 MLO and SLO devices for both downlink (Fig. 6a) and uplink (Fig. 6b). For SLO, the uplink is far more restrictive than the downlink, and a minimum bandwidth of 40 MHz and an extremely high MCS, i.e., 4K QAM with coding rate 3/4, are required. In the downlink, VR requirements can be met even with 20 MHz and 1024-QAM with coding rate 3/4. For the rest of bandwidth configurations, the UL in SLO requires much higher MCS values than the downlink, and even for 320 MHz, the lowest MCS cannot achieve the 2 ms requirement at 90th percentile. On the contrary, for MLO, the uplink and downlink resemble each other much more, generally requiring similar MCS values, which are also lower than those required by SLO.

Takeaway: The UL requirements are harder to meet than those for the DL despite there being a much lower UL traffic. MLO offers a clear advantage over SLO even in situations where the total used bandwidth is comparable with SLO (e.g., SLO using 160 MHz and MLO using two links of 80 MHz), as MLO implicitly relieves the self-contention created by both the UL and the DL by taking advantage of transmitting over multiple links. In addition, MLO provides a higher flexibility than SLO because it requires a lower MCS for the same total bandwidth. This is very useful to support VR streaming under more challenging propagation conditions, such as being far away from the AP or not in direct line of sight.

B. Increasing Number of VR Users

We now study the capacity of MLO to serve a certain number of VR streams and compare it to SLO. We set a single AP transmitting multiple VR streams of 100 Mbps and 90 FPS to end users. All STAs have the same MCS of 1024-QAM 5/6, and 80 MHz channels are used on all links.

Fig. 7a shows the 75th, 95th and 99.9th percentiles (shown as stacked bars) of the packet delay suffered by DL traffic, for both SLO (in blue) and MLO (in red), and for an increasing number of VR streams. The dashed lines highlight the DL delay thresholds for each percentile (as defined in Table 1). We can observe that SLO's delay increases much faster than MLO's, allowing three streams before exceeding the 75th percentile threshold of 5 ms (black dashed line). MLO comfortably allows up to six streams, and offers delay improvements at lower loads (e.g., for three streams, SLO has a 99.9% delay of 7.3 ms, while MLO has a delay of 6 ms for six streams). Generally, the number of extra streams that can be added with MLO is directly proportional to the extra number of links enabled. Notably, in all cases, MLO guarantees a lower delay in both DL and UL compared to SLO while supporting twice the number of VR streams, showing that the opportunistic nature of MLO offers a slight improvement in network delay.

Just like for the DL, Fig. 7b shows the 90th and 99.9th percentiles of the packet delay for the UL. It can be observed that, for SLO, the UL does not support more than two users, exceeding the 2 ms threshold for the third user, once again showing that the UL limits SLO connections. In contrast,

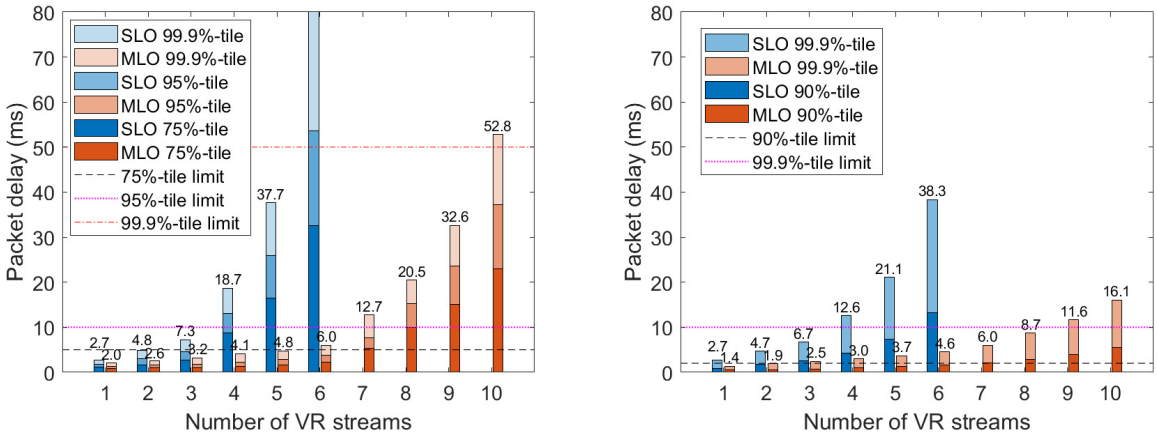

(a) Downlink packet delay

(b) Uplink packet delay

FIGURE 7: Packet delay as number of users increases. The straight lines indicate the respective WFA requirements for DL and UL.

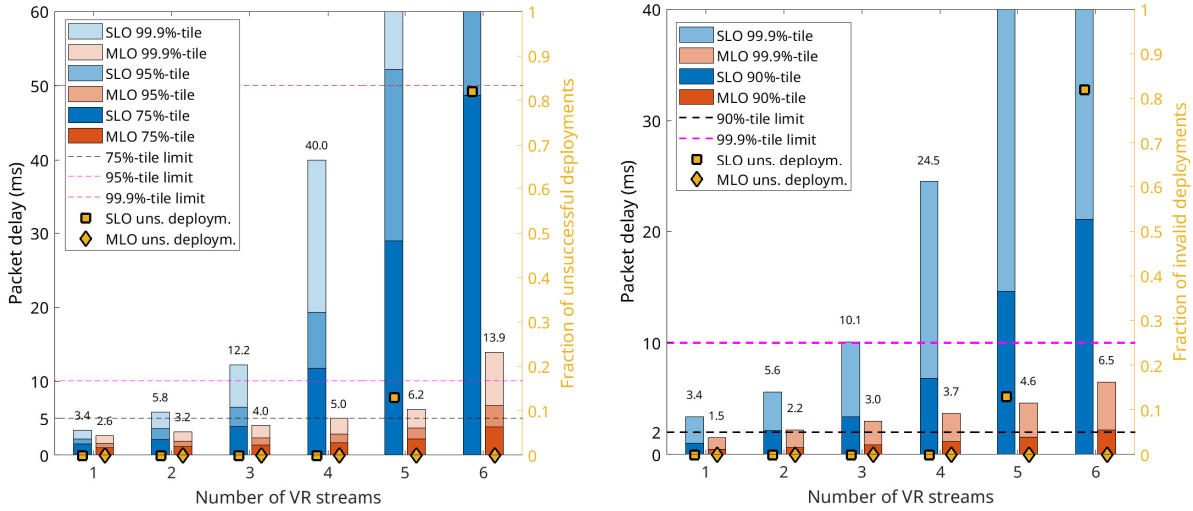

(a) Downlink packet delay

(b) Uplink packet delay

FIGURE 8: Packet delay and fraction of unsuccessful deployments versus the number of randomly placed STAs in the considered deployment.

MLO can sustain six users, the same number as in the DL. Additionally, the 99.9th percentile delay for six MLO users is 4.6 ms, which is lower than the 4.7 ms achieved by SLO with two users. This indicates that, even if we had two SLO networks to match the bandwidth, SLO would only support four VR users, while MLO allows for an additional 50%.

Takeaway: MLO provides an advantage over SLO and allows supporting more VR users. This is a consequence of the increase in the number of independent channel access instances over the available links, rather than the increased bandwidth, as we have shown that using two MLO links leads to more users than using two SLO links.

C. Randomly placed VR users within the AP Coverage Area

We replicate the previous methodology, but now focusing on random user placement. STAs are placed randomly within a radius of 10 meters around the AP, guaranteeing the necessary MCS for good SLO performance with 80 MHz links according to our results in Fig. 6b.

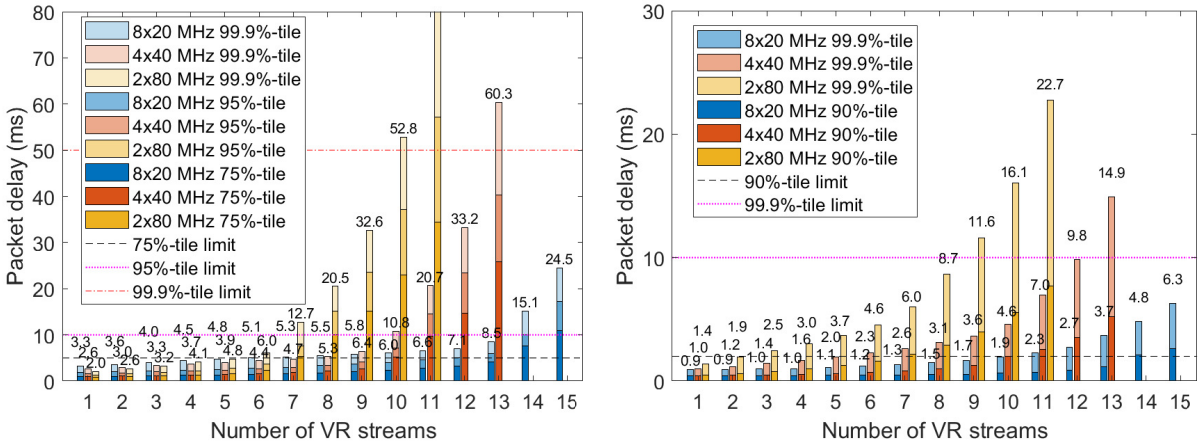

(a) Downlink

(b) Uplink

FIGURE 9: Packet delay for different configurations of links and bandwidth.

Fig. 8a shows the DL delay as the number of XR streams increases, with the fraction of unsuccessful deployments for each scenario (i.e., simulations in which packets are lost because the buffer was full). Now that users are placed randomly, we observe a general increase in the delay, as well as the network having issues earlier than in Fig. 7a. In particular, four SLO users lose packets in 13% of simulations, while in the previous section we showed that up to six users could be properly supported. Fig. 8b shows the UL delay, where we find that SLO can fulfill the delay requirements of only one user (as opposed to the two users supported previously), and that MLO cannot meet the requirements for six users anymore.

Takeaway: SLO suffers from the instability generated by considering different user positions and their consequent diversity in MCS values employed. SLO can barely support one VR user, and despite MLO's performance decreases compared to the case where fixed locations are assumed, it can still enable 5× more users than SLO.

D. Bespoke MLO for VR Applications

We now attempt to further increase the number of VR users in the network by reducing contention. For that, we distribute the same bandwidth over a different number of links, thus increasing the opportunities for streams to be transmitted in parallel. We define MLO configurations of two, four and eight links, each one using 80 MHz, 40 MHz and 20 MHz channels, respectively. By doing this, we maintain the same overall bandwidth (160 MHz in total) used across all the experiments, but we spread it over an increasing number of links.

Fig. 9a shows the DL packet delay for all configurations. We can observe that adding more links, even if each has a lower bandwidth, can increase the number of VR users supported. With two links, we get to serve up to six users, up to nine users with four links, and up to thirteen with eight links. Note that if we focus on the cases where all configurations meet the requirements, for up to three users, two links of 80 MHz result in lower delays overall. Then from four to eight users, four links of 40 MHz is the best option. Beyond nine users, it is better to use eight links.

Fig. 9b shows the UL packet delay for all three configurations, in which we can observe that increasing the number of links improves the delay for all cases. These results show there is a clear trend between the bandwidth used per channel and the links-per-user ratio. When we have few users but many links, as the downlink arrives in batches, most links end up being idle while a small subset is used for transmitting all the data. In these cases, a higher bandwidth allows for higher data capacity and lower transmission times. Once the number of users increases, we have a higher chance of transmitting simultaneously on all the links, and having fewer links results in increased waiting times in the queue. If we have four users and two links, we can only support up to two users simultaneously. Once there are packets from more users than links, some packets must wait for the ongoing transmission to finish before being transmitted, thus leading to increased delay.

The UL behaves differently from the DL for two reasons: the first is the lower traffic load required per STA, and the second is the timing between packets. As UL packets do not arrive in batches, the buffer does not fill as quickly as the DL, resulting in minimal aggregation. Transmissions are always short, thus not benefiting from higher capacity, and having more links allows all packets to be sent as soon as possible.

Takeaway: There is a relationship in the DL between channel access opportunities and transmission time (more independent links vs. more bandwidth per link). This is also affected by the number of VR users in the network.

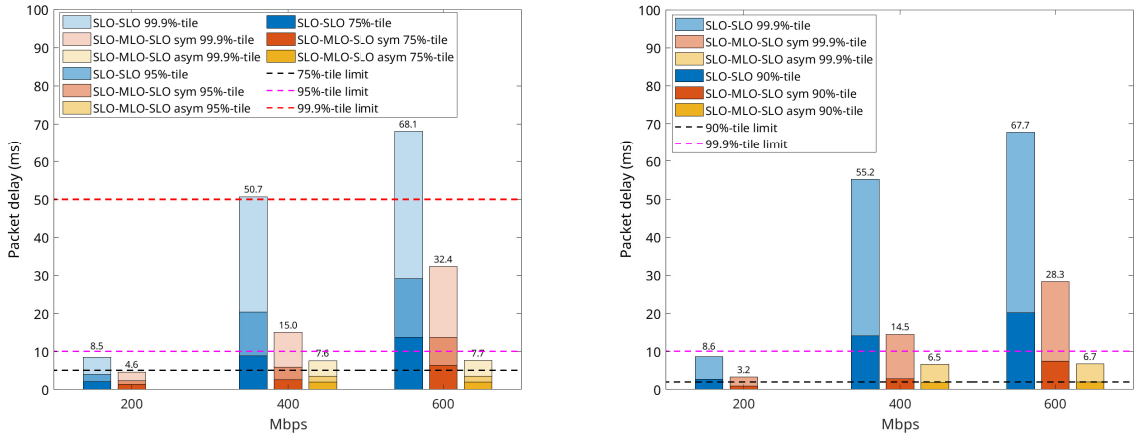

(a) Downlink

(b) Uplink

FIGURE 10: VR packet delay in the main BSS vs. total BG traffic load, with symmetrical and asymmetrical BG load. Note that asymmetric setup omitted in the 200 Mbps case, as it is equal to the symmetric setup.

Under the assumption of using the same total bandwidth, the number of configured links should be large enough to guarantee that the delay-sensitive UL transmissions are not blocked by channel access contentions and, at the same time, maintain a sufficient bandwidth to support the high DL throughput demand of VR applications.

VI. Results for OBSSs scenarios (VR Coexistence)

In this section, we study the impact of OBSS contention on VR streams. VR traffic is very sensitive to external contention, as UL packets need to be delivered in a timely manner to inform the VR server of the movement and inputs of the user, thus if the UL is interrupted, the user experience is impacted. We consider the interplay between BG and VR streams, as well as mutual interference between VR streams.

A. MLO contending with SLO BSSs

We start by studying the coexistence of VR streams with OBSS contention. We consider a main BSS with two contending BG streams in SLO BSSs. Both SLO and MLO are considered for the main BSS, with their respective configurations being adjusted to ensure a fair comparison (i.e., as SLO has only one link, they only perceive the contention from one BSS instead of two). For the MLO scenario, we set an MLO AP delivering VR content to a single user while contending with two SLO APs, each sharing one of the MLO channels. We consider two types of OBSS activity: (i) symmetric, in which both BG APs carry the same traffic load, or (ii) asymmetric, with different loads in each of the BG APs. The position of BG STAs is fixed to ensure consistent contention, and the VR STA is placed randomly in a 10 meter radius around the MLO AP.

Fig. 10a shows the DL packet delay for the VR stream vs. the BG traffic load (200, 400 or 600 Mbps). For symmetric cases (in red), the traffic is split evenly, while for the asymmetric cases (in yellow), one of the BG APs has a fixed load of 100 Mbps while the other transmits the rest (i.e. 100, 300 and 500 Mbps). As in the previous plots, the different delay requirements are shown using dashed lines. Starting with SLO (in blue), delays are high in all cases, only maintaining the delay requirements for 200 Mbps BG traffic. MLO achieves better results for both symmetric and asymmetric BG traffic, but performance is better for asymmetric scenarios, as MLO can use the link with lighter load of 100 Mbps to opportunistically transmit its traffic. All scenarios successfully deliver the full VR traffic. However, BG traffic is not always fully transmitted. In the SLO baseline, the BG AP fails to fully transmit its load in 5% of the 400 Mbps scenarios and becomes fully saturated for 600 Mbps. In the MLO setup with two SLO APs, the symmetric 600 Mbps scenario results in saturation of the BG AP, while in the asymmetric case, the overloaded AP reaches saturation at both 400 and 600 Mbps.

Fig. 10b shows the UL packet delay, in which SLO fails to deliver the minimum requirement for the 90th percentile even at the lowest BG traffic load. MLO works for 200 Mbps, and MLO meets the delay requirements for both 400 Mbps and 600 Mbps but only when the BG traffic is asymmetric. For a symmetric BG load, MLO cannot meet the delay requirements for a total BG traffic of 200 Mbps.

In this scenario, we also consider the buffer occupancy (ρ), which we define as the time average fraction of packets in the buffer normalized by the buffer size. It serves us as an indicator of how aggressively a device contends to access the channel (i.e., high ρ values mean higher contention). Table 3 shows the ρ values for the scenario with a main BSS using

SLO, showing that as BG traffic increases, so does ρ on both DL and UL. Table 4 shows the ρ values for an MLO main BSS and symmetric BG traffic. We observe that both BG APs have a high ρ when symmetric load is considered, even with 200 Mbps BG load. This means that they are both contending for the channel often, blocking the channel for the VR streams in the main AP, thus increasing its own ρ. In contrast, for the asymmetric case in Table 5, there is one AP with a much lower ρ, allowing the main VR AP to transmit through that channel more often, leading to an up to 45% decrease in ρ for both DL and UL VR streams.

TABLE 3: Buffer occupancy (ρ) values for the coexistence scenario of SLO VR main BSS and SLO BG OBSS. BG traffic is only DL.

Link	200 Mpbs		400 Mbps		600 Mbps	
	DL	UL	DL	UL	DL	UL
VR	0.4321	0.3509	0.7465	0.6120	0.8856	0.7334
BG	0.7181	-	0.7935	-	0.9369	-

TABLE 4: Buffer occupancy (ρ) values for the coexistence scenario between main MLO VR BSS and **symmetric** BG SLO OBSS.

Link	200 Mpbs		400 Mbps		600 Mbps	
	DL	UL	DL	UL	DL	UL
VR	0.2803	0.1761	0.4441	0.3011	0.7154	0.4795
BG 1	0.7135	-	0.7360	-	0.9080	-
BG 2	0.7082	-	0.7429	-	0.9064	-

TABLE 5: Buffer occupancy (ρ) values for the coexistence scenario between main MLO VR BSS and **asymmetric** BG SLO OBSS.

Link	400 Mbps		600 Mbps	
	DL	UL	DL	UL
VR	0.3785	0.2544	0.3891	0.2631
BG 1	0.7163	-	0.7118	-
BG 2	0.8770	-	0.9399	-

Takeaway: SLO struggles to maintain the necessary low delay for VR gaming when OBSS activity increases. MLO can use its multiple degrees of freedom to maintain a low delay even when OBSS interference is high, especially for cases in which asymmetric interference is perceived on the different channels. In these situations, MLO opportunistically selects the least loaded link.

B. MLO against MLO BSSs

In this section, we consider a fully overlapping scenario of 2 MLO APs with VR users, and study the interactions between VR streams. We consider two distinct scenarios, one in which the APs share all links, meaning that they use the same two channels for their two links (see Fig. 5a), and a second scenario in which the APs share one channel in one link, but have a dedicated separate channel on their second link (see Fig. 5b). As before, the VR STAs are placed randomly in a 10 meter radius arount their AP.

Fig. 11a shows the DL packet delay as the number of STAs increases in each AP, as well as the fraction of unsuccessful deployments. Delay increases quickly when the two BSSs share both links, with three STAs being the breaking point for meeting the 75th delay requirement. This matches the results in Fig. 8a (Section V C), since six users (three per AP) can be served with an MLO AP. Fig. 11b shows the uplink delay, and following the previously observed trends, the UL is more restrictive, limiting the shared channel case to two users per AP. Using a dedicated channel decreases the delay considerably, allowing the network to serve up to 4 STAs per AP. We can also observe that sharing two channels results in packet losses once we reach 5 users, while all scenarios are successful when using three channels. These results are in accordance with the 'MLO Performance Anomaly' [10], which appears when a single MLO BSS wins access and transmits in all its links at the same time.

The average buffer occupancy for each case is shown in Table 6. It can be observed that the buffer is significantly less occupied when using a dedicated channel for their second link, both for the DL and UL. For two channels, we can observe that the DL buffer is always full for four and five STAs, which has a severe impact on the UL ρ, making it increase quickly. In contrast, for three channels the ρ is high in the DL, but the UL can use a different link and keep its ρ low.

TABLE 6: Buffer occupancy (ρ) values for the Multiple MLO BSS VR scenarios.

STAs per AP	2 channels of 80 MHz		3 channels of 80 MHz	
	DL	UL	DL	UL
1	0.2351	0.1424	0.2219	0.1266
2	0.4975	0.2151	0.4277	0.1552
3	0.8523	0.3823	0.6367	0.2074
4	0.9997	0.6240	0.8511	0.2929
5	1.0000	0.8056	0.9978	0.4526

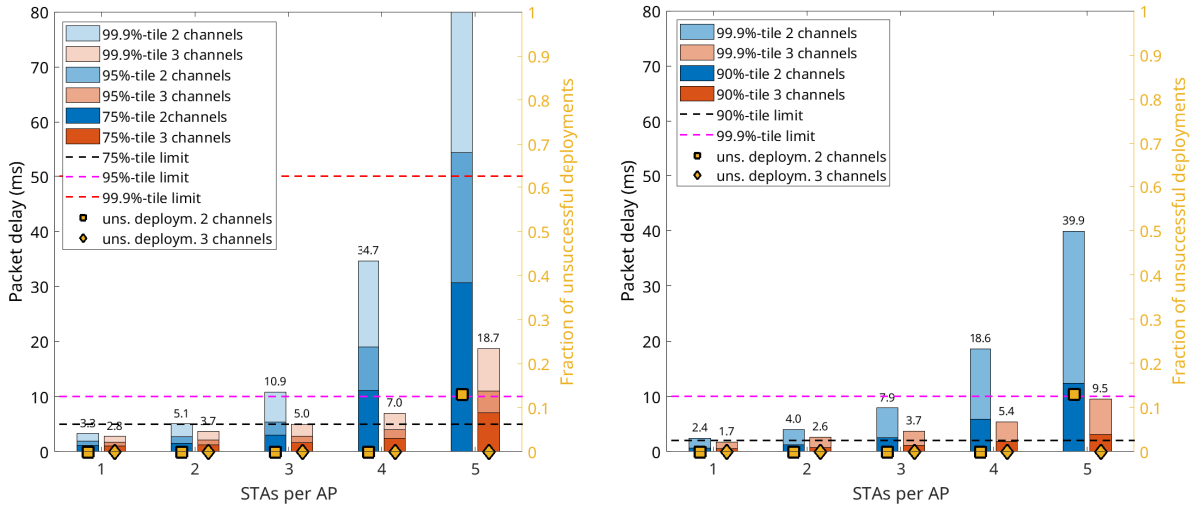

(a) Downlink (b) Uplink

FIGURE 11: Packet delay vs. number of associated STAs per AP while using two and three total channels.

Takeaway: Using the same channels for multiple overlapping MLO APs results in the APs blocking each other from transmitting if they have similar traffic loads. To improve MLO performance, it is necessary to adjust channel assignment strategies to avoid using the exact set of channels as another OBSS. Using at least one link with a separated channel results in better performance, allowing the APs to opportunistically transmit through the lighter loaded one.

C. Bespoke MLO for VR applications with OBSS interference

We repeat the previous experiment, but this time, we explore splitting the channel bandwidth into multiple smaller channels for each considered link, as described in Section V D. This approach serves as a potential alternative to our earlier proposal of using a third link when bandwidth constraints prevent the allocation of an additional 80 MHz channel.

Fig. 12 shows the DL and UL packet delay as the number of VR users increases in each AP. We use two links of 80 MHz, four links of 40 MHz, and eight links of 20 MHz. For the DL in Fig. 12a we observe the same trend that we saw in Fig. 9a (Section V D), where a few links is the best for a small number of users, but more links are needed when the number of users is big: two links of 80 MHz result in lower delays for one or two users, while four links leads to better performance for three users, and eight links is better for four users or more. In the UL in Fig. 12a, much like in Fig. 9b, it is always better to have more links. We also note that two links allow us to have two VR users, four links allow four users, and eight links only five users. Thus, in this case, having more links of 20 MHz does not provide a significant advantage. We also observe that fraction of unsuccessful deployments increases rapidly after 5 users, with 7 users being unsustainable even with the smallest links.

The buffer occupancy (ratio between packets in the buffer and max. capacity) for this scenario is presented in Table 7, in which we can observe that using 40 MHz and 80 MHz channels leads to lower ρ values for up to three STAs, but it increases rapidly for more STAs, leading to 20 MHz having lower ρ. The UL works as expected, with lower ρ for the lower channel widths.

Takeaway: As randomness in user placement and OBSS activity increases, maintaining the low delay required for good VR performance becomes more challenging. In such cases, increasing the number of links—even with reduced bandwidth per link—helps improve network capacity. However, with eight links, we observe diminishing returns, enabling only one additional user compared to the four-link scenario. This is because, in cases where low MCSs are used, each narrow link can transmit only a few packets per transmission, requiring significantly more transmissions to deliver the same load.

VII. Conclusions

In this paper, we modeled real VR traffic traces to test Wi-Fi MLO capabilities to support the stringent requirements of VR traffic in terms of MCS-bandwidth pairs and number of links. We showed that MLO can support VR applications with lower bandwidth and lower MCS than SLO, providing more robustness over a wider range of propagation conditions. We also showed that MLO offers lower delays due to increasing the number of independent channel access instances and that using an equivalent number of links, MLO allows an extra 50% of users per network over SLO. In order

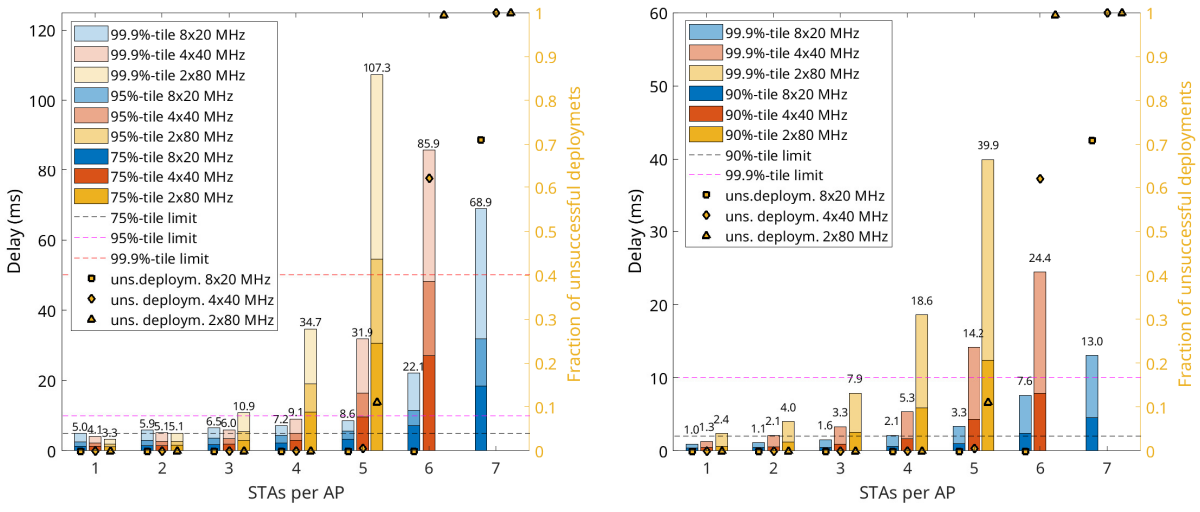

(a) Downlink

(b) Uplink

FIGURE 12: Packet delay for different configurations of links and bandwidth with multiple MLO BSSs and random STA placement.

TABLE 7: ρ values for each scenario.

STAs per AP	8 links 20 MHz		4 links 40 MHz		2 links 80 MHz	
	DL	UL	DL	UL	DL	UL
1	0.3024	0.1070	0.2666	0.1154	0.2351	0.1424
2	0.5467	0.1136	0.5014	0.1368	0.4975	0.2132
3	0.7462	0.1260	0.7239	0.1818	0.8523	0.3852
4	0.8861	0.1459	0.9166	0.2763	0.9997	0.6236
5	0.9830	0.1981	0.9999	0.5373	1.0000	0.8059
6	0.9998	0.3644	1.0000	0.7432	1.0000	0.8747
7	1.0000	0.6027	1.0000	0.7645	1.0000	0.8780

to accommodate a higher number of VR users, a proper configuration of links and channel bandwidth is required. Spreading the same bandwidth over more links can allow for more users to contend in the network without exceeding delay requirements, but for a lower user count, having an excess of links may not result in any gains.

Further, we studied OBSS scenarios with SLO background traffic, showing how MLO can outperform SLO by opportunistically using its links. We also demonstrate how proper channel allocation policies can improve performance in overlapping MLO scenarios, mitigating the 'MLO Performance Anomaly' which negatively affects worst case latency.

Acknowledgement

M. Carrascosa-Zamacois and B. Bellalta were supported in part by grants MAX-R HEu-CL4-MAX-R-101070072, Wi-XR PID2021-123995NB-I00, and by MCIN/AEI under the Maria de Maeztu Units of Excellence Programme (CEX2021-001195-M). G. Geraci was in part supported by the Spanish Research Agency through grants PID2021- 123999OB-I00, CEX2021-001195-M, and CNS2023-145384. L. Galati Giordano was in part supported by UNITY- 6G project, funded from European Union's Horizon Europe Smart Networks and Services Joint Undertaking (SNS JU) research and innovation programme under the Grant Agreement No 101192650.

REFERENCES

[1] E. Oughton, G. Geraci, M. Polese, M. Ghosh, W. Webb, and D. Bubley, "The future of wireless broadband in the peak smartphone era: 6G, Wi-Fi 7, and Wi-Fi 8," *techrxiv.173221414.41000458*, 2024.

[2] L. Galati-Giordano, G. Geraci, M. Carrascosa, and B. Bellalta, "What will Wi-Fi 8 be? A primer on IEEE 802.11 bn ultra high reliability," *IEEE Communications Magazine*, vol. 62, no. 8, pp. 126–132, 2024.

[3] G. Li, "Wi-Fi 7 market momentum: Fueling the evolution of XR," https://www.wi-fi.org/beacon/guoqing-li/wi-fi-7-market-momentum-fueling-the-evolution-of-xr, 2024.

[4] T. beacon, "The state of connectivity: Wi-Fi momentum in 2024," https://www.wi-fi.org/beacon/the-beacon/the-state-of-connectivity-wi-fi-momentum-in-2024, 2024.

[5] W.-F. Alliance, "Wi-Fi brings immersive XR experience to life," https://www.wi-fi.org/system/files/Wi-Fi%20brings%20immersive%20XR%20experiences%20to%20life.pdf, 2024.

[6] M. Carrascosa-Zamacois, L. Galati-Giordano, F. Wilhelmi, G. Fontanesi, A. Jonsson, G. Geraci, and B. Bellalta, "Performance Evaluation of MLO for XR Streaming: Can Wi-Fi 7 Meet the Expectations?" in *Proc. IEEE CAMAD*, 2024, pp. 1–6.

[7] "IEEE Draft Standard for Information technology–Telecommunications and information exchange between systems Local and metropolitan area networks–Specific requirements - Part

11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment: Enhancements for Extremely High Throughput (EHT)," *P802.11be/D7.0, Aug 2024*, pp. 1–999, 2024.

[8] A. Garcia-Rodriguez, D. López-Pérez, L. Galati-Giordano, and G. Geraci, "IEEE 802.11be: Wi-Fi 7 Strikes Back," *IEEE Communications Magazine*, vol. 59, no. 4, pp. 102–108, 2021.

[9] E. Khorov, I. Levitsky, and I. F. Akyildiz, "Current status and directions of IEEE 802.11 be, the future Wi-Fi 7," *IEEE access*, vol. 8, pp. 88 664–88 688, 2020.

[10] M. Carrascosa-Zamacois, G. Geraci, L. Galati-Giordano, A. Jonsson, and B. Bellalta, "Understanding Multi-link Operation in Wi-Fi 7: Performance, Anomalies, and Solutions," in *Proc. IEEE PIMRC*, 2023, pp. 1–6.

[11] C. Chen, X. Chen, D. Das, D. Akhmetov, and C. Cordeiro, "Overview and performance evaluation of Wi-Fi 7," *IEEE Communications Standards Magazine*, vol. 6, no. 2, pp. 12–18, 2022.

[12] G. Naik, D. Ogbe, and J.-M. J. Park, "Can Wi-Fi 7 support real-time applications? On the impact of multi link aggregation on latency," in *ICC 2021-IEEE International Conference on Communications*. IEEE, 2021, pp. 1–6.

[13] M. Carrascosa-Zamacois, G. Geraci, E. Knightly, and B. Bellalta, "Wi-Fi multi-link operation: An experimental study of latency and throughput," *IEEE/ACM Transactions on Networking*, 2023.

[14] A. A. Abdalhafid, S. K. Subramaniam, Z. A. Zukarnain, and F. H. Ayob, "Multi-Link Operation in IEEE802. 11be Extremely High Throughput: A Survey," *IEEE Access*, 2024.

[15] M. Jansen, J. Donkervliet, A. Trivedi, and A. Iosup, "Can My WiFi Handle the Metaverse? A Performance Evaluation Of Meta's Flagship Virtual Reality Hardware," in *Companion of the 2023 ACM/SPEC International Conference on Performance Engineering*, 2023, pp. 297–303.

[16] G. Gonçalves, P. Monteiro, M. Melo, J. Vasconcelos-Raposo, and M. Bessa, "A comparative study between wired and wireless virtual reality setups," *IEEE access*, vol. 8, pp. 29 249–29 258, 2020.

[17] C. Michaelides, M. Casasnovas, D. Marchitelli, and B. Bellalta, "Is Wi-Fi 6 Ready for Virtual Reality Mayhem? Aa Case Study Using One AP and Three HMDs," *Authorea Preprints*, 2023.

[18] C. Michaelides, M. Casasnovas, D. Nunez, and B. Bellalta, "Lessons Learned From a Large-Scale Virtual Reality Experience Over Wi-Fi," 2025.

[19] A. Hazarika and M. Rahmati, "Towards an evolved immersive experience: Exploring 5G-and beyond-enabled ultra-low-latency communications for augmented and virtual reality," *Sensors*, vol. 23, no. 7, p. 3682, 2023.

[20] I. F. Akyildiz and H. Guo, "Wireless communication research challenges for extended reality (XR)," *ITU Journal on Future and Evolving Technologies*, vol. 3, no. 1, pp. 1–15, 2022.

[21] A. Mallik, J. Xie, and Z. Han, "A Performance Analysis Modeling Framework for Extended Reality Applications in Edge-Assisted Wireless Networks," *arXiv preprint arXiv:2405.07033*, 2024.

[22] J. Ahn, Y. Y. Kim, and R. Y. Kim, "Virtual reality-wireless local area network: Wireless connection-oriented virtual reality architecture for next-generation virtual reality devices," *Applied Sciences*, vol. 8, no. 1, p. 43, 2018.

[23] H. Zhang, A. Elmokashfi, and P. Mohapatra, "Wifi and multiple interfaces: Adequate for virtual reality?" in *2018 IEEE 24th International Conference on Parallel and Distributed Systems (ICPADS)*. IEEE, 2018, pp. 220–227.

[24] M. Alsakati, C. Pettersson, S. Max, V. N. Moothedath, and J. Gross, "Performance of 802.11 be Wi-Fi 7 with multi-link operation on AR applications," in *2023 IEEE Wireless Communications and Networking Conference (WCNC)*. IEEE, 2023, pp. 1–6.

[25] J. Perez-Ramirez, O. Seijo, and I. Val, "Time-critical IoT applications enabled by Wi-Fi 6 and beyond," *IEEE Internet of Things Magazine*, vol. 5, no. 3, pp. 44–49, 2022.

[26] C. Michaelides, M. Casanovas, D. Marchitelli, and B. Bellalta, "VR Gaming Dataset - Multi-users tests," https://doi.org/10.5281/zenodo.8169785, 2023.

[27] Wi-Fi Alliance, "Wi-Fi delivers immersive VR gaming," https://www.wi-fi.org/system/files/VR_Gaming_Highlights_20231215_0.pdf, 2023, last accessed 17/06/2024.

[28] IEEE 802.11 TGax, "TGax Simulation Scenarios," https://mentor.ieee.org/802.11/dcn/14/11-14-0980-16-00ax-simulation-scenarios.docx, 2015, last accessed 17/06/2024.