

Admission control and scheduling of periodic real-time flows in IEEE 802.11be networks

Mattia Pirri^{*}, Zenepe Satka[†], Luca Leonardi^{*}, Mohammad Ashjaei[†], Gaetano Patti^{*}, Lucia Lo Bello^{*}

^{*}University of Catania, Catania, Italy

[†]Mälardalen University, Västerås, Sweden

mattia.pirri@phd.unict.it, {luca.leonardi, gaetano.patti, lucia.lobello}@unict.it, {zenepe.satka, mohammad.ashjaei}@mdu.se

Abstract—This work explores how to leverage two novel features provided by the IEEE 802.11be (Wi-Fi 7) standard, such as the Restricted Target Wake Time and Multi-Link Operation, to offer support to real-time periodic flows. In fact, the standard does not define mechanisms to effectively integrate these features in admission control and scheduling policies for real-time flows. The work first addresses the design challenges behind providing bounded latency guarantees and efficient link utilization to real-time uplink flows and then outlines two strategies, called Dedicated and Shared R-TWT. Finally, the work compares the results of a preliminary analytical assessment of the proposed strategies.

Index Terms—IEEE 802.11be, real-time traffic scheduling, Restricted Target Wake Time (R-TWT)

I. INTRODUCTION

The rapid evolution of Industry 4.0 is driving an increasing demand for wireless networks that support high data rates and bounded latencies in time-critical industrial applications, such as predictive maintenance, real-time monitoring and control of machines, mobile robots, and automated guided vehicles (AGVs). To address the requirements of these time-critical applications, the emerging IEEE 802.11be amendment [1] (marketed as Wi-Fi 7) introduces innovative features, such as the Restricted Target Wake Time (R-TWT) and Multi-Link Operation (MLO). R-TWT provides the ability to protect time-critical traffic from contention-based traffic reserving specific time windows for real-time traffic transmission. MLO enables a Multi-Link Device (MLD) to transmit and receive simultaneously across multiple links operating on different frequency bands (e.g., 2.4, 5, and 6 GHz). By dynamically switching or aggregating traffic across multiple links, MLO enhances throughput, reduces latency, and increases reliability.

However, while the potential benefits of R-TWT and MLO are clear, the standard does not specify the logic required to exploit them in an admission control and scheduling policy for real-time flows. While recent literature has extensively investigated R-TWT to guarantee bounded latencies and balance energy consumption [2]–[5], the design of admission control strategies that schedule R-TWT on multiple links in a bandwidth-efficient way is still largely unexplored [6]. The main challenge is to schedule real-time flows in non-overlapping guaranteed time windows. In this context, this work makes a step towards novel solutions for admission control and scheduling of periodic real-time flows in IEEE 802.11be networks leveraging the R-TWT and MLO mechanisms. In particular, the work explores different solutions for admission and




Fig. 1: Trigger-enabled R-TWT operation.

scheduling of periodic real-time uplink flows in IEEE 802.11be networks, discusses the relevant design challenges, and presents the results of preliminary performance assessments.

II. BACKGROUND ON R-TWT

R-TWT envisages that an Access Point (AP) and a Station (STA) negotiate specific time windows, known as Service Periods (SPs), for data exchange. As shown in Fig. 1, the negotiation defines the *Wake Interval* (determining how often the STA wakes up), the *Wake Duration* (i.e., the duration of the SPs), and the *Target Wake Time* (TWT) (identifying the absolute start time of the first SP). The Wake Duration is expressed in discrete units (called wake duration unit) of either a Time Unit (TU), i.e., 1024 μ s, or a fraction of a TU (i.e., 256 μ s), whereas the TWT is strictly quantized in increments of 1024 μ s [1].

To control the channel access within an SP, an R-TWT can be configured to be a trigger-enabled TWT (as shown in Fig. 1). In the trigger-enabled TWT configuration, transmissions are initiated by the AP via Trigger Frames (TF), thus avoiding the contention and collision for real-time flows.

To achieve real-time traffic isolation, the AP advertises the schedule by embedding specific *TWT Elements* within its periodic management frames (i.e., Beacons). These R-TWT elements broadcast to the network the exact temporal boundaries of each active R-TWT schedule, i.e., the timing parameters TWT, Wake Duration, and Wake Interval. During an advertised R-TWT SP, an STA that is not a member of the specific R-TWT schedule must defer a transmission if the latter cannot be completed before the SP begins.

III. DESIGN CHALLENGES

This Section outlines the key design challenges for admission control and scheduling of real-time uplink traffic in IEEE 802.11be networks.

A. Admission control and scheduling

When a new admission request for a real-time uplink flow arrives, the AP must allocate R-TWT resources on an MLO link without violating the requirements of the already admitted flows. The AP can perform either global or incremental scheduling. In global scheduling, the AP recalculates the entire schedule, potentially shifting the SPs of all active flows or regrouping them to share the same SP (see Sect. III-B), to maximize bandwidth efficiency and reduce resource fragmentation. However, this process can be time-consuming and requires broadcasting the updated R-TWT elements in the following Beacon frames. Conversely, with incremental scheduling, the AP tries to accommodate the new flow in the existing temporal gaps between scheduled SPs. If this is not feasible, the AP may combine the new flow with an existing flow with a compatible period and deadline to share the same SP. In particular, the SP may be shared using either Multi-User Orthogonal Frequency Division Multiple Access (OFDMA) or Time Division Multiple Access (TDMA). This work focuses on the TDMA approach, as discussed in Section IV. Flows that cannot be scheduled in compliance with their requirements by either the global or the incremental scheduling procedure are rejected. Finally, a hybrid strategy can also be adopted to maximize the flow admission rate. If the incremental scheduling fails, the AP switches to the global scheduling. The schedule is updated if a feasible global solution is found, otherwise the request is rejected.

B. Grouping

To minimize protocol overhead and improve the channel utilization, a key design choice in the R-TWT allocation process is how to schedule the admitted traffic, i.e., whether to allocate a dedicated SP per flow or to aggregate multiple flows within a shared R-TWT SP.

A dedicated R-TWT SP provides full temporal isolation. However, since each R-TWT schedule must be announced in a dedicated R-TWT element of the beacon frame, a high number of concurrent R-TWT schedules increases the beacon overhead. Moreover, as the SP duration is expressed in discrete units, each dedicated SP must be rounded up to the next multiple of the wake duration unit, thus causing bandwidth waste.

A shared R-TWT SP significantly reduces the overhead due to the Beacon by advertising a single R-TWT element per group of flows rather than broadcasting a dedicated element for each flow. Moreover, using shared R-TWT SPs maximizes bandwidth efficiency by applying the wake duration rounding only once for the entire group rather than for each flow. However, this approach requires more sophisticated scheduling logic at the AP, as it must also compute the shared SP duration and ensure that the expanded SP never violates the deadline for the flows belonging to the same group.

C. Beacon handling

The beacon is a periodic management frame transmitted by the AP to support multiple functions, including network synchronization, advertisement of R-TWT SPs, and coordination of multi-link configurations among associated MLDs. As beacon transmissions are assigned the highest priority among

all flow types, including the real-time flows [1], a violation of the flow requirements may occur. To mitigate this issue, the scheduler must consider the beacon as an independent flow and reserve a time interval for it. The time interval starts from the time instant at which the AP is expected to transmit the beacon and spans the length of the Beacon duration plus a guard interval to account for the delays potentially imposed by the medium access contention.

IV. AN OVERVIEW OF THE PROPOSED STRATEGIES

In the IEEE 802.11be network here considered, each flow f is characterized by a specific set of requirements (T, D, L) denoting its period, deadline, and payload size. To schedule incoming real-time flows, the AP assigns each flow to a specific R-TWT schedule on one of the available MLO links, taking into account the flow's timing constraints. Two admission control and R-TWT scheduling algorithms are proposed, as follows.




Fig. 2: Scheduling example of the proposed strategies.

1) *Dedicated R-TWT*: This approach represents a baseline solution in which a separate R-TWT schedule is assigned to each admitted flow, thus providing full temporal isolation. As shown in Fig. 2, this strategy suffers from two distinct forms of bandwidth waste, i.e., (i) an *offset waste*, as the AP is forced to align each new R-TWT SP to a multiple of the TU (i.e., 1024 μ s), thus preventing the allocation of a new flow between two consecutive R-TWT, and (ii) an *SP duration rounding waste*, as the duration of each R-TWT SP is rounded up to the next multiple of the wake duration unit.

Algorithm 1 outlines the process of admission control and scheduling of a dedicated R-TWT for an incoming real-time flow f on the available links $\mathcal{L}$. First, the AP calculates the airtime (t_{raw}) required for the uplink transmission, taking into account the payload size (L), the Modulation and Coding Scheme (MCS), and the MAC/PHY overhead associated with the trigger-based exchange (Alg. 1 Line 2). Next, the Wake Duration (WD) is calculated as a multiple of the wake duration unit, i.e., 256 μ s (Alg. 1 Line 3). To meet the deadline (D), the AP computes the maximum Wake Interval WI_{max} . This bound is defined as $D - t_{raw}$, thus ensuring that, even in the worst-case scenario (i.e., when a packet is generated immediately after its designated SP has started), the delay does not exceed the deadline. In addition to meeting the requirements of the new flow, the admission control assigns the Wake Interval (WI) ($\leq WI_{max}$) in a way that maximizes the Greatest Common Divisor (GCD) with both the Beacon Interval (BI_l) and the Wake Intervals of all active R-TWT schedules (S_l) on that link. In fact, under worst-case conditions, to avoid collisions the total airtime of all periodic flows (including the beacon) must fit within a single GCD interval. Consequently, a larger

GCD provides a wider collision-free time window that allows additional flows to be accommodated within a single GCD.

Algorithm 1 Dedicated R-TWT Admission Control

Require: Flow $f(T, D, L)$, Links $\mathcal{L}$
Ensure: Admission decision, Dedicated R-TWT schedule $\langle TWT, WI, WD \rangle$ on link l

```

1: for each link  $l \in \mathcal{L}$  do
2:    $t_{raw} \leftarrow \text{airtime}(L, MCS)$   $\triangleright$  Incl. all overhead
3:    $WD \leftarrow \lceil t_{raw}/256 \rceil \times 256 \mu s$   $\triangleright$  TWT duration unit
4:    $WI_{max} \leftarrow D - t_{raw}$   $\triangleright$  Maximum safe WI
5:    $S_l \leftarrow$  R-TWT schedules on link  $l$ 
6:    $WI \leftarrow \arg \max_{w \leq WI_{max}} \text{GCD}(w, BI_l, \{\tau.WI \mid \tau \in S_l\})$ 
7:   Search free  $TWT$  on link  $l$  avoiding Beacon and  $S_l$ 
8:   if valid  $TWT$  found then
9:      $S_l.add(\langle TWT, WI, WD \rangle)$   $\triangleright$  Create R-TWT schedule
10:    return SUCCESS
11:   end if
12: end for
13: return REJECT

```

To maximize the GCD, the proposed algorithm applies a deadline-aware oversampling. This step provides sufficient time span to assign distinct, non-overlapping offsets. Finally, the AP searches for a free temporal offset TWT using a 1024 μs resolution (see Sect. II). A valid offset must perfectly accommodate the dedicated R-TWT without colliding with the Beacon window or any pre-existing R-TWT SP. If a free offset is found, the flow is successfully admitted, the new schedule is added to S_l , and the algorithm returns a success status. Conversely, the algorithm proceeds to evaluate the next available link. If all links lack sufficient bandwidth, the admission request is rejected.

2) *Shared R-TWT*: This approach aims to mitigate the dedicated R-TWT limits by aggregating into a shared R-TWT SP the flows with periods and deadlines that satisfy specific requirements, e.g., the ones with periods multiple of each other. As shown in Fig. 2, this approach exploits the idle gaps between SPs and minimizes the SP duration rounding waste. If no compatible R-TWT schedules are available for grouping, the strategy falls back to the Dedicated R-TWT strategy.

Algorithm 2 attempts to merge the incoming flow f into an existing R-TWT schedule, for the sake of bandwidth efficiency. The procedure iterates on all available links $\mathcal{L}$ evaluating the pre-existing schedules S_l . For each candidate schedule τ , the AP first verifies the *isCompatible* function. This function ensures that the new flow and the existing R-TWT schedule may be grouped. If an R-TWT schedule is found, the AP calculates the airtime t_{raw} for the new flow and computes WD. For the sake of simplicity, in this work a TDMA approach is adopted and the AP sequentially polls the aggregated flows within the shared SP using dedicated trigger frames. The WD is calculated as the sum of the previously allocated time ($\tau.t_{raw}$) and the airtime of the new flow (t_{raw}). The 256 μs rounding is then applied once to this sum.

Since accommodating multiple flows increases the duration of the shared SP, the admission control must verify that the deadlines of both the existing and new flows are never violated. This is mathematically enforced by checking that the R-TWT schedule Wake Interval ($\tau.WI$) plus the new

Algorithm 2 Shared R-TWT Admission Control

Require: Flow $f(T, D, L)$, Links $\mathcal{L}$
Ensure: Admission decision, Updated shared R-TWT or Dedicated R-TWT schedule $\langle TWT, WI, WD \rangle$ on link l

```

1: for each link  $l \in \mathcal{L}$  do
2:    $S_l \leftarrow$  R-TWT schedules on link  $l$ 
3:   for each R-TWT schedule  $\tau \in S_l$  do
4:     if isCompatible( $\tau, f$ ) then
5:        $t_{raw} \leftarrow \text{airtime}(L, MCS)$ 
6:        $WD \leftarrow \lceil (\tau.t_{raw} + t_{raw})/256 \rceil \times 256 \mu s$ 
7:       if  $\tau.WI + WD \leq \min_{i \in \tau.F \cup \{f\}}(D_i)$  then
8:         if  $WD \leq \text{gap}(\tau, S_l \setminus \{\tau\})$  then
9:            $\tau.WD \leftarrow WD$ ;  $\tau.F.add(f)$ 
10:          return SUCCESS
11:        end if
12:      end if
13:    end for
14:  end for
15: end for
16: // Fallback to dedicated R-TWT strategy
17: return Dedicated R-TWT(f)

```

$\triangleright$ Alg. 1

WD does not exceed the most stringent application deadline among all flows mapped to that schedule, i.e., $\tau.WI + WD \leq \min_{i \in \tau.F \cup \{f\}}(D_i)$, where $\tau.F$ denotes the set of flows already assigned to the R-TWT schedule τ .

Once the deadline constraint is satisfied, the admission control must ensure that the SP does not overlap with adjacent R-TWT schedules. As the R-TWT start times must be aligned to integer multiples of 1024 μs , a baseline collision-avoidance check consists in strictly limiting the overall SP duration to less than 1024 μs . This approach is computationally very simple, but it limits the aggregation capacity. To overcome this limitation and increase the number of grouped flows, a more sophisticated strategy evaluates the actual idle time on the link. In particular, the scheduler dynamically checks that the WD does not exceed the idle time before the next scheduled event (e.g., another R-TWT SP or a Beacon), denoted as $\text{gap}(\tau, S_l \setminus \{\tau\})$. If the SP satisfies this constraint, the flow is successfully admitted, the schedule τ is updated with the new flow parameters, and the algorithm returns a success status. Conversely, if no R-TWT schedule can accommodate the flow on any MLO link (e.g., due to deadline violations), the algorithm falls back to the baseline strategy.

V. PRELIMINARY ASSESSMENT

To assess the proposed strategies and their impact on the flow admission rate and bandwidth efficiency, a preliminary analytical assessment was performed. In particular, a mix of high-frequency control loops (Robots) and lower-frequency telemetry flows (Vehicles), inspired by [4] and shown in Table I, was considered.

TABLE I: The considered flows [4]

Flow type	Period	Payload	Deadline
Robot	8 ms	50 B	8 ms
Vehicle	100 ms	100 B	20 ms

The traffic is modeled as periodic UDP flows. Three scenario were considered, i.e., Scenario L (Low traffic load) with 5 Robots and 5 Vehicles flows, Scenario M (Medium traffic

TABLE II: Admitted Flows and R-TWT Schedules: Dedicated vs. Shared R-TWT

Scenario	R-TWT Scheduling	Admitted Flows			R-TWT Schedules		
		Median	Min	Max	Median	Min	Max
L (10 flows)	Dedicated	10	10	10	10	10	10
	Shared	10	10	10	2	2	2
M (20 flows)	Dedicated	12	12	17	12	12	17
	Shared	20	20	20	2	2	2
H (50 flows)	Dedicated	12	12	17	12	12	17
	Shared	50	50	50	6	6	6

TABLE III: Link Utilization: Dedicated vs. Shared R-TWT (Mean $\pm$ Std. Dev over 100 runs)

Scenario	R-TWT Scheduling	Per-Link Utilization U_l [%]			Avg. Util. $\bar{U}$ [%]
		2.4 GHz	5 GHz	6 GHz	
L (10 flows)	Dedicated	13.9 $\pm$ 2.4	14.2 $\pm$ 2.3	7.4 $\pm$ 1.8	11.8 $\pm$ 0.4
	Shared	7.2 $\pm$ 0.4	0.0 $\pm$ 0.0	0.0 $\pm$ 0.0	2.4 $\pm$ 0.1
M (20 flows)	Dedicated	14.9 $\pm$ 3.0	15.0 $\pm$ 3.0	14.7 $\pm$ 2.8	14.9 $\pm$ 2.0
	Shared	7.0 $\pm$ 0.4	0.0 $\pm$ 0.0	0.0 $\pm$ 0.0	2.3 $\pm$ 0.1
H (50 flows)	Dedicated	14.8 $\pm$ 3.2	14.8 $\pm$ 3.0	14.9 $\pm$ 2.9	14.8 $\pm$ 1.9
	Shared	13.8 $\pm$ 0.9	7.1 $\pm$ 0.4	0.0 $\pm$ 0.0	7.0 $\pm$ 0.3

load) with 10 Robots and 10 Vehicle flows, and Scenario H (High traffic load) 25 Robots and 25 Vehicles flows. Regarding the network configuration, an MLO architecture spanning the $\mathcal{L} = \{2.4, 5, 6\}$ GHz frequency bands is considered, with channel bandwidths of 20, 40, and 160 MHz, respectively. Moreover, the physical layer is configured with a Guard Interval of 0.8 μ s, while the MCS index is set to 0 for control frames and 7 for data frames [3]. The metrics considered for the assessment are the network admission rate, the MAC-layer management overhead (measured as the number of established R-TWT schedules), and the bandwidth utilization in the MLO.

To quantify the bandwidth utilization, here the per-link channel utilization (U_l) is defined as the sum of the duty cycles of all R-TWT schedules active on a given link $l \in \mathcal{L}$, i.e., $U_l = \sum_{\tau \in \mathcal{S}_l} \frac{WD_\tau}{WI_\tau} \times 100$, where $\mathcal{S}_l$ represents the set of R-TWT schedules on the link l , while WD_τ and WI_τ are the reserved SP duration and wake interval of schedule τ , respectively. Consequently, the overall network utilization ($\bar{U}$) is calculated as the arithmetic mean of the utilization over all available links, i.e., $\bar{U} = \frac{1}{|\mathcal{L}|} \sum_{l \in \mathcal{L}} U_l$.

To ensure statistical significance, 100 runs were performed for each scenario. In each run, the order of the input flow requests was randomized to assess the robustness of the admission control logic with respect to the arrival sequence. Tables II and III summarize the results of the considered performance metrics for all evaluated scenarios.

In Scenario L, both algorithms successfully schedule 100% of the traffic flows. Conversely, in the other two scenarios, the limitations of the Dedicated R-TWT strategy become evident. For instance, in Scenario H, the Dedicated strategy only admits a median of 12 flows (this value was obtained in 89% of the runs) out of 50, while the Shared R-TWT strategy effectively supports the highest considered traffic load, achieving a 100% admission rate even with 50 flows. The high number of times the admission control rejects a flow

when the Dedicated strategy is adopted is due to the schedule fragmentation. In fact, this strategy assigns a separate offset for each flow, quickly exhausting all available offsets. As a result, new flows are rejected when the utilization of the three bands is around 15%. Instead, the Shared R-TWT strategy prevents this fragmentation and improves bandwidth utilization. For example, in Scenario H, this strategy accommodates all 50 flows with an average network utilization of only 7.0%. In particular, all real-time flows were scheduled on the 2.4 and 5 GHz bands. Consequently, the 6 GHz band remains unused, thereby being fully available to accommodate additional real-time flows or best-effort traffic. Finally, grouping significantly reduces the MAC-layer management overhead. The Dedicated strategy requires one schedule per flow, thus reaching saturation at 12 schedules (in 89% of the runs) due to the limited number of available offsets. In contrast, the Shared strategy consolidates the 50 flows of Scenario H into six schedules. Consequently, this approach minimizes the number of R-TWT elements and reduces the management payload transmitted in Beacon frames.

VI. CONCLUSIONS AND FUTURE WORK

This work investigated admission control and scheduling strategies for real-time periodic uplink flows in IEEE 802.11be networks. Two MLO-aware R-TWT allocation strategies, namely the Dedicated and Shared R-TWT, were evaluated through an analytical assessment. The Dedicated strategy leads to rapid network saturation due to bandwidth inefficiencies, while flow grouping in Shared R-TWT effectively mitigates the problem. The Shared approach also significantly reduces management overhead, improves bandwidth utilization and increases the network admission rate under high traffic loads. Ongoing work is addressing extensive performance evaluations of the Shared R-TWT strategy through simulations and experiments on a real testbed leveraging Software-Defined Radio (SDR) platforms. Further work will investigate a way to deal with aperiodic real-time flows.

REFERENCES

- [1] "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 802.11be-2024*, pp. 1–1020, 2025.
- [2] A. Belogaev, X. Shen, C. Pan, X. Jiang, C. Blondia, and J. Famaey, "Dedicated Restricted Target Wake Time for Real-Time Applications in Wi-Fi 7," in *2024 IEEE Wireless Communications and Networking Conference (WCNC)*. IEEE, 2024, pp. 1–6.
- [3] E. Mozaffariharar, F. Wilhelmi, L. Galati-Giordano, P. Imputato, M. Menth, and S. Avallone, "R-TWT in Wi-Fi 7 and Beyond: Enabling Bounded Latency, Energy Efficiency, and Reliability," in *2025 IEEE 30th International Conference on Emerging Technologies and Factory Automation (ETFA)*, 2025, pp. 1–8.
- [4] C. Barroso-Fernández, J. Martín-Pérez, C. Ayimba, and A. De La Oliva, "Time-Sensitive IIoT Flows over Wi-Fi: a Network Calculus Approach," *IEEE Internet of Things Journal*, 2025.
- [5] D. V. Bankov, A. I. Lyakhov, E. A. Stepanova, and E. M. Khorov, "Performance evaluation of Wi-Fi 7 networks with restricted target wake time," *Problems of Information Transmission*, vol. 60, no. 3, pp. 233–254, 2024.
- [6] D. Ergenç, M. Zakrzewski, and F. Dressler, "Multi-Link Scheduling with Restricted Target Wake Time in Wi-Fi 7," in *21th IEEE/IFIP Wireless On-demand Network systems and Services Conference (WONS 2026)*, 2026.