

The Role of MLO in Achieving Deterministic Wi-Fi

Roger Sanchez-Vital^{*†} and Eduard Garcia-Villegas[†]

^{*}IDLab, University of Antwerp – imec, Antwerp, Belgium
roger.sanchezvital@uantwerpen.be

[†]Universitat Politècnica de Catalunya – BarcelonaTech (UPC), Barcelona, Spain
eduardo.garcia@upc.edu

Abstract—The emergence of next-generation wireless applications has shifted the focus of Wireless Local Area Networks (WLANs) from peak throughput toward deterministic performance and Ultra High Reliability (UHR). While IEEE 802.11be (Wi-Fi 7) introduced Multi-Link Operation (MLO) primarily for bandwidth aggregation, its potential role in achieving Time-Sensitive Networking (TSN) requirements is critical for industrial applications. In this paper, we propose an analytical model to evaluate MLO-based frame duplication, focusing on the interplay between Packet Delivery Ratio (PDR), transmission delay, and energy efficiency. Our results, obtained across four representative deployment scenarios, demonstrate that MLO is essential to approach the “five nines” (99.999%) reliability target required for industrial control and can reduce latency up to 6% compared to single-link configurations. However, these improvements come with a measurable increase in energy per bit of up to 4x, revealing a critical trade-off for battery-constrained devices. These results motivate the use of the proposed model as a basis for developing dynamic MLO management algorithms aimed at meeting delay and reliability requirements while minimizing energy costs.

Index Terms—Wi-Fi, IEEE 802.11, MLO, TSN

I. INTRODUCTION

Industrial automation, immersive applications, and other delay-sensitive services increasingly require Wireless Local Area Networks (WLANs) to provide high reliability and bounded latency in addition to high throughput [1]. Traditionally, wireless links, especially in the context of spectrum sharing in license-free bands, have been considered too unreliable for Time-Sensitive Networking (TSN) environments, which require guaranteed delivery within strict temporal bounds [2], [3]. IEEE 802.11be (Wi-Fi 7) addresses part of this challenge through Multi-Link Operation (MLO), which allows a device to operate over multiple links and exploit frequency diversity and redundancy.

The implementation and benefits of MLO depend heavily on the radio capabilities of the station (STA) and the access point (AP). In the ideal configuration, known as Simultaneous Transmit and Receive (STR), a device features multiple independent radio chains, enabling it to transmit on one link while receiving on another simultaneously [1]. This full-duplex-like operation prevents the unavailability that occurs when a radio cannot listen while transmitting, which is crucial for minimizing

latency in real-time applications [4]. Conversely, for size- and power-constrained devices (e.g., smartphones), non-STR modes like Enhanced Multi-link Single-radio (EMLSR) allow nodes to listen on multiple links simultaneously to contend for access, but restrict data transmission or reception to a single link at a time [5]. This allows constrained devices to select the best available channel dynamically, improving reliability over static single-link operations without the hardware costs of STR.

While Wi-Fi 7 introduced MLO primarily to aggregate bandwidth and reduce latency, the upcoming IEEE 802.11bn (Wi-Fi 8) shifts the paradigm toward Ultra High Reliability (UHR) to ensure quasi-deterministic connectivity [6]. A key innovation in this evolution is Distributed MLO. By placing multiple APs under the control of the same Multi-Link Device (MLD), this framework allows a multi-link STA to switch links with minimal signaling overhead and delay. This implicitly enables seamless transitions between APs, maintaining UHR performance in dense deployments by ensuring consistent latency as users move across coverage areas.

Recent studies show that simultaneous transmission across multiple links can reduce latency compared with single-link operation [4], [5]. QoS-aware work has also considered cooperative link allocation for industrial service-level agreements [7] and multi-link scheduling for high-priority, low-latency traffic [8]. However, redundant operation such as frame duplication introduces a reliability–energy trade-off that remains less explored than throughput and latency.

In this paper, we propose an analytical model to evaluate the performance of MLO-based frame duplication. The model captures the interplay between reliability, transmission delay, duplicated packet probability, and energy consumption as a function of the number and characteristics of active links and retransmission limits. Beyond characterizing the benefits and costs of MLO, the proposed model can assist the design of dynamic MLO management mechanisms that adapt the number of links in use and the retransmission configuration to satisfy reliability and delay requirements while minimizing energy costs. We evaluate the model across representative deployment scenarios and quantify the resulting trade-off between ultra-high reliability, latency reduction, and energy efficiency.

The rest of this paper is organized as follows. Section II presents the analytical model, Section III evaluates its behavior, Section IV analyzes the considered deployment scenarios, and Section V concludes the paper.

This research was funded in part by the European Union’s Horizon Europe research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 101282240, and by the Spanish MCIU/AEI/10.13039/501100011033/FEDER/ UE through project PID2023-146378NB-I00.




Fig. 1. System diagram of a multi-link communication between a transmitter and a receiver.

II. ANALYTICAL MODEL

This study investigates the performance of MLO, specifically focusing on the deployment of frame duplication mechanisms to enhance link reliability. The analytical model considers a scenario in which an STA communicates with an AP over multiple links, as depicted in Fig. 1. To accurately reflect real-world wireless dynamics, the model explicitly incorporates the effects of signal fading and external interference. Ultimately, the primary objective is to evaluate the system's overall efficacy by analyzing the interplay between the individual Frame Error Rates (FER), expected delays, and aggregate energy per bit across the available links.

To evaluate the performance of the system, we model the packet delivery ratio (PDR), duplicated packet probability (P_{dup}), delay, and energy per bit.

A. Packet Delivery Ratio (PDR)

PDR is derived from FER and Packet Error Rate (PER). The FER represents the error probability of a given link measured at the frame level, whereas the PER signifies the probability of an unsuccessful packet delivery after r retransmission attempts. In multi-link topologies, the aggregate PER across all links can be calculated alongside the corresponding number of lost frames. For an individual link l , the packet loss probability is determined by the link's FER_l and the maximum number of permitted retransmissions, r_l , as expressed in Equation (1).

$$PER_l = FER_l^{r_l+1} \quad (1)$$

Consequently, the probability of a successful transmission on that specific link is: $PDR_l = 1 - PER_l$. Finally, considering the set of L independent links connecting the two STAs, the aggregate PER (PER_{total}) is derived from the individual error probabilities of each link l (cf. Equation (2)).

$$PER_{total} = \prod_l PER_l \quad (2)$$

With which we can define $PDR_{total} = 1 - PER_{total}$, representing the overall PDR of the multi-link system. These equations assume statistically independent frame errors across retransmission attempts and across links, respectively.

B. Duplicated Packet Probability

When employing frame duplication techniques, it is essential to evaluate the duplicated packet probability, defined as the likelihood of a given frame being successfully delivered multiple times. Because receiving a packet across multiple interfaces introduces unnecessary network overhead, analyzing the total number of duplicated frames is of significant interest.

Based on the individual link error rates, the probability that exactly one copy of a frame successfully reaches the destination is given by Equation (3).

$$P_{1copy} = \sum_{l=1}^L \left[(1 - PER_l) \prod_{\substack{j=1 \\ j \neq l}}^L PER_j \right] \quad (3)$$

Subsequently, the probability that a packet is duplicated (i.e., successfully delivered via more than one link) can be determined as follows: $P_{dup} = 1 - PER_{total} - P_{1copy}$

C. Delay

In addition to evaluating PDR and P_{dup} , transmission delay is a critical metric for evaluating system performance. The overall delay depends heavily on the specific timing intervals defined by the IEEE 802.11 standard, varying based on both fixed protocol constants and the dynamic number of retransmissions required for a successful delivery [9].

The total time required for a single frame transmission attempt, T_{frame} , is a composite of fixed intervals (such as DCF Interframe Space (DIFS) and Short Interframe Space (SIFS)) and dynamic durations. For a given retransmission attempt r , this time can be computed using the expression in Equation (4).

$$T_{frame}(r) = DIFS + T_{backoff}(r) + T_{data} + SIFS + T_{ACK} \quad (4)$$

The calculation of the transmission time requires a detailed breakdown of its dynamic components. Specifically, the data transmission time, denoted as T_{data} (in μs), represents the total duration required to transmit the frame over the medium. This encompasses PHY layer parameters, such as the preamble duration ($T_{Preamble}$), and packet extension (T_{PE}). By factoring in the OFDM symbol duration (T_{SYM}) and the total number of symbols, T_{data} is formally expressed in Equation (5):

$$T_{data} = 20 + T_{Preamble} + N_{SYM} \cdot T_{SYM} + T_{PE} \quad (5)$$

Central to this calculation is N_{SYM} , which defines the total number of Orthogonal Frequency-Division Multiplexing (OFDM) symbols in the data field. For a single-user Extremely High Throughput (EHT) Physical Protocol Data Unit (PPDU), N_{SYM} is computed as:

$$N_{SYM} = \left\lceil \frac{8 \cdot L_{payload} + N_{service}}{N_{DBPS}} \right\rceil \quad (6)$$

where N_{DBPS} represents the number of data bits per symbol, which depends on the configured PHY rate. Comprehensive specifications and parameter details can be found in Clause 36.4.3 of the IEEE 802.11be standard in [9].

Similarly, T_{ACK} refers to the time spent transmitting the acknowledgment frame. It is calculated using the same base formula as T_{data} , but adjusted for the specific ACK frame size. Another critical dynamic component is the backoff duration, $T_{backoff}$, which represents the wait time applied after the channel is sensed busy or a collision occurs. This value is intrinsically linked to the number of retransmissions, r . The system must first determine the current Contention Window, $CW(r)$, based on the minimum contention window, CW_{min} , and the current retransmission attempt (cf. Equation (7)).

$$CW(r) = \min(2^r \cdot CW_{min}, CW_{max}) \quad (7)$$

In practice, the backoff time is determined by randomly choosing a number of slots uniformly distributed between 0 and CW . Therefore, the average backoff time after r retransmissions is given by Equation (8).

$$T_{backoff}(r) = \frac{(CW(r) - 1)}{2} \cdot T_{slot} \quad (8)$$

To evaluate the total time elapsed after multiple failed attempts, T_{frame_y} is defined in Equation (9) as the cumulative time spent after y retransmissions of the same packet. Because the contention window increases with each failed attempt, this cumulative time grows with the number of retransmissions and with the channel access time, T_{access} , which depends on the level of activity in the channel and denotes the average time the STA spends deferring due to a busy medium while waiting to send a frame. As a reference, under saturated conditions with multiple contending STAs, T_{access} can be obtained from Bianchi's model [10]. In practice, T_{access} can be measured at the Wi-Fi device driver level.

$$T_{frame_y} = \sum_{r=0}^{r=y} T_{frame}(r) + T_{access} \quad (9)$$

Finally, the theoretical expected delay of a successfully delivered frame relies on this cumulative transmission time (T_{frame_y}), combined with the frame error rate of the link (FER_l) and the maximum allowed number of retransmissions (r_l). For a standard one-link system ($L = 1$), the theoretical probabilistic delay for link i (in μs) is expressed in Equation (10).

$$D_i = \frac{[(1 - FER_l) \cdot \sum_{y=0}^{r_l} FER_y \cdot T_{frame_y}] \mu s}{\sum_{x=0}^{r_l} (1 - FER_l) \cdot FER_x} \quad (10)$$

For frame replication over multiple links, the same per-link retransmission process can be applied to each active link. Since a packet is delivered as soon as the first copy is successfully

received, the multi-link delay corresponds to the expected earliest successful reception time across links. The closed-form derivation, which requires enumerating all combinations of failed and successful retransmission attempts across links, is omitted due to space constraints and because it provides limited additional insight.

D. Energy per bit

Building upon delay analysis, evaluating the energy consumption of the network requires mapping the previously defined temporal intervals to the corresponding operational states of the radio interface. The power consumption profile of an STA varies significantly depending on whether it is actively transmitting (P_{tx}), actively receiving (P_{rx}), or in an idle state listening to the channel (P_{idle}).

To determine the energy consumed during a single transmission attempt y , the standard IEEE 802.11 time components must be multiplied by their respective power states. The data frame transmission (T_{data}) draws P_{tx} , while the reception of the ACK (T_{ACK}) consumes P_{rx} . All remaining intervals, including mandatory waiting periods (DIFS, SIFS) and dynamic delays ($T_{backoff}(y)$ and T_{access}), operate at the idle power level, P_{idle} . Thus, Equation (11) denotes the energy consumed during a single attempt y .

$$E_{attempt}(y) = T_{data}P_{tx} + T_{ACK}P_{rx} + (DIFS + SIFS + T_{backoff}(y) + T_{access})P_{idle} \quad (11)$$

When calculating the energy consumption for a specific link l , it is crucial to highlight that energy is expended during every transmission attempt, regardless of whether the frame is ultimately delivered successfully or dropped after exceeding the retry limit. For MLO configurations, Equation (11) sums the energy consumed by each active radio, including its channel-access and retransmission periods. Energy associated with inactive radios, MLO control signaling, and background operation outside the considered packet transaction is not included. Since an STA continues to transmit until it either succeeds or reaches the maximum number of retransmissions r_l , a subsequent transmission attempt y only occurs if all previous y attempts have failed. The probability of executing attempt y is exactly FER_l^y . Consequently, the total energy consumed by link l across all potential retransmissions is defined in Equation (12).

$$E_l = \sum_{y=0}^{r_l} FER_l^y \cdot E_{attempt}(y) \quad (12)$$

To provide a standardized metric of network efficiency, the overall energy performance is evaluated in terms of the energy per successfully delivered bit, denoted as E_{bit} . This metric is calculated by dividing the sum of the energy consumed across all l links by the number of bits successfully delivered to the destination, as defined in Equation (13), where $L_{payload}$ is the length of the payload in bytes.

TABLE I
SYSTEM'S CONSTANT AND LINK-SPECIFIC PARAMETERS.

Constant Parameters	Value
CW_{min}/CW_{max}	16 / 1024
$DIFS$	34 μs
$SIFS$	16 μs
T_{slot}	9 μs
$T_{preamble}$	20 μs
T_{SYM}	14.4 μs
T_{PE}	12 μs
$L_{payload}$	1500 bytes
$N_{service}$	16 bits
Frames per simulation	1000

Link-specific Parameters	2.4 GHz	5 GHz	6 GHz
Data Rate (Mbps)	65	130	272
Retransmissions	5	3	1
Tx Power (W)	0.308	0.308	0.308
Rx Power (W)	0.110	0.154	0.220
Idle Power (W)	0.055	0.066	0.0825

$$E_{bit} = \frac{\sum_{l=1}^L E_l}{L_{payload} \cdot PDR_{total}} \quad (13)$$

III. MODEL EVALUATION

To verify the numerical implementation of the analytical model, we developed a Monte Carlo simulator for direct data exchange between two MLDs. The link parameters follow the guidelines in [11]. As summarized in Table I, the framework covers the 2.4, 5, and 6 GHz bands, each assigned specific operational characteristics to reflect realistic network conditions. To ensure precise energy consumption modeling throughout the direct data exchange, the simulator incorporates band-specific power metrics for Tx, Rx, and idle states. Because the simulator implements the same probabilistic assumptions as the analytical model, this comparison is intended as an implementation-consistency check rather than independent validation; validation with measured, temporally correlated channels remains future work.

A. Packet Delivery Ratio (PDR)

Fig. 2 illustrates PDR against FER. The data captures how individual frame errors, which may result from varying signal quality or environmental interference, impact the overall success of packet transmissions. The results are categorized into three distinct scenarios involving one, two, and three active links. As the FER increases from 0.0 toward 1.0, the PDR remains stable at a value of 1.0 before eventually declining. This decline occurs at different thresholds depending on the number of links utilized. In a single-link configuration, the PDR begins to drop significantly once the FER exceeds approximately 0.4. However, the multi-link configurations exhibit much higher resilience; the two-link setup maintains near-perfect delivery until an FER of 0.6, while the three-link setup remains robust until the FER approaches 0.8. This clear rightward shift in the curves demonstrates that increasing the number of links through MLO effectively mitigates packet loss

by providing redundant paths for data. The plotted simulation points reproduce the analytical curves, providing the intended implementation-consistency check.




Fig. 2. PDR vs FER for one, two, and three links.

B. Duplicated Packet Probability

Fig. 3 represents the duplicated packet probability as a function of the FER. Under ideal channel conditions where the FER is 0.0, the duplication probability is 1.0, indicating that all available links are successfully delivering the redundant copies of the data. As the FER increases, the likelihood of receiving multiple copies of the same packet begins to decline, though the rate of this decline depends heavily on the number of links in use. The three-link configuration demonstrates significantly higher resilience than the two-link setup; for example, at an FER of 0.6, the duplication probability for two links drops to approximately 0.6, whereas the three-link scenario maintains a much higher probability near 0.8. The simulation points reproduce the analytical curves under the same assumptions and illustrate the redundancy cost accompanying the reliability gain.

C. Delay

Fig. 4 analyzes transmission delay as a function of FER, the number of active links, and the initial access delay Acc . An access delay of 0 represents an idealized baseline, while $T_{access} = 579 \mu s$ and $865 \mu s$ correspond to 10 and 15 contending STAs, respectively, according to the Bianchi model.

Results demonstrate that, for any given access time, the configuration with three links consistently exhibits the lowest delay, followed by two links and a single link. Configurations with a high access time (e.g., $865 \mu s$) start off with a much higher baseline delay and struggle more as errors increase compared to the configurations with zero access time. The value of MLO becomes especially noticeable as channel conditions worsen; while the single-link delay climbs steeply toward 4 ms as the FER approaches 1.0, the multi-link configurations exhibit a much shallower curve, effectively




Fig. 3. Duplicated packet probability vs FER for one, two, and three links.

buffering the system against high error rates. For readability, Figs. 4 and 5 show only simulation results; the analytical curves are omitted because they coincide with the simulation results in all evaluated cases.




Fig. 4. Delay vs FER for one, two, and three links.

D. Energy consumption

Fig. 5 illustrates the energy per bit, measured in Joules per bit and scaled by 10^{-7} , relative to the FER. A critical trade-off is revealed here: while adding links reduces latency, it explicitly increases the energy cost per delivered bit. For example, the three-link configuration with an access delay of $865 \mu\text{s}$ consistently consumes the most energy across all FER values, starting near 0.3×10^{-7} J/bit. In contrast, the single-link simulation under ideal $0.0 \mu\text{s}$ access conditions represents the most energy-efficient baseline.

Results show that energy consumption remains relatively stable at low error rates but experiences an exponential surge

as the FER exceeds 0.8. This sharp vertical climb indicates that when channel conditions are extremely poor, the energy required for constant retransmissions across multiple active radios becomes massive, reaching peaks near 2.5×10^{-7} J/bit. Furthermore, results highlight that higher access delays increase the minimum required energy, meaning that the longer a device must wait to access the medium, the more energy it wastes, regardless of the number of links used.




Fig. 5. Energy per bit vs FER for one, two, and three links.

IV. TO WHAT EXTENT CAN MLO ENHANCE RELIABILITY, DELAY, AND ENERGY CONSUMPTION?

This section evaluates four representative deployment scenarios (Ultra-Dense, Dense, Moderate, and Sparse) to determine if configurations meet industrial TSN reliability and latency requirements. They span progressively lighter contention and link-error conditions: U is the number of saturated contending STAs used to obtain T_{access} from Bianchi's model, while the FER values represent increasingly favorable band-specific link qualities. Simulations use 1000-byte packets, at most four retransmissions, and 100,000 packets; the 2.4, 5, and 6 GHz rates are 73.1, 146.3, and 306.2 Mbps. The QoS targets are 99.999% PDR and 5 ms maximum delay (late frames are discarded) [3]. The MLO manager exhaustively selects the active links and retry limits that minimize energy per bit while meeting these constraints. The configurations selected by the MLO manager are reported in Table II. The one-link configuration uses the 6-GHz link, the two-link configuration uses the 5- and 6-GHz links, and the three-link configuration uses all available links.

To assess adaptive selection, Table III compares the manager-selected configuration with two simple baselines: the best single link and simultaneous use of all three links. Under the assumed error model, the single-link baseline misses the PDR target in all four operating points. In the Ultra-Dense case it yields 99.795%, whereas the selected two-link configuration reaches the tabulated 99.999% level at a lower energy cost than the all-links baseline.

TABLE II
SCENARIO LINK CONFIGURATIONS

Scenario	Total U	Per-band U/FER		
		2.4 GHz	5 GHz	6 GHz
Ultra-Dense	33	15/0.20	10/0.15	8/0.13
Dense	27	13/0.13	8/0.09	6/0.08
Moderate	26	12/0.13	8/0.10	6/0.09
Sparse	20	10/0.11	6/0.09	4/0.08

TABLE III
SCENARIO PERFORMANCE METRICS. BOLD ROWS ARE SELECTED BY THE MLO MANAGER; ONE AND THREE LINKS ARE THE SINGLE-LINK AND ALWAYS-ALL-LINKS BASELINES.

Scenario	Links	Retr.	Delay (ms)	E_{bit} (J/bit $\times 10^{-8}$)	PDR
Ultra-Dense	1	4	1.609	1.903	0.99795
	2	2	1.511	4.325	0.99999
	3	2	1.515	8.731	0.99999
Dense	1	4	1.164	1.421	0.99990
	2	2	1.127	3.325	1.00000
	3	1	1.127	6.833	0.99999
Moderate	1	4	1.184	1.447	0.99970
	2	2	1.135	3.366	1.00000
	3	2	1.136	6.692	1.00000
Sparse	1	4	0.825	1.072	0.99998
	2	4	0.801	2.598	1.00000
	3	1	0.802	5.361	0.99999

Latency results reveal lighter differences among deployment scenarios. The Ultra-Dense environment presents the highest delays due to severe contention and retransmissions, reaching 1.609 ms in the single-link case. With MLO, delay decreases to around 1.51 ms, corresponding to a reduction of approximately 6%. In the Dense, Moderate, and Sparse scenarios, latency reductions are smaller but still consistent, ranging between 3% and 4%. Although the latency gains are moderate compared to the reliability improvements, the results confirm that MLO can reduce retransmission overhead and improve transmission responsiveness.

Energy efficiency remains the main trade-off associated with MLO operation. Across all scenarios, energy consumption increases significantly as additional links are activated. In the Ultra-Dense scenario, the energy per bit increases from 1.903×10^{-8} J/bit in single-link mode to 4.325×10^{-8} J/bit and 8.731×10^{-8} J/bit with two and three links, corresponding to increases of approximately 127% and 359%, respectively. Similar trends are observed in the remaining scenarios, where energy consumption increases by roughly 130–140% with two links and up to 400% with three links. This reflects the additional cost of operating multiple radios and duplicating transmissions to maximize reliability.

V. CONCLUSIONS

This paper highlights the role of MLO in transitioning Wi-Fi from a throughput-centric technology toward a near-

deterministic framework capable of supporting TSN. The proposed analytical model and subsequent simulations demonstrate that MLO-based frame duplication can be an effective mechanism for industrial applications requiring ultra-high reliability. While single-link configurations often fall short in congested environments, the evaluated multi-link configurations reach the 99.999% PDR target across all tested scenarios. Parallel transmissions reduce average delay by up to 6%, but using multiple radios and redundant traffic increases energy per bit by up to 4x. Consequently, while MLO provides the redundancy needed to overcome the stochastic nature of the wireless medium, future research must focus on adaptive resource management to balance these reliability gains with the power constraints of mobile industrial devices as the wireless industry moves toward the UHR targets of IEEE 802.11bn.

REFERENCES

- [1] 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*, vol. 32, no. 1, pp. 308–322, 2024.
- [2] D. Cavalcanti, C. Cordeiro, M. Smith, and A. Regev, "Wi-Fi TSN: Enabling Deterministic Wireless Connectivity over 802.11," *IEEE Communications Standards Magazine*, vol. 6, no. 4, pp. 22–29, 2022.
- [3] M. K. Atiq, R. Muzaffar, O. Seijo, I. Val, and H.-P. Bernhard, "When IEEE 802.11 and 5G Meet Time-Sensitive Networking," *IEEE Open Journal of the Industrial Electronics Society*, vol. 3, pp. 14–36, 2022.
- [4] 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.
- [5] 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 *2023 IEEE 34th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, 2023, pp. 1–6.
- [6] L. Galati-Giordano, G. Geraci, M. Carrascosa, and B. Bellalta, "What will Wi-Fi 8 Be? A Primer on IEEE 802.11bn Ultra High Reliability," *IEEE Communications Magazine*, vol. 62, no. 8, pp. 126–132, 2024.
- [7] S. Kumar, D. Camps-Mur, and E. Garcia-Villegas, "C-SLA-MLO: Enhancing SLA Compliance in Industrial Wi-Fi through Cooperative Multilink Operation," *Internet of Things*, vol. 27, p. 101269, 2024.
- [8] L. Gan, B. Li, M. Yang, and Z. Yan, "Low-latency guarantee protocol based on multi-links scheduling random access in the next generation wlan: Ieee 802.11be," in *Smart Grid and Internet of Things*, Y.-B. Lin and D.-J. Deng, Eds. Cham: Springer International Publishing, 2021, pp. 286–296.
- [9] *IEEE 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 2: Enhancements for Extremely High Throughput (EHT)*, IEEE Std. IEEE 802.11be-2024, 2024, accessed: 2026-05-12.
- [10] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 3, pp. 535–547, 2000.
- [11] IEEE 802.11 Working Group, "Simulation Scenarios," <https://mentor.ieee.org/802.11/dcn/14/11-14-0980-16-00ax-simulation-scenarios.docx>, 2014, IEEE 802.11-14/0980r16.