An Experimental Study of Latency for IEEE 802.11be Multi-link Operation

Marc Carrascosa[⋆], Giovanni Geraci[⋆], Edward Knightly[♭], and Boris Bellalta[⋆]

[⋆]Dept. Information and Communication Technologies, Universitat Pompeu Fabra, Barcelona
[♭]Dept. Electrical and Computer Engineering and Computer Science, Rice University, Houston, TX

Abstract—**Will Multi-Link Operation (MLO) be able to improve the latency of Wi-Fi networks? MLO is one of the most disruptive MAC-layer techniques included in the IEEE 802.11be amendment. It allows a device to use multiple radios simultaneously and in a coordinated way, providing a new framework to improve the WLAN throughput and latency. In this paper, we investigate the potential latency benefits of MLO by using a large dataset containing 5 GHz spectrum occupancy measurements. Experimental results show that when the channels are symmetrically occupied, MLO can improve latency by one order of magnitude. In contrast, in asymmetrically occupied channels, MLO can sometimes be detrimental and increase latency. This is a result of packets being assigned to an interface before carrying out the backoff, which is more likely to be interrupted on the busier link. We overcome this issue by allowing multiple backoffs to run in parallel, assigning the packet to the particular interface where the backoff expires first, which also achieves lower latency overall.**

I. INTRODUCTION

The importance of wireless connectivity in a globalized society is unquestionable, and forced lockdowns reminded us how dependable Wi-Fi is. We resorted to Wi-Fi to be in touch with our loved ones, to make online purchases, and to get work done and keep the economy afloat. In a post-pandemic world, Wi-Fi technologies will be vital for accessing fair and remote-friendly education, medical care, and business opportunities in the unlicensed spectrum. There will be nearly 628 million public Wi-Fi hotspots by 2023 [1], one out of ten equipped with Wi-Fi 6 based on the IEEE 802.11ax amendment [2].

As the popularity of Wi-Fi grows, so does the demand for augmented data rates, higher reliability, and lower latency, driving the development of a new Wi-Fi 7 generation based on the IEEE 802.11be Extremely High Throughput (EHT) specification [3]–[7]. Despite its name, Wi-Fi 7 will be chasing much more than peak throughput. Indeed, the 802.11be Task Group acknowledges the need for lower delays to enable delay-sensitive networking use cases, including augmented and virtual reality, cloud computing, and cross-factory floor communications in next-generation enterprises [8]–[12].

In a quest for lower delays, one of the most disruptive features being proposed for 802.11be is Multi-Link Operation (MLO) [13]–[16]. In MLO, devices can make simultaneous use of different channels or bands, potentially allowing delay-sensitive traffic to be transmitted through multiple links to ensure its timely reception. With its standardization process being consolidated, and prompted by the increasing interest from the research community [17]–[20], a fundamental question arises as to whether and to what extent MLO can reduce Wi-Fi latency in real-world scenarios.

In this paper, capitalizing on over-the-air measurements of spectrum occupancy for the entire 5 GHz band recently collected [21], [22] and freely available in open source[1], we experimentally investigate the latency[2] performance of 802.11be MLO. Atop these traces, which include scenarios with high Access Point (AP) density and crowded environments and span multiple hours, we develop an emulation tool that fuses a Wi-Fi MLO state machine with the high-resolution spectrum measurements. Besides legacy Wi-Fi Single-Link Operation (SLO), we study the two MLO channel access modes currently under consideration by the IEEE 802.11be Task Group [3], [16]: (i) MLO-STR, where two radio interfaces are operated independently, and (ii) MLO-NSTR, where one interface acts as primary and the other as secondary. Our main contributions can be summarized as follows:

- We show that when using two links with statistically symmetrical occupancy, MLO reduces 95th percentile latency by up to an order of magnitude with respect to SLO by availing of a second radio interface.
- In contrast, we surprisingly discover that when using two links with asymmetrical occupancy, MLO-STR can sometimes worsen the latency performance with respect to SLO. In the worst case, we observe an increase of up to 112% in terms of 95th percentile latency.
- To overcome the aforementioned issue, we consider a minor variation of STR, denoted MLO-STR+, that allows to run in parallel as many backoff instances as interfaces. Then, MLO-STR+ simply allocates the first packet waiting for transmission to the interface whose backoff expires first. This way, STR+ guarantees same delay as

M. Carrascosa and B. Bellalta were supported by WINDMAL PGC2018-099959-B-I00 (MCIU/AEI/FEDER,UE) and Cisco.

G. Geraci was supported by MINECO's Project RTI2018-101040 and by a "Ramón y Cajal" Fellowship from the Spanish State Research Agency.

E. Knightly's research was supported by Cisco, Intel, and NSF Grants CNS-1955075 and CNS-1923782.

---

[1] WACA dataset: https://github.com/sergiobarra/WACA_WiFiAnalyzer.

[2] The terms latency and delay are used interchangeably throughout the paper.

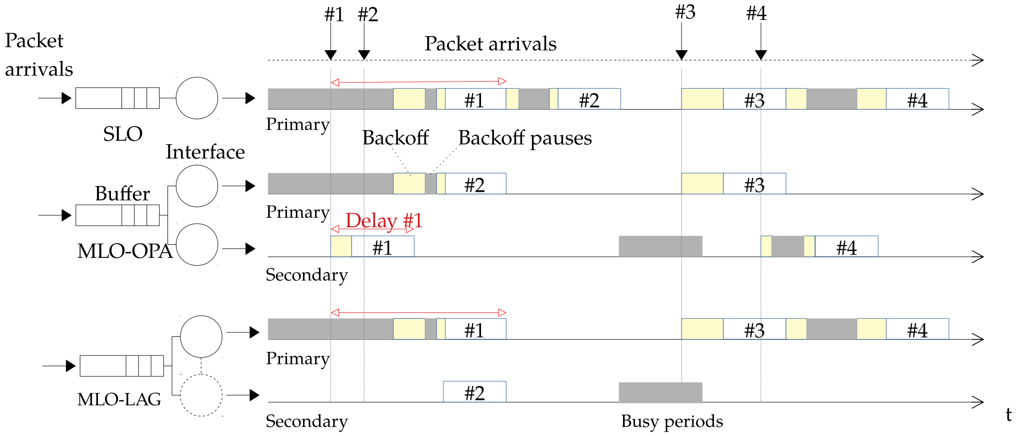

Fig. 1: Illustration of SLO, MLO-STR, and MLO-NSTR operations. Grey, yellow, and white bars respectively indicate occupied channels, random backoffs, and packet transmissions. Packet transmissions include both the data part and the corresponding ACK, as well as DIFS and SIFS inter-frame spaces.

or lower than SLO, with reductions of up to 70% in the best observed case.

II. MULTI-RADIO MULTI-LINK OPERATION

IEEE 802.11be considers two main channel access methods to support Multi-link Operation: Simultaneous Transmit and Receive (MLO-STR), and Non-simultaneous Transmit and Receive (MLO-NSTR) [3], [16]. We introduce them in the following, from the perspective of an AP equipped with two radio interfaces and thus able to operate on two different channels simultaneously:

- **MLO-STR:** The two radio interfaces operate independently and asynchronously, and a packet waiting for transmission is allocated to a radio interface as soon as the latter becomes available. If both radio interfaces are available, the packet is randomly allocated to either. Once an interface is allocated a packet, it starts channel contention by initializing a backoff instance.
- **MLO-NSTR:** One interface acts as primary, and the other as secondary. When there are packets waiting for transmission, the primary interface undergoes contention to access the channel through a backoff counter. Once the backoff counter reaches zero, packets are sent through both interfaces if the secondary one has been idle for at least a PIFS interval. Otherwise, only the primary interface is used to transmit.

Besides the MLO modes, IEEE 802.11be also considers the conventional Single-link Operation, where an AP is equipped with only one radio interface.

Figure 1 exemplifies SLO, MLO-STR, and MLO-NSTR operations. SLO follows default Wi-Fi operations, where packets are sequentially transmitted. In the case of MLO-STR, arriving packets are allocated to whichever interface becomes available first. This results in a significant delay reduction for packets #1, #2 and #4. In the case of MLO-NSTR, the secondary channel's dependence on the primary sometimes prevents efficiently using the two radio interfaces. As a result, and unlike MLO-STR, the delay for packets #1 and #4 cannot be reduced with respect to SLO.

III. EXPERIMENTAL SETUP

In this work, we consider a target WLAN Basic Service Set (BSS) consisting of one AP and one station (STA), both equipped with two Wi-Fi interfaces each operating in the 5 GHz band on channels 36 and 100, respectively denoted *primary* and *secondary*. We refer to this BSS as MLO-BSS. On these channels, the target MLO-BSS observes the environment activity, i.e., the transmissions generated by Orthogonal Basic Service Sets (OBSS). The MLO-BSS and OBSS under consideration are illustrated in Figure 2 in blue and red, respectively. For this setup, we consider the three modes of operation described in Section II, namely: (i) SLO, where only the primary channel interface is available; (ii) MLO-STR, where both interfaces are available and work independently; and (iii) MLO-NSTR, where both interfaces are available but usage of the secondary channel is conditioned on the primary also being unoccupied.

For the above scenario, we consider downlink traffic, i.e., from the AP to the STA. We assume packet arrivals to follow a Poisson process, and transmitted packets to have a constant size of $L = 12000$ bits. Table I summarizes the main parameters used in the Wi-Fi state machine.

A. WACA Dataset

In order to evaluate the latency of 802.11be MLO in a real-world setting, we employ the WACA dataset, containing over-the-air measurements of the 5 GHz band occupancy that we have recently collected and made publicly available. This dataset was obtained by conducting extensive measurement campaigns on different days and in multiple locations, including a sold-out football stadium (F. C. Barcelona's Camp Nou). In this paper, we employ the football stadium measurements since they range from completely idle to fully occupied

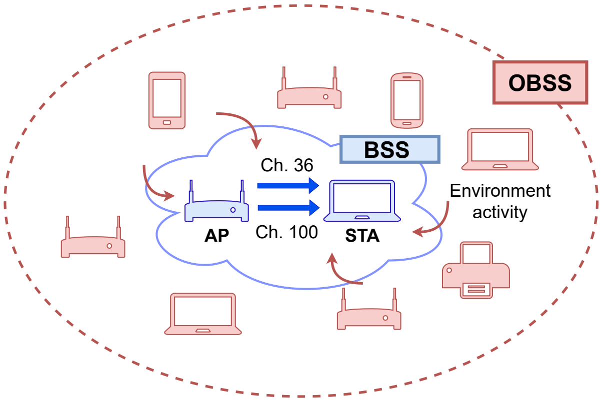

Fig. 2: Scenario considered. The WACA dataset is used to characterize the environment activity (red) observed by the target BSS (blue) on channels 36 and 100 in the 5 GHz band.

channels. In the dataset, spectrum samples consist of 1 s of consecutive, 10 μs receive signal strength indicator (RSSI) measurements. We refer the reader to [21], [22] for further details on the dataset. Compared to [21], [22], in this work we have implemented a new Wi-Fi state machine, capable of (i) fully characterizing the temporal dynamics of the system under finite traffic loads, i.e., non-full buffer conditions, and (ii) supporting multiple Wi-Fi interfaces and packet buffers.

In what follows, we employ the FCB-WACA dataset to investigate how different combinations of primary and secondary channel occupancies affect the MLO-BSS performance. In particular, we assume that the MLO-BSS perceives the same spectrum activity as the one captured in the WACA dataset, and it contends for channel access accordingly. As for the OBSS, we adopt the same hinder interaction model as in [22], assuming that the OBSS sense the MLO-BSS channel access whenever this takes place and therefore defer their transmissions.

B. Trace-based Simulations Methodology

In order to study the effect of channel occupancy on latency, we partition the available traces in our dataset for both primary and secondary channels into different average channel occupancy regimes: $\{10\%, 20\%, \ldots, 90\%\}$. Then, we run each simulation as follows:

1) We select the occupancy regime of interest for the primary and secondary channels, e.g., 10% and 40%, respectively;
2) We combine uniformly at random one spectrum sample each for the primary and secondary channels;
3) For each spectrum sample pair and given a particular traffic load of interest, we compute the packet arrival times at the AP;
4) We execute the Wi-Fi state machine for SLO, MLO-STR, and MLO-NSTR access policies. The same packet arrival times are considered in all cases to allow a direct comparison.
5) We store the individual delay experienced by each packet over all spectrum samples.

For a fair comparison between SLO and MLO, we guarantee that all results are obtained in non-saturation conditions and thus we discard any simulations where less than 95% of all the transmitted packets are received.

IV. DELAY PERFORMANCE

This section investigates the delay performance of both SLO and MLO modes for different combinations of channel occupancies and traffic loads.

To evaluate the gains of MLO in terms of delay, we consider the same traffic load for both SLO and MLO modes. In addition, the primary channel considered in the MLO mode is the same channel used in SLO. Therefore, we could expect that adding another channel in MLO mode, regardless of its occupancy, should yield lower delays.

A. Symmetrically Occupied Channels

Here we study the case of *symmetric* channel occupancies in which both MLO interfaces have channels with similar occupancy levels. In particular, we study the delay performance with pairs of channels in the ranges of 10%, 40%, and 70% occupancy. In those cases, the average full-buffer throughput under SLO is 37, 22, and 6.8 Mbps, respectively. For these three scenarios (symmetric low, medium, and high occupancy), we feed the Wi-Fi state machine with Poisson traffic and vary the intensity as a fraction of this SLO average full-buffer throughput, namely 0.2, 0.4, 0.6, 0.8.

Figure 3 shows the average and 95th percentile delay for all channel access modes and the different channel occupancies. First, we observe that when both channels have 10% occupancy (Figure 3a), the three schemes have strikingly different scaling with increasing traffic load as MLO delay does not increase at the same rate as SLO delay. For example, at 20% traffic load, STR and NSTR offer a modest decrease in average delay compared to Single Link Operation of 17% and 9% respectively. In contrast, when the traffic load is 80%, STR and NSTR reduce the average delay by 69% and 62%. This scaling is even more pronounced analyzing the 95th percentile of delay, in which MLO achieves up to a 78% delay reduction.

Name	Variable	Value
Legacy preamble	$T_{\text{PHY-legacy}}$	20 μs
HE single-user preamble	$T_{\text{PHY-HE-SU}}$	52 μs
OFDM symbol duration	σ	16 μs
OFDM legacy symbol dur.	σ_{Legacy}	4 μs
Short InterFrame Space	SIFS	16 μs
DCF InterFrame Space	DIFS	30 μs
Slot time	T_0	10 μs
Service field	L_{SF}	32 bits
MAC header	L_{MH}	272 bits
Tail bits	L_{TB}	6 bits
ACK bits	L_{ACK}	112 bits
Frame size	L	12000 bits

TABLE I: Notation and Wi-Fi state machine parameters.

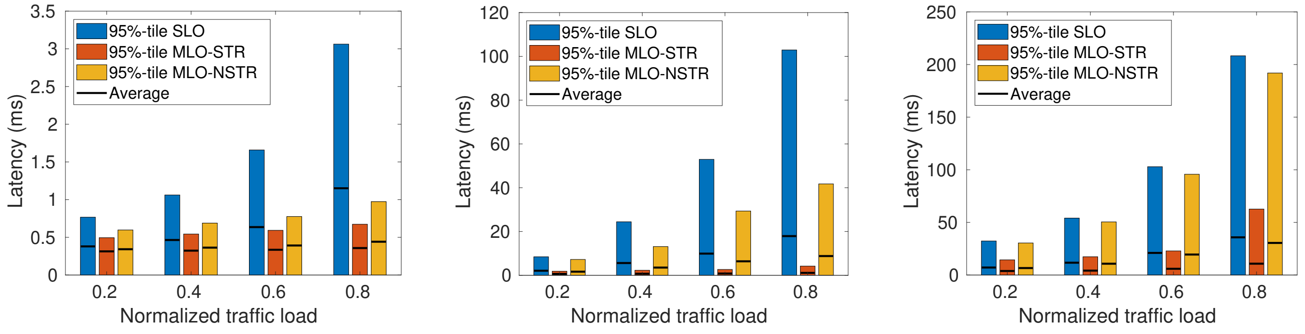

(a) 10% occupancy on both channels

(b) 40% occupancy on both channels

(c) 70% occupancy on both channels

Fig. 3: Latency for symmetrically occupied channels vs. variable normalized traffic load.

Thus, for both average and 95th percentile delay, the benefits of MLO are increasingly pronounced under higher traffic load as in this case, there are often multiple packets in the buffer such that both interfaces can be used. Moreover, with a relatively low channel occupancy of 10%, both channels are often available.

Next, we consider the case that both channels have symmetrical medium (40%) occupancy (Figure 3b). Here, while SLO's *average* delay increases only modestly with traffic (i.e., from 2 to 18 ms), the 95th percentile delay increases much more rapidly, exceeding 100 ms. In contrast, STR can yield a staggering order of magnitude reduction in 95th percentile delay compared to SLO. The reason is that STR avails usage of two channels and can access either or both of them. STR, therefore, realizes delay benefits compared to SLO unless both channels are occupied.

Unfortunately, unlike STR, the benefits of NSTR over SLO are limited and mostly confined to how the average delay scales with traffic load. Indeed, NSTR can only gain access to the secondary channel if the primary channel is also idle, implying that the average delay is guaranteed to be lower than the average delay under SLO. However, the 95th percentile delays are triggered by long periods of occupancy of the primary channel, thus making any availability of the secondary channel during this time irrelevant. As a result, the 95th percentile delay under NSTR rapidly grows as the normalized traffic load increases.

Lastly, when both channels have high (70%) occupancy (Figure 3c), STR again has the most favorable 95th percentile delay scaling with traffic load, providing substantial reductions as compared to both SLO and NSTR. Nonetheless, at such high channel occupancies, even STR has difficulty finding transmission opportunities on either channel, so both mean and 95th percentile delays are increasing. Additionally, NSTR provides negligible benefits in both average and 95th percentile delay compared to SLO.

Findings: When both channel occupancies are symmetrically medium to high load, NSTR fails to provide significant 95th percentile delay benefits compared to SLO. The key reason is that NSTR is only able to realize a benefit compared to SLO if both channels are simultaneously unoccupied, an increasingly unlikely occurrence in this scenario. Fortunately, STR yields significant 95th percentile latency benefits (compared to both SLO and NSTR) even in the challenging regime of increasing occupancies and traffic. This is because STR can utilize either available channel, and reduce the packets waiting time even if it cannot simultaneously utilize both available channels.

B. *Asymmetrically Occupied Channels*

Here, we employ the same normalized traffic loads from the previous section, but change the channel occupancy of our interfaces so that they lie in different ranges. Between the two channels, we always assume the primary to be the less occupied one. Note that the opposite case favors both MLO modes in this comparison: SLO would always incur a high delay, and both MLO modes would take advantage of a more idle secondary channel.

Figure 4a depicts the case of a low (10%) primary and medium (40%) secondary channel occupancy. As expected, NSTR offers deterministically lower delays than SLO, with the highest benefits occurring under higher traffic loads. However, STR surprisingly incurs a higher average and 95th percentile delay than SLO for the lowest traffic load of 0.2. Indeed, STR starts contention by initializing the backoff counter as soon as a channel is detected to be idle. Unfortunately, such channel may be occupied before the backoff timer expires, thus pausing the backoff counter. If the backoff is paused too often (or for long intervals), the packet could incur even higher delays than it would have if the other channel—initially busy—had been selected.

In Figure 4b, this effect is exacerbated due to the even higher occupancy of the secondary channel, as selecting an idle secondary channel incurs the risk of the latter being occupied before the backoff counter expires. When this occurs, the 95th percentile delay can be twice as high as that with SLO, albeit still confined to below 10 ms. However, STR average and 95th percentile delays grow at a lower rate than those of SLO as the traffic load increases. Indeed, STR can still take advantage of a secondary channel (even when highly occupied) to reduce congestion and curb the latency when it is caused

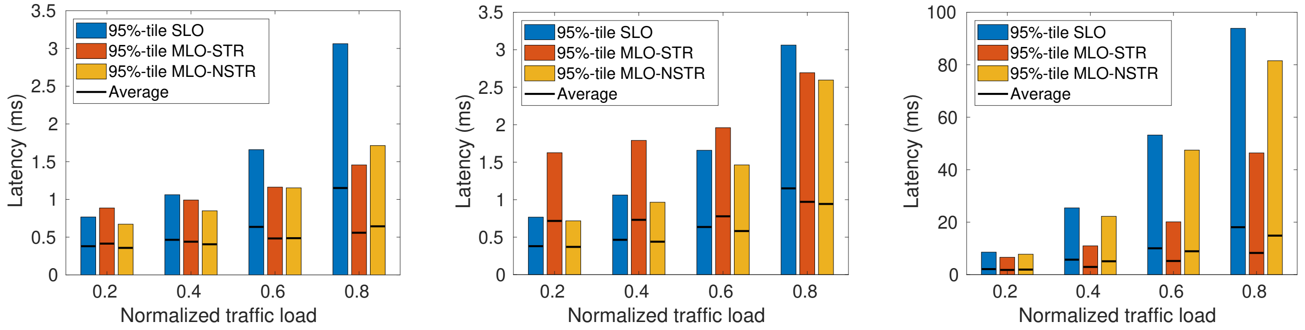

(a) Primary of 10% and secondary of 40% (b) Primary of 10% and secondary of 70% (c) Primary of 40% and secondary of 70%

Fig. 4: Latency for non-symmetrically occupied channels vs. variable normalized traffic load.

not only by the channel occupancy patterns but also by the amount of traffic.

Finally, Figure 4c considers the more symmetrical case of primary and secondary channel occupancy of 40% and 70%, respectively. Similar to the prior case of Figures 3b and 3c, STR scales well with the increasing traffic load, keeping the average delay below that of SLO and the decreasing the 95th percentile by up to a half. Compared to Figure 4b, the primary channel occupancy has grown, leading to a faster increase in the SLO delay vs. traffic load. However, STR is capable of leveraging both links and thus achieves lower delays.

Findings: Channel occupancy is a crucial factor to account for when selecting a secondary channel in MLO mode. For STR, specifically, using a secondary channel that is much busier than the primary can lead to even higher delays than using SLO. This is owed to packets being suboptimally assigned to an interface before carrying out the backoff, with the latter likely to be interrupted on the busier channel. This effect is exacerbated when the difference between channel occupancies increases.

C. MLO-STR with parallel backoffs

We have shown that MLO-STR can lead to even higher delays than using SLO in the case of channels with different occupancies. To better understand the design space of MLO channel access, we now define a minor variation—denoted *Opportunistic MLO-STR* (MLO-STR+)—which nonetheless allows to overcome the limitations of MLO-STR and to evaluate the ultimate capabilities of MLO.

- **MLO-STR+:** When both interfaces are idle, one backoff instance is started on each. Packet allocation is deferred until either backoff counter expires, and the first waiting packet is allocated to the interface whose counter expires first. This approach differs from MLO-STR, where a packet is assigned to a channel as soon as the latter is idle, without waiting for its backoff counter to expire.

The main advantage of MLO-STR+ lies in the fact that, if one channel becomes occupied during the backoff, a transmission opportunity may be found on the other, avoiding unnecessarily delaying the waiting packet. In practice, implementing MLO-STR+ only requires a minor firmware update on the current Wi-Fi state machine: the ability to control when an interface can initiate, pause, and complete the backoff countdown without actually being allocated a packet.

Figure 5 shows the average and 95th percentile delay for the same cases studied in Figure 4. We still take SLO as the baseline and compare MLO-STR and MLO-STR+ modes. In Figure 5a, STR+ consistently outperforms STR and SLO in both average and 95th percentile delay, since packets are transmitted either at the same time as in SLO, or faster via the secondary interface. In Figure 5b, when the secondary channel has a 70% occupancy, we encounter the worst scenario for STR. In this case, STR selects the secondary channel when it undergoes a short idle periods. However, since the latter are typically followed by longer intervals of occupancy, the backoff counter often remains frozen, leading to 95th percentile delays more than twice as high as those with SLO. This shortcoming is avoided altogether by the proposed STR+, assigning a packet to either interface only after ensuring that the corresponding backoff counter has expired. Finally, Figure 5c depicts the case of 40% and 70% occupancy on the primary and secondary channels, respectively. As the former has increased, the SLO delay grows rapidly. STR already outperforms SLO in average and 95th percentile delay, and STR+ slightly reduces these values further.

Findings: MLO-STR+ improves over MLO-STR by delaying the allocation of the packet at the head of the queue until one of the backoff counters expires, allowing to leverage up-to-date information on the channel state, and thus to ultimately make better decisions.

V. CONCLUSIONS

In this paper, we provided an experimental study of latency for IEEE 802.11be MLO. Using the WACA dataset, which contains real-world channel occupancy measurements in the 5 GHz spectrum, we cast light upon the latency performance of two MLO channel access modes, namely (i) MLO-STR, where two radio interfaces are operated independently, and (ii) MLO-NSTR, where one interface acts as primary and the other as secondary.

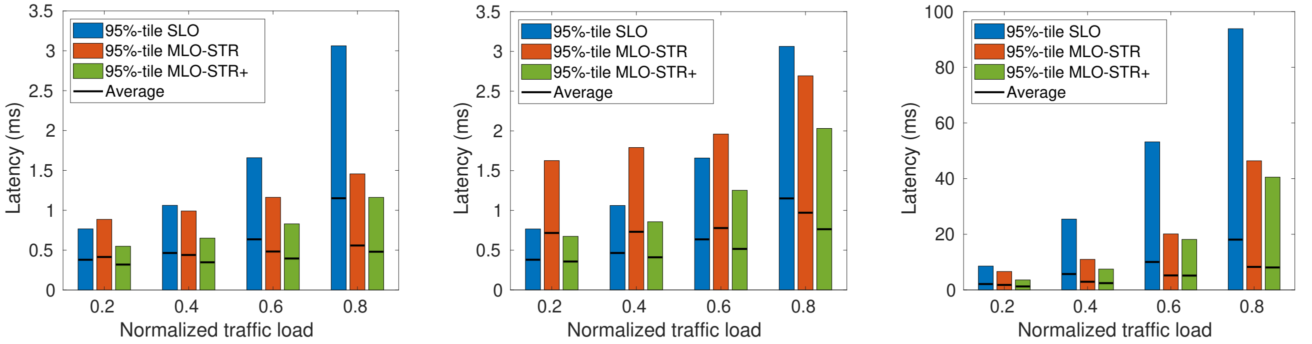

(a) Primary of 10% and secondary of 40% (b) Primary of 10% and secondary of 70% (c) Primary of 40% and secondary of 70%

Fig. 5: Latency for non-symmetrically occupied channels vs. variable normalized traffic load. MLN-STR vs MLN-STR+.

We showed that when both channels are on average equally occupied, both MLO modes can reduce the 95th percentile latency by nearly one order of magnitude as they avail of a second radio interface. In contrast, in asymmetrically occupied channels, we surprisingly found the use of MLO-STR to be detrimental and cause even higher latency values than SLO. We define MLO-STR+ to show that this issue can be overcome by delaying the packet assignment until the expiration of the backoff, which also achieves lower latency overall.

REFERENCES

[1] Cisco Annual Internet Report (2018–2023) White Paper, Mar. 2020.

[2] IEEE 802.11, "P802.11ax - 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 High Efficiency WLAN," 2020.

[3] "IEEE P802.11be/D1.0 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 8: Enhancements for extremely high throughput (EHT)," May 2021.

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

[5] 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.

[6] R. P. F. Hoefel, "IEEE 802.11 be: Throughput and Reliability Enhancements for Next Generation WI-FI Networks," in *2020 IEEE 31st Annual International Symposium on Personal, Indoor and Mobile Radio Communications*. IEEE, 2020, pp. 1–7.

[7] M. Yang and B. Li, "Survey and perspective on extremely high throughput (EHT) WLAN—IEEE 802.11 be," *Mobile Networks and Applications*, vol. 25, no. 5, pp. 1765–1780, 2020.

[8] D. Lopez-Perez, A. Garcia-Rodriguez, L. Galati-Giordano, M. Kasslin, and K. Doppler, "IEEE 802.11be Extremely High Throughput: The Next Generation of Wi-Fi Technology Beyond 802.11ax," *IEEE Communications Magazine*, vol. 57, no. 9, pp. 113–119, 2019.

[9] B. Bellalta, "On the Low-Latency Region of Best-Effort Links for Delay-Sensitive Streaming Traffic," *IEEE Communications Letters*, vol. 25, no. 3, pp. 970–974, 2020.

[10] T. Adame, M. Carrascosa-Zamacois, and B. Bellalta, "Time-sensitive networking in IEEE 802.11 be: On the way to low-latency WiFi 7," *Sensors*, vol. 21, no. 15, p. 4954, 2021.

[11] M. Carrascosa and B. Bellalta, "Cloud-gaming: Analysis of Google stadia traffic," *arXiv preprint arXiv:2009.09786*, 2020.

[12] C. Deng, X. Fang, X. Han, X. Wang, L. Yan, R. He, Y. Long, and Y. Guo, "IEEE 802.11be Wi-Fi 7: New Challenges and Opportunities," *IEEE Communications Surveys Tutorials*, vol. 22, no. 4, pp. 2136–2166, 2020.

[13] I. Levitsky, Y. Okatev, and E. Khorov, "Study on Simultaneous Transmission and Reception on Multiple Links in IEEE 802.11 be networks," in *2020 International Conference Engineering and Telecommunication (En&T)*. IEEE, 2020, pp. 1–4.

[14] M. Yang, B. Li, Z. Yan, and Y. Yan, "AP Coordination and Full-duplex enabled Multi-band Operation for the Next Generation WLAN: IEEE 802.11 be (EHT)," in *2019 11th International Conference on Wireless Communications and Signal Processing (WCSP)*. IEEE, 2019, pp. 1–7.

[15] Á. López-Raventós and B. Bellalta, "IEEE 802.11be multi-link operation: When the best could be to use only a single interface," *CoRR*, vol. abs/2105.10199, 2021. [Online]. Available: https://arxiv.org/abs/2105.10199

[16] Á. López-Raventós and B. Bellalta, "Multi-link Operation in IEEE 802.11 be WLANs," *arXiv preprint arXiv:2201.07499*, 2022.

[17] T. Song and T. Kim, "Performance Analysis of Synchronous Multi-Radio Multi-Link MAC Protocols in IEEE 802.11 be Extremely High Throughput WLANs," *Applied Sciences*, vol. 11, no. 1, p. 317, 2021.

[18] 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*, 2021, pp. 1–6.

[19] H. Park and C. You, "Latency Impact for Massive Real-Time Applications on Multi Link Operation," in *2021 IEEE Region 10 Symposium (TENSYMP)*, 2021, pp. 1–5.

[20] G. Lacalle, I. Val, O. Seijo, M. Mendicute, D. Cavalcanti, and J. Perez-Ramirez, "Analysis of Latency and Reliability Improvement with Multi-Link Operation over 802.11," in *2021 IEEE 19th International Conference on Industrial Informatics (INDIN)*, 2021, pp. 1–7.

[21] S. Barrachina-Muñoz, B. Bellalta, and E. Knightly, "Wi-Fi All-Channel Analyzer," in *Proceedings of the 14th International Workshop on Wireless Network Testbeds, Experimental evaluation & Characterization*, 2020, pp. 72–79.

[22] S. Barrachina-Muñoz, B. Bellalta, E. W. Knightly *et al.*, "Wi-Fi Channel Bonding: An All-Channel System and Experimental Study From Urban Hotspots to a Sold-Out Stadium," *IEEE/ACM Transactions on Networking*, 2021.