# Can Wi-Fi 7 Support Real-Time Applications? On the Impact of Multi Link Aggregation on Latency

Gaurang Naik, Dennis Ogbe, Jung-Min (Jerry) Park
Bradley Department of Electrical and Computer Engineering, Virginia Tech, Arlington, VA, USA
{gaurang, dogbe, jungmin}@vt.edu

***Abstract*—Multi Link Aggregation (MLA) is a feature likely to be introduced in Wi-Fi 7, the next-generation of Wi-Fi, which will be based on the IEEE 802.11be specifications. MLA will allow Wi-Fi devices that support multiple bands (such as the 2.4 GHz, 5 GHz, and 6 GHz bands) to operate on them simultaneously. The resulting throughput and latency gains are likely to bring Wi-Fi one step closer to supporting emerging real-time applications like augmented and virtual reality. While throughput gains resulting from the use of MLA are mostly linear, the latency gains exhibit interesting characteristics and are the subject of this paper. We use our in-house simulator to study the latency enhancements resulting from MLA and seek to answer whether Wi-Fi 7 devices can meet the challenging latency requirements demanded by most real-time applications. In this pursuit, we observe that allowing Wi-Fi devices to contend on even a single additional link without changing any physical layer parameters can lead to an order of magnitude improvement in the worst-case latency in many scenarios. In addition, we highlight that even in dense conditions, MLA can help Wi-Fi devices meet the challenging latency requirements of most real-time applications.**

## I. Introduction

The global augmented reality (AR) and virtual reality (VR) market is estimated to be worth over $570 billion by 2025 [1]. In the US alone, the market is estimated to be worth over $85 billion [2]. AR/VR and other real-time applications (RTA) place stringent quality of service (QoS) constraints on the underlying communication technology, often demanding high throughput, low latency, and high reliability simultaneously [3]. While 5G is likely to be the main driving technology for wireless AR and VR [4] in outdoor environments, the Wi-Fi community has accelerated its efforts to develop features and techniques that support wireless RTA in next-generation Wi-Fi. This is best demonstrated in the Project Authorization Report of IEEE 802.11be [5], which outlines two of its objectives (among others) to be to define at least one mode of operation that supports (i) a maximum throughput of at least 30 Gbps, and (ii) improved worst-case latency and jitter.

One of the critical features being designed with the aforementioned objectives for next-generation Wi-Fi 7 systems is Multi Link Aggregation (MLA)[1] [6]. MLA belongs to a class of techniques in IEEE 802.11be that are referred to as Multi Link Operations (MLO). The design of MLO is motivated by the fact that most of today's Wi-Fi devices are capable of operating in multiple bands such as the 2.4 GHz band and different parts of the 5 GHz bands. Additionally, the 6 GHz bands (5.925–7.125 GHz) were recently added to the pool of unlicensed bands in the US, while similar considerations are ongoing in Europe [7]. Thus, future Wi-Fi devices are likely to support operations in at least three bands—namely 2.4 GHz, 5 GHz, and 6 GHz. Even though the current 802.11 architecture includes provisions for devices to operate in multiple bands[2], Wi-Fi device operations are currently limited to a single band at a given time. Thus, Wi-Fi devices that support multiple bands cannot use them simultaneously even if channels across different bands are idle and available for transmission, resulting in sub-optimal spectral utilization.

In contrast, IEEE 802.11be devices that support MLA will be allowed to transmit (or receive) frames at the same time on all their supported links. Such devices will be referred to as Multi Link Devices (MLDs). Links aggregated in MLA can be two or more channels in the same band (such as in the 2.4 GHz band) or two or more channels in different bands.

Since MLA enables IEEE 802.11be devices to operate over multiple links at the same time, the operational bandwidth of MLDs increases. Thus, MLA is expected to yield significant throughput gains and these gains are expected to be proportional to the aggregated bandwidth [9]. However, a far more interesting consequence of MLA is the resulting latency reduction. The traditional Medium Access Control (MAC) layer of Wi-Fi uses a truncated binary exponential back-off protocol [10]. Even with MAC enhancements like Enhanced Distributed Channel Access (EDCA)—which prioritizes packets that belong to certain traffic categories—Wi-Fi devices are often unable to meet the stringent latency requirements of RTA. This is especially true for the worst-case latency requirements [5], which are on the order of 10 msec for most RTA [3]. Allowing Wi-Fi devices to concurrently compete for channel access on multiple links enables them to flush the queued packets as soon as any of the links are available, and is expected to yield substantial latency reductions [11].

While the fact that MLA can facilitate lower end-to-end latencies is known, there is a lack of research in how impactful these latency reductions are, and whether they will be sufficient for supporting RTA. To address these gaps, the questions that we seek to answer in this paper are (i) *How much can MLA contribute to decreasing the worst-case*

This work was partially supported by NSF through Grants 1563832, 1642928, and 1822173, and by the industry affiliates of the Broadband Wireless Access and Applications Center (BWAC).

[1]MLA is also referred to as multi-channel aggregation or multi-band aggregation in some references.

[2]This provision was included to support Fast Session Transfer (FST) [8]—a feature that allows IEEE 802.11ad devices to rapidly switch between mmWave bands and sub-6 GHz bands.

*latency in Wi-Fi networks?* and (ii) *Is the latency reduction sufficient to support RTA?* To answer these questions, we have built a custom MATLAB-based simulator that implements the physical (PHY) and MAC layer functionalities of MLA that are consistent with their descriptions in IEEE 802.11be documents. Using this simulator, we have performed a comprehensive study on the performance of MLA under various network configurations and report our findings in this paper. We make the following key observations: (i) unlike throughput gains, latency reductions from MLA seem to be non-linear and less sensitive to parameters such as the aggregated bandwidth and/or the modulation and coding scheme (MCS), (ii) the addition of a second link alone leads to an order of magnitude reduction in the worst-case latency in many scenarios, and (ii) while the addition of subsequent links is beneficial, the resulting reductions in latency seem to follow the law of diminishing marginal returns. We summarize our main contributions below.

- We analyze the MLA performance in the context of its ability to improve the latency in future Wi-Fi networks.
- Using our simulator, we investigate the impact of traditional throughput-enhancing techniques like increasing the bandwidth and/or MCS on Wi-Fi latency and compare their latency reductions with that obtained from MLA.
- We show that even in dense traffic, it is possible for MLA to help future Wi-Fi devices meet the worst-case latency requirement of 10 msec demanded by most RTA.

## II. Related Work

At the time of writing this paper, studies on the performance of MLA are largely limited to industry contributions within TGbe. For example, [9], [12] highlight MLA throughput gains, while [11], [13] report significant latency gains. While the central message of these contributions—*MLA brings throughput and latency gains*—is clear, their quantification remains somewhat unclear due to the lack of consistency across simulation settings and the choice of network parameters.

In addition to the aforementioned references, the basic principles of MLA and its schemes have been described in [6], [7], [14], [15]. However, discussions in these papers are brief and there are no MLA-related performance evaluation studies. Yang et al. [16] study the performance of MLA in 802.11be but restrict their focus to throughput improvements. On the other hand, while Adame et al. [17] study the low-latency enabling features in IEEE 802.11be and describe MLA to be one of them, simulation results in the paper are limited to single-link operations. Furthermore, Avdotin et al. [18] study the ability of IEEE 802.11be to support RTA. However, they restrict their focus to Multi-User Orthogonal Frequency Division Multiple Access (MU OFDMA) in the uplink and propose modifications to the channel access rules for uplink MU OFDMA to provision better support for RTA.

## III. Background

### A. Latency Components in Wi-Fi

Wi-Fi uses Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) as its MAC protocol. CSMA/CA is fundamentally a *listen-before-talk* protocol, wherein a device with pending packets senses the channel and transmits if the channel is idle. The channel is sensed for a minimum duration referred to as the *inter-frame spacing* (IFS)[3] followed by a random backoff procedure. During the backoff phase, the Wi-Fi device picks the value of the *backoff counter* (BO) randomly in [0, CW-1], where CW refers to the parameter *Contention Window*. The device then decrements the value of BO for each idle slot (9 $\mu$sec) and transmits when BO decrements to 0. If the channel is busy, the BO value is frozen until the channel is idle and the device re-enters the backoff phase. CW lies in the range [$CW_{min}$, $CW_{max}$] and is reset/doubled each time the packet transmission succeeds/fails. The success of a particular transmission is inferred when the *acknowledgment frame* (ACK) is received. If transmission attempt(s) for a given packet fail, the packet is re-transmitted until a maximum retry limit is reached. If the packet is not successfully delivered to the receiver even after this limit is reached, it is dropped and the CW of the device is reset to $CW_{min}$. We refer the interested readers to [19] for additional details on the CSMA/CA protocol.

The end-to-end latency in Wi-Fi systems is composed of four major components. These are the queuing, channel access, transmission, and re-transmission latencies[4] as illustrated in Fig. 1. The queuing latency is the difference between the time at which a queued packet moves to the front of the queue and the time at which that packet was queued. Thereafter, the duration until the packet is attempted for the first time constitutes the channel access latency. If the transmission attempt succeeds, the re-transmission latency for the packet is zero. In this case, the only other latency component is the transmission latency, which is the time it takes for the transmitter to send the packet over-the-air and receive the ACK frame from the receiver. However, if the transmission attempt fails the time duration between the start of the first transmission attempt and the start of the final (i.e., successful) transmission attempt constitutes the re-transmission latency.

### B. Channel Contention Rules for MLA in IEEE 802.11be

In principle, MLA in Wi-Fi 7 will be similar to carrier aggregation supported in cellular Long Term Evolution (LTE) [20]. However, unlike cellular networks, Wi-Fi devices operate in the unlicensed bands where channel availability across different links cannot be guaranteed. Consequently, coordinating channel access across different links is a challenging task in Wi-Fi. In the following discussions, we describe how IEEE 802.11be extends the channel contention rules used by contemporary Wi-Fi devices to provision support for MLA.

The protocol for channel contention in MLDs depends on the choice of the MLA scheme. At present, there are three MLA schemes under consideration for IEEE 802.11be [7]. These are (i) independent MLA, (ii) synchronous single-primary MLA, and (iii) synchronous multi-primary MLA.

[3]The value of IFS differs with the type of packet at the transmitter.

[4]We ignore propagation and processing delays because they are much smaller than the four listed latency components.

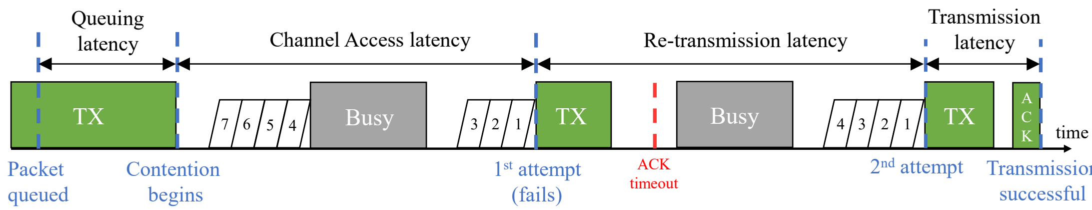


Fig. 1: Illustration of latency components in Wi-Fi; ($CW_{min}$, BO) = (16, 7) for attempt 1, and (32, 4) for attempt 2.

MLDs that support independent MLA perform channel contention independently (i.e., without any coordination) on all links. The contention parameters BO, CW, $CW_{min}$, and $CW_{max}$ are separately maintained and independently updated for each link. This implies that even if an MLD is in the transmit state on one or more links if its queue is non-empty it will continue to contend for channel access on the remaining links. Independent MLA operations for a three-link scenario are illustrated in Fig. 2. Observe that because there is no information exchange across the links, transmissions are unlikely to be synchronized. Therefore, independent MLA is also referred to as asynchronous MLA. Due to this asynchronous transmit and receive behavior of independent MLA, it is suitable when there is a large frequency separation between the aggregated links. If independent MLA is used when the frequency separation between the aggregated links is small, the in-device interference from the transmitting link(s) at the receiving link(s) can be significantly high and can impair the reception of packets on the receiving link(s).

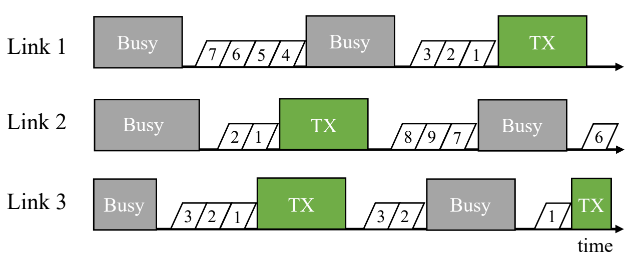


Fig. 2: Independent MLA for three links in IEEE 802.11be.

Synchronous MLA schemes attempt to synchronize transmission attempts across the aggregated links [7]. This is done by extending the channel bonding rules [10] in the IEEE 802.11 specifications, whereby when the MLD gains access to one *primary* link, the status of the other links is checked. If the other links are idle, they are aggregated with the primary link. Synchronous MLA schemes are suitable in configurations where the frequency separation between the links is small. Nevertheless, there are several challenges associated with these schemes—as highlighted in [21], [22]—that must be resolved before they can be adopted in the IEEE 802.11be specifications. Independent MLA is the more likely of the three schemes that will be adopted in Draft 1.0 of IEEE 802.11be specifications, which is expected to be released by May 2021 [23]. Thus, for the remainder of this paper, we limit our attention to independent MLA.

## IV. System Model

### A. System Model

In this paper, we limit our attention to MLA operations in a single Wi-Fi Basic Service Set (BSS). We consider a BSS with one IEEE 802.11be AP, $M_{SL}$ single-link STAs, and $M_{MLD}$ MLDs. These single-link devices can be Wi-Fi STAs that pre-date the IEEE 802.11be (i.e., IEEE 802.11ax and older) specifications or IEEE 802.11be devices that do not support MLO. We assume that all MLDs can contend using independent MLA on $N_L$ links. All STAs are randomly placed in a square region of length 10 m, with the AP placed at the center. A sample realization of the simulation topology with $M_{SL} = M_{MLD} = 5$ is shown in Fig. 3.

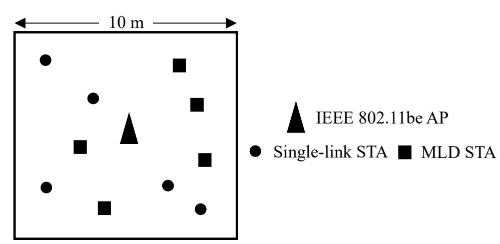


Fig. 3: Simulation Topology

TABLE I: Simulation parameters

| Parameter | Value | Parameter | Value |
|---|---|---|---|
| $N_L$ | 1 to 5 | Max. retry limit | 10 |
| Power (AP) | 24 dBm | Path loss model | WINNER A1 |
| Power (STA) | 14 dBm | Thresh SIR (MCS 2) | 8 dB |
| MCS | 2,10 | Thresh SIR (MCS 10) | 31 dB |
| ED thresh | -62 dBm | Sensitivity (MCS 2) | -90 dBm |
| PD thresh | -82 dBm | Sensitivity (MCS 10) | -67 dBm |
| $L_{MLD}$ | 1-50 Mbps | Load (single-link) | full-buffer |
| Packet size | 1500 Bytes | Rate Adaptation | No |

Arguably, the latency experienced by a particular Wi-Fi device increases when other Wi-Fi devices in the network aggressively compete for the channel. Since the focus of our study is to evaluate MLA performance in such *worst-case* scenarios, we assume that all the single-link devices in the BSS have saturated traffic, i.e., they always have a packet to transmit in their respective queues. Further, we also assume saturated traffic at the AP. To model real-time traffic patterns at the MLD STAs, we assume constant bit-rate (CBR) traffic of $L_{MLD}$ Mbps at the MLD STAs (consistent with references such as [11], [13]), where fixed-sized packets are queued with a fixed inter-arrival time. This inter-arrival time is computed from the offered load at the MLDs. For example, if the fixed packet size is 1500 Bytes and $L_{MLD}$ is 10 Mbps, packets are queued at the MLDs once every 1.2 msec. Finally, for simplicity, we assume that all the relevant parameters (e.g., $N_L$, $L_{MLD}$, etc.) are the same for all MLDs. Simulation parameters used in the following section are outlined in Table I.

## V. Performance Evaluation

### A. The Simulator

To simulate the impact of independent MLA on the end-to-end packet latency, we implemented a modular and configurable software simulator in MATLAB. The traffic generated at all devices and queued at the MAC layer can be individually customized. We model the Wi-Fi MAC layer mechanisms such as the CSMA/CA protocol, channel bonding, frame aggregation, and traffic class prioritization using EDCA [10] for single-link Wi-Fi operations as well as independent MLA. At the PHY layer, we model channel sensing (both physical and virtual [19]), and packet collisions. For each packet transmitted by a Wi-Fi device, the signal-to-interference ratio (SIR) is computed at the receiver. The effects of the PHY layer are abstracted such that a packet is forwarded to the MAC layer if the received signal strength indicator (RSSI) of the packet is above the receiver sensitivity and the SIR of the packet is above an SIR threshold. These RSSI and SIR threshold values are obtained from [24]. If either of these conditions fails (i.e., the RSSI is below the receiver sensitivity, or the SIR is below the threshold value) the packet is dropped at the PHY layer.

In the following subsections, we restrict our performance evaluation to the cumulative distribution function (CDF) of the end-to-end latency experienced by the MLDs' packets. Since the focus of this paper is on the performance of MLA, we do not discuss the performance of the single-link STAs in the BSS. Unless explicitly stated otherwise, in the CDF plots we refer to the *90 percentile latency* as the *worst-case latency*.

### B. Single MLD competing with single-link devices

In this subsection, we assume that there is no traffic class prioritization for the MLD STA's packets, i.e., the MLD STA, the single-link STAs, and the AP use the same channel access parameters (corresponding to Best Effort traffic).

Fig. 4(a) and Fig. 4(b) shows the impact of traditional throughput-enhancing mechanisms such as increasing the channel bandwidth and/or the MCS on the distribution of packet latency and compares them to the gains achieved using MLA. We first look at a scenario where the offered load at the MLD STA is small ($L_{MLD}$ = 1 Mbps) in Fig. 4(a). Observe that when $N_L$ = 1, the 90 percentile latency obtained using a 20 MHz channel and MCS 2 is extremely high (800 msec). While increasing the bandwidth (to 80 MHz) and the MCS (to MCS 10) helps in reducing this worst-case latency (to 550 sec and 300 msec, respectively), we see that adding a second 20 MHz link (with MCS 2) alone yields a far greater latency reduction (the 90 percentile latency is reduced to $\approx$100 msec). Furthermore, when $N_L$ = 2, the additional latency reduction obtained from increasing the bandwidth and the MCS is small.

Fig. 4(b), on the other hand, shows that if the offered load at the MLD STA is high ($L_{MLD}$ = 20 Mbps), adding a second low-bandwidth link does very little to reduce the latency tail. However, a substantial latency gain is observed when the bandwidth and the MCS are increased to 80 MHz and MCS 10, respectively. This is despite the fact that increasing the MCS (from 2 to 10) decreases the resilience of the transmissions to packet collisions (see SIR threshold and sensitivity values in Table I). Adding a second 80 MHz, MCS 10 link results in substantial further reduction in the worst-case latency. Overall, the 90 percentile latency shows an order-of-magnitude improvement from 5 seconds[5] for $N_L$ = 1, 20 MHz, and MCS 2 to 100 msec for $N_L$ = 2, 80 MHz, and MCS 10.

TABLE II: MLD throughput for cases shown in Fig. 4(b).

| **Case I** | **Case II** | **Case III** | **Case IV** | **Case V** |
|---|---|---|---|---|
| **$N_L$ = 1**<br>20 MHz<br>MCS 2 | **$N_L$ = 1**<br>80 MHz<br>MCS 2 | **$N_L$ = 1**<br>80 MHz<br>MCS 10 | **$N_L$ = 2**<br>20+20 MHz<br>MCS 2 | **$N_L$ = 2**<br>80+80 MHz<br>MCS 10 |
| 1.9 Mbps | 7.8 Mbps | 19.1 Mbps | 3.8 Mbps | 19.8 Mbps |

The fundamental difference between the two aforementioned cases is queuing latency. When the MLD competes for channel access with other Wi-Fi devices, the link capacity is divided across all contenders. When $L_{MLD}$ = 1 Mbps, the MLD's share of the link with a single 20 MHz channel and MCS 2 is sufficient to flush the queued packets at a fast enough rate. This can be verified from the throughput of the MLD, which is $\approx$1 Mbps. The predominant latency components, in this case, are the channel access and re-transmission latencies, which are substantially lowered by adding a second 20 MHz link. Consequently, adding a second 20 MHz link alone results in substantial reduction in the worst-case latency.

On the other hand, when $L_{MLD}$ = 20 Mbps, the average MLD throughput in the five cases presented in Fig. 4(b) is given in Table II. Observe in Table II that in Cases I, II, and IV, the capacity requirement of the MLD is not met and MLD throughput $\ll L_{MLD}$ (i.e., 20 Mbps). The only two cases where the MLD throughput $\approx$ 20 Mbps are those where the worst-case latency in Fig. 4(b) is substantially lower (i.e., Cases III and V). Furthermore, observe that even though the MLD throughput is $\approx$ 20 Mbps in Cases III and V, adding a second link (with the same bandwidth and MCS) yields substantial reduction in the end-to-end latency through reductions in the channel access and re-transmission latencies. The 90 percentile latency is reduced from 620 msec when $N_L$ = 1 to 95 msec when $N_L$ = 2. In summary, as long as the capacity requirement of the MLD is met (by increasing the bandwidth, MCS, or both), MLA can yield substantial latency improvements by lowering the channel access and re-transmission latencies.

Next, in Fig. 4(c), we look at the impact of the number of aggregated links, i.e., $N_L$ on the latency. From the above discussions, we know that the queuing latency can be controlled only if the offered load at the MLD, $L_{MLD}$, is smaller than its share of the channel capacity. Thus, in Fig. 4(c), we minimize the impact of queuing by setting $L_{MLD}$ = 1 Mbps.

[5]The root cause for such high latency values is the large number of packet collisions due to many full-buffer contenders (which is 8 including the AP).

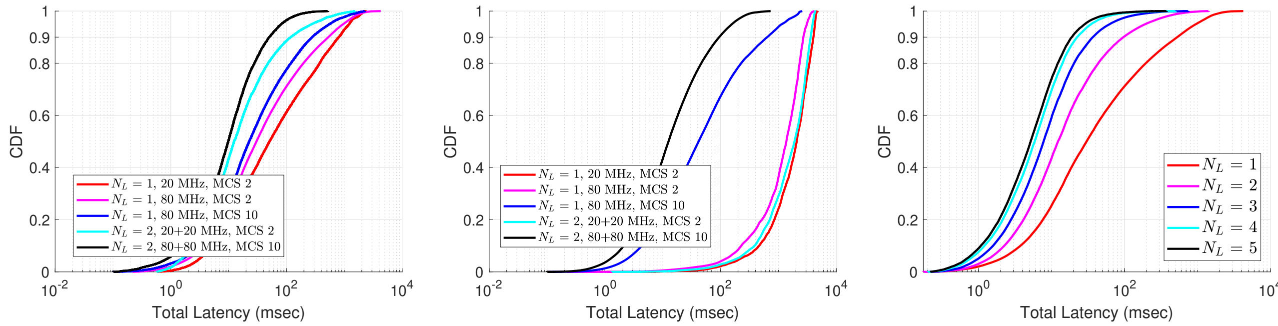


(a) $M_{SL}$ = 7, $M_{MLD}$ = 1, $L_{MLD}$ = 1 Mbps (b) $M_{SL}$ = 7, $M_{MLD}$ = 1, $L_{MLD}$ = 20 Mbps (c) $M_{SL}$ = 7, $M_{MLD}$ = 1, $L_{MLD}$ = 1 Mbps

Fig. 4: CDF of the end-to-end latency of the MLD's packets. One MLD competing for channel access with single-link STAs.

The bandwidth of all but one links is set to 80 MHz and the MCS on all links is set to MCS 2. For $N_L \geq 2$, we set the bandwidth of one link to 20 MHz to emulate a 2.4 GHz link[6]. We observe in Fig. 4(c) that while the addition of each link contributes positively toward the reduction of latency, the relative gains decrease with the addition of each successive link. The 90 percentile latency is reduced from 600 msec in the single-link scenario to 40 msec in the three-link scenario, which is a 93% reduction. However, increasing $N_L$ to 5 yields only a 50% further reduction in the 90 percentile latency.

While increasing the number of links is indeed beneficial in scenarios where the offered load at the MLD is high, the reduction in latency in such cases can be primarily attributed to a reduction in the queuing latency. Addressing capacity issues, which is the fundamental cause for the increased queuing latency, is far more appealing by increasing the channel bandwidth (which is available in abundance in the 6 GHz bands [7]), modulation order, and the number of spatial streams. This is due to two main reasons, (i) increasing the number of radio frequency front-ends available in future Wi-Fi 7 devices to more than 3 may be infeasible, especially in power-limited devices like smartphones and VR headsets, and (ii) the number of links that have large enough frequency gaps and support simultaneous transmission and reception, which is a fundamental requirement for independent MLA, is likely to be three (one each in the 2.4 GHz, 5 GHz, and 6 GHz bands).

### *C. Multiple MLD competing with single-link devices*

While the aforementioned results and discussions provide interesting insights on the worst-case performance of MLA, the setup used in these simulations is not entirely practical for two reasons, (i) the number of MLDs in a BSS is likely to be more than one, especially with the widespread adoption of Wi-Fi 7 and Wi-Fi 7-based RTA, and (ii) in practice, RTA packets are likely to be transmitted with a higher priority than others to satisfy the QoS requirements of such applications. Therefore, in the following discussions, we introduce more than one MLD in the network and observe the resulting impact on the end-to-end latency of the MLDs' packets. At the same time, packets generated at all the MLDs are assigned the highest priority by transmitting these packets using the channel access parameters corresponding to voice traffic.

Fig. 5(a) shows the distribution of the end-to-end latency when three MLDs share the channel with seven single-link STAs. We observe that even in the presence of 7 full-buffer STAs and $L_{MLD}$ = 20 Mbps at each of the three MLDs, the 99 percentile latency of MLDs' packets is just over 10 msec for the single-link scenario. Compared to the results shown in Sec. V-B, this is a substantial improvement and arises from traffic class prioritization used by the MLDs in the transmission of their packets. Furthermore, while the availability of additional links for contention reduces the latency, this reduction is not significant. Thus, in scenarios where the number of MLDs with RTA traffic is small, as long as MLDs use traffic prioritization to transmit their packets the QoS requirements of RTA can be met even when the number of full-buffer contenders (with lower priority traffic) is high.

The benefits of traffic class prioritization notwithstanding, the gains achieved by using this prioritization come at a cost when the density of high-priority traffic increases. This is evident in Fig. 5(b) and Fig. 5(c), which show the latency distribution of MLDs' packets when there are seven MLDs and seven single-link full-buffer STAs in the BSS for $L_{MLD}$ = 20 Mbps and $L_{MLD}$ = 50 Mbps, respectively.

First, we see from Fig. 5(a) and Fig. 5(b) that by increasing the number of MLDs from 3 to 7 for the same offered load, the 90 percentile latency when $N_L$ = 1 jumps 300% from 5 msec to 20 msec. The corresponding increase in the latency when $L_{MLD}$ is increased to 50 Mbps is much worse (Fig. 5(c)). These effects can be attributed to the increase in the number of packet collisions among the MLDs, which lead to very high re-transmission latencies. It must be noted that voice traffic packets are assigned a higher priority through the reduction of the channel access parameters (i.e., IFS, $CW_{min}$, and $CW_{max}$) values, which inevitably leads to a large number of collisions when many such STAs share the channel. Fig. 5(b) shows that 35% of the MLD packets experience 90 percentile latency of greater than 10 msec when $N_L$ = 1 and $L_{MLD}$ = 20 Mbps. However, the addition of a second link alone reduces the 90 percentile latency to 6 msec. On the other hand, even when $N_L$ = 3, more than 90% of the packets experience an end-to-

[6]The typical channel bandwidth in the 2.4 GHz band is 20 MHz.

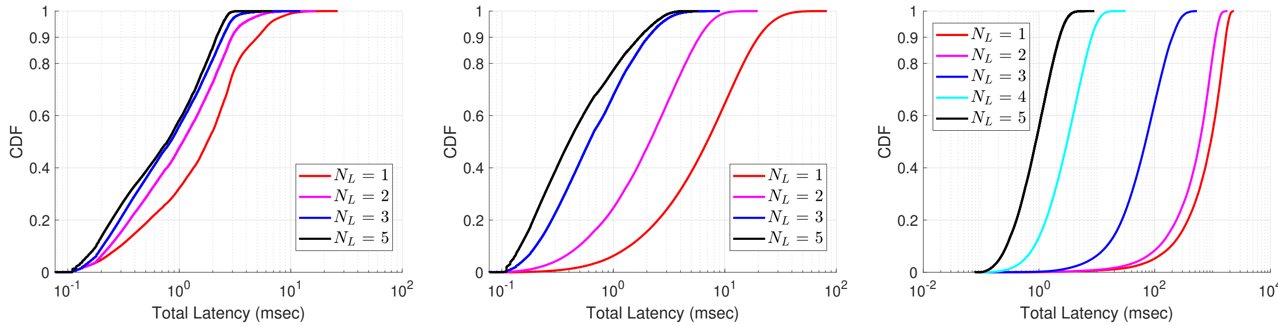


(a) $M_{MLD}$ = 3, $L_{MLD}$ = 20 Mbps (b) $M_{MLD}$ = 7, $L_{MLD}$ = 20 Mbps (c) $M_{MLD}$ = 7, $L_{MLD}$ = 50 Mbps

Fig. 5: Impact of traffic prioritization and multiple MLDs on the MLDs' end-to-end latency, $M_{SL}$ = 7.

end latency of more than 10 msec when $L_{MLD}$ = 50 Mbps. Nevertheless, even at such high RTA traffic densities, the 90 percentile latency can be brought to under 10 msec if $N_L \geq 4$.

## VI. Conclusions & Future Work

In this paper, we investigate the impact of Multi Link Aggregation—a key feature likely to be introduced in Wi-Fi 7—on the packet latency in Wi-Fi networks. The simulation results shown in this paper demonstrate that the use of MLA in future Wi-Fi networks can yield an order-of-magnitude reduction in the worst-case latency experienced by Wi-Fi devices. Furthermore, we highlight that even in dense traffic conditions, MLA can help Wi-Fi devices meet the challenging worst-case latency requirement of 10 msec demanded by most real-time applications. There are three natural extensions to the study conducted in this paper. First, we look at the performance of MLA in a single Wi-Fi 7 BSS. Thus, the exposed node problem, which is known to negatively affect Wi-Fi performance in dense environments, is ignored. Second, we consider only CSMA/CA-based Wi-Fi transmissions, whereas IEEE 802.11ax additionally allows uplink and downlink access using MU OFDMA. Finally, we evaluate the performance of only one of the three MLA schemes, i.e., independent MLA. We will address these shortcomings in our future work.

## References

[1] Allied Market Research, "Augmented and Virtual Reality Market Report." Online: https://www.alliedmarketresearch.com/augmented-and-virtual-reality-market, December 2018.

[2] Allied Market Research, "U.S. Augmented and Virtual Reality Market Report." Online: https://www.alliedmarketresearch.com/us-augmented-and-virtual-reality-market-A06735, August 2020.

[3] K. Meng, A. Jones, and D. Cavalcanti, "IEEE 802.11-19/0065r6: RTA report summary." contribution to the IEEE RTA TIG, January 2019 (accessed February 8, 2021).

[4] Qualcomm Technologies, Inc., "VR and AR pushing connectivity limits." Online: https://www.qualcomm.com/media/documents/files/vr-and-ar-pushing-connectivity-limits.pdf, October 2018.

[5] L. Cariou, "IEEE 802.11-18/0540r2: IEEE P802.11 Wireless LANs: 802.11 EHT Proposed PAR." IEEE 802.11be Project Authorization Report, January 2019 (accessed February 8, 2021).

[6] D. López-Pérez, A. Garcia-Rodriguez, L. Galati-Giordano, M. Kasslin, and K. Doppler, "IEEE 802.11 be extremely high throughput: The next generation of Wi-Fi technology beyond 802.11 ax," *IEEE Communications Magazine*, vol. 57, no. 9, pp. 113–119, 2019.

[7] G. Naik, J.-M. Park, J. Ashdown, and W. Lehr, "Next Generation Wi-Fi and 5G NR-U in the 6 GHz Bands: Opportunities & Challenges," *IEEE Access*, 2020.

[8] E. Perahia, C. Cordeiro, M. Park, and L. L. Yang, "IEEE 802.11 ad: Defining the next generation multi-Gbps Wi-Fi," in *2010 7th IEEE consumer communications and networking conference*, pp. 1–5, IEEE, 2010.

[9] A. Patil, G. Cherian, A. Asterjadhi, and D. Ho, "IEEE 802.11-19/0764r1: Multi-Link Aggregation: Peak Throughput Gains." contribution to the IEEE TGbe, May 2019 (accessed February 8, 2021).

[10] G. Naik, J. Liu, and J.-M. J. Park, "Coexistence of wireless technologies in the 5 ghz bands: A survey of existing solutions and a roadmap for future research," *IEEE Communications Surveys & Tutorials*, vol. 20, no. 3, pp. 1777–1798, 2018.

[11] A. Patil, G. Cherian, A. Asterjadhi, and D. Ho, "IEEE 802.11-19/1081r1: Multi-Link Aggregation: Latency Gains." contribution to the IEEE TGbe, August 2019 (accessed February 8, 2021).

[12] Y. Seok, W. Tsao, G. Bajko, J. Yee, J. Liu, P. Cheng, and T. Pare, "IEEE 802.11-19/0766r1: Enhanced Multi-band/Multi-channel Operation." contribution to the IEEE TGbe, May 2019 (accessed February 8, 2021).

[13] Y. Seok, W. Tsao, G. Bajko, J. Yee, J. Liu, P. Cheng, and T. Pare, "IEEE 802.11-19/0979r2: EHT Multi-Link Operation Follow-up." contribution to the IEEE TGbe, September 2019 (accessed February 8, 2021).

[14] E. Khorov, I. Levitsky, and I. F. Akyildiz, "Current Status and Directions of IEEE 802.11 be, the Future Wi-Fi 7," *IEEE Access*, 2020.

[15] C. Deng, X. Fang, X. Han, X. Wang, L. Yan, R. He, Y. Long, and Y. Guo, "IEEE 802.11 be–Wi-Fi 7: New Challenges and Opportunities," *IEEE Communications Surveys & Tutorials*, 2020.

[16] 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)*, pp. 1–7, IEEE, 2019.

[17] T. Adame, M. Carrascosa, and B. Bellalta, "Time-Sensitive Networking in IEEE 802.11 be: On the Way to Low-latency WiFi 7," *arXiv preprint arXiv:1912.06086*, 2019.

[18] E. Avdotin, D. Bankov, E. Khorov, and A. Lyakhov, "Enabling Massive Real-Time Applications in IEEE 802.11 be Networks," in *2019 IEEE 30th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, pp. 1–6, IEEE, 2019.

[19] M. Burton and G. Hill, "802.11 Arbitration," *White Paper, Certified Wireless Network Professional Inc., Durham, NC*, 2009.

[20] G. Yuan, X. Zhang, W. Wang, and Y. Yang, "Carrier aggregation for LTE-advanced mobile communication systems," *IEEE communications Magazine*, vol. 48, no. 2, pp. 88–93, 2010.

[21] Y. Seok, J. Yee, J. Liu, and T. Pare, "IEEE 802.11-19/1305r4: Synchronous Multi-Link Operation." contribution to the IEEE TGbe, January 2020 (accessed February 8, 2021).

[22] S. Naribole, S. Kandala, W. Lee, and A. Ranganath, "IEEE 802.11-19/1405r7: Multi-link Channel Access Discussion." contribution to the IEEE TGbe, November 2019 (accessed February 8, 2021).

[23] A. Garcia-Rodriguez, D. Lopez-Perez, L. Galati-Giordano, and G. Geraci, "IEEE 802.11 be: Wi-Fi 7 Strikes Back," *arXiv preprint arXiv:2008.02815*, 2020.

[24] D. J. Urban, "The Importance Of Wi-Fi 6 Technology For Delivery Of gbps Internet Service," *SCTE ISBE Cable-Tec Expo*, September 2019.