Adaptive Scheduling for Low-Latency Coordination in Wi-Fi 8 Multi-AP Networks

Hiya Shah
School of Engineering and Applied Science, Ahmedabad University, Gujarat, India
hiya.s5@ahduni.edu.in

***Abstract*—Reliable low-latency communication is a critical requirement in enterprise wireless networks such as hospitals, offices, and campuses. This paper proposes an earliest deadline first (EDF)-Lyapunov-Robbins-Monro (ELR), a stochastic scheduling algorithm for IEEE 802.11bn (Wi-Fi 8) Multi-Access Point Coordination Coordinated-Spatial Reuse (MAPC C-SR) networks that jointly accounts for queue stability and deadline-aware latency regulation under bursty traffic. A Lyapunov drift-based criterion for a group is adopted to ensure queues remain stable under varying traffic loads. Since the optimal balance between queue backlog and deadline urgency cannot be determined a priori under bursty traffic, EDF term is incorporated into the selection metric with a tunable balance parameter α, governed by Robbins-Monro stochastic approximation scheme. The proposed algorithm addresses the inability of existing schedulers to track sudden congestion under bursty traffic, by dynamically adjusting α to suppress sharp delay spikes. Simulations over a four-access point (AP) enterprise deployment under bursty Markov-Modulated Poisson Process (MMPP) traffic demonstrate that ELR achieves 14.23%, 13.26%, and 7.97% reduction in 99th percentile delay over maximum number of packets (MNP), oldest packet (OP), and traffic alignment tracker (TAT) respectively under high load with 16 stations (STAs).**

***Index Terms*—IEEE 802.11bn, Coordinated Spatial Reuse, Multi-Access Point Coordination, Stochastic Approximation**

I. Introduction

With the continuous evolution of wireless local area networks (WLANs), the demand for ultra-low latency, and reliable connectivity across dense environments has grown substantially [1]. In wireless networks, multiple stations contend for channel access using exponential random backoff. Multiple simultaneous transmissions cause collisions leading to retransmissions and increased energy consumption per bit. This energy efficiency problem has been addressed using contention-adaptive mechanisms [2] through Enhanced Distributed Channel Access (EDCA) WLANs.

As the network density increases, collision, and backoff overheads increase. This limits the achievable gain through a single-AP deployment. To improve the packet transmission rates, coordination among APs is addressed through MAPC framework standardized by 802.11bn Task Group (TGbn). This enables APs to share received signal strength indicator (RSSI) measurements and transmit power constraints to coordinate transmissions through C-SR. Within this framework, Chemrov [3] has proposed Parameterized Spatial Reuse (PSR) which enables simultaneous transmissions by reducing transmission power below controlled interference thresholds for real-time applications (RTAs). A throughput analysis in coordinated multi-AP WLANs was carried out in [4]. The study considered multiple deployments with varying AP-AP distances and established significant spatial reuse gains, thereby motivating the need for schedulers that can exploit these gains while adhering to the latency constraints. Subsequently, authors have proposed algorithms within the MAPC framework to achieve specific goals such as maximizing number of packets delivered, scheduling the oldest packet, and a combination of both of these [5]. However, these algorithms are heuristic and are not based on proofs of convergence or queue stability.

To address limited adaptivity in goal-specific schedulers, multi-armed bandit (MAB) algorithms [6], [7] enable adaptive media access control (MAC) layer scheduling. This is achieved by dynamically selecting transmission groups to maximize aggregated throughput. However, MAB-based approaches fail to incorporate delay constraints or queue stability guarantees due to throughput-only optimization. This limits their applicability in latency-sensitive deployments. Stochastic network utility maximization and renewal optimization frameworks have been proposed to provide analytical guarantees for online scheduling, extending the results from classical Bianchi model [8] under saturated traffic conditions. These frameworks are however, analytical models of fixed MAC protocols and cannot be completely extended to multi-AP coordination setting.

To address the need for convergent online parameters, a variant of Robbins-Monro incorporates prior information to find the convergence rate in optimization tasks [9]. Scheduling algorithms with analytical convergence guarantee [10], and ergodic convergence under convex settings [11] form a theoretical basis for the convergence of online parameters. However, application of Robbins-Monro based parameter adaptation within a Lyapunov-stable MAPC C-SR scheduler remains uninvestigated. While MAPC C-SR improves spectral efficiency, simultaneous transmissions introduce inter-basic service set (BSS) interference, causing queue buildup and degraded latency.

Existing schedulers lack the ability to respond to such sudden traffic bursts, resulting in unstable delay behavior. ELR addresses this by dynamic α adaptation in scheduling framework. To the best of the authors' knowledge, no prior work has jointly considered queue stability, convergent adaptive parameter tuning and deadline-aware scheduling within a unified MAPC C-SR framework. The main contributions of our work are as follows:

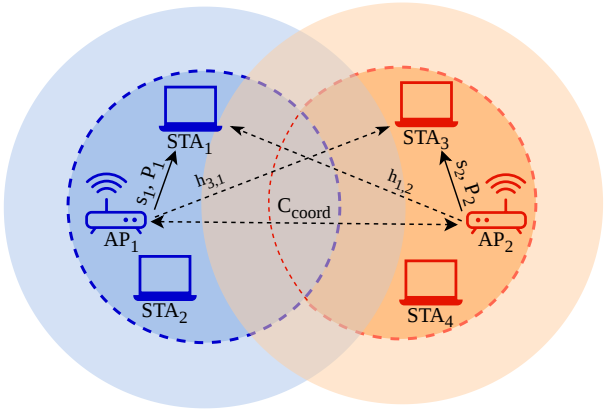

Fig. 1. Overlapping BSS interference model for IEEE 802.11bn C-SR. Outer circle denotes the BSS coverage, while the inner dashed circle represents C-SR feasibility region.

1) We develop a simulation-based model for dense Wi-Fi 8 multi-AP networks incorporating C-SR to evaluate performance over multiple realizations with randomly placed STAs and bursty traffic dynamics.
2) We propose an adaptive downlink scheduling algorithm considering queue backlog-aware group selection with a Robbins-Monro stochastic approximation mechanism to regulate delay around a desired target. Lyapunov drift analysis is provided to establish queue stability under ergodicity and bounded second moment assumptions.

The remainder of the paper is structured as follows: Section II introduces the system model and describes information transmission across different parameters under consideration. Section III discusses the Lyapunov drift-based Robbins-Monro adaptive scheduling algorithm mathematically. The simulation results are given in Section IV. Section V concludes the paper.

II. System Model

A. System Description

We consider a coordinated multi-AP network consisting of N APs at positions $\mathbf{x}_i \in \mathbb{R}^2$, $i \in \{1, \ldots, N\}$, and K STAs at positions $\mathbf{y}_j \in \mathbb{R}^2$, $j \in \{1, \ldots, K\}$. Fig. 1 illustrates AP_1-STA_1 and AP_2-STA_3 as simultaneous transmissions s_1, and s_2 respectively, where $h_{j,k}$ denotes the channel coefficient from AP k to STA j, and P_1, P_2 denote the corresponding transmit powers. Inter-BSS interference arises via cross-links $h_{1,2}$ and $h_{3,1}$. Inter-AP coordination is achieved through Multi-Access Point Coordination Intention-to-Coordinate Frame (MAPC-ICF) as shown in Fig. 2. The coordinating AP selects STA group for transmission opportunity (TXOP) allocation based on transmission parameters (modulation and coding scheme (MCS) and channel bandwidth). The scenario represents an enterprise network of four BSSs, each with one AP at the center and two associated STAs. MAC protocol data units are aggregated into A-MPDUs to reduce per-frame overhead, with failed subframes retransmitted in subsequent TXOPs through the block acknowledgment mechanism. Packet arrivals follow a continuous-time MMPP, while scheduling decisions are made at discrete TXOP epoch boundaries.

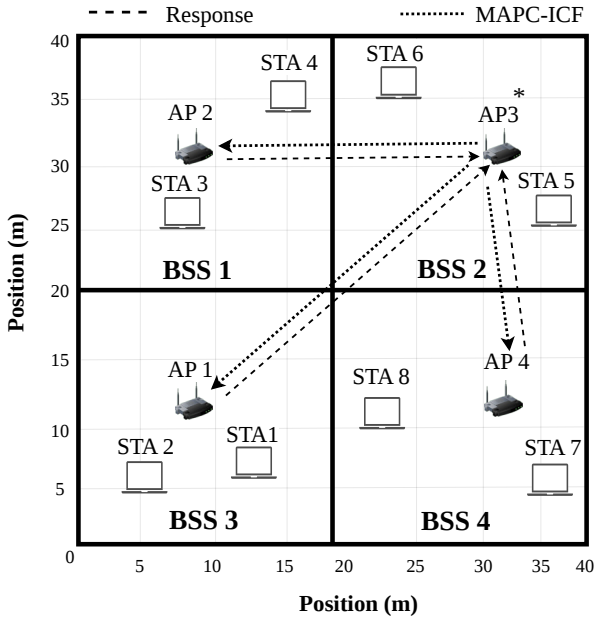

Fig. 2. A four BSS enterprise scenario, each BSS occupying 20 m×20 m quadrant. AP_3 acts as a coordinating AP and is denoted by (*).

B. Channel Model

To quantify the interference experienced at each STA during simultaneous transmissions, we model channel gain between every AP-STA pair using the IEEE 802.11ax TGax enterprise indoor pathloss model. AP i transmits per-spatial-stream ℓ with power $P_{i,\ell}$, $\ell \in \{1, \ldots, N_{\text{ss}}\}$, subject to an Equivalent Isotropic Radiated Power (EIRP) constraint, $P_{\text{EIRP}}^{\max}$, and spectral density limits. The Euclidean separation between AP i and STA j is $d_{j,i} = \|\mathbf{x}_i - \mathbf{y}_j\|_2 > 0$. The channel coefficient is modeled at $f_c = 6$ GHz as $h_{j,i} = 10^{-PL(d_{j,i})/10}$, where

$$PL(d_{j,i}) = \begin{cases} PL_0 + 20\log_{10}(d_{j,i}) \\ \quad + 7n_w + X_\sigma, & d_{j,i} < d_{\text{BP}}, \\ PL_0 + 20\log_{10}(d_{\text{BP}}) \\ \quad + 35\log_{10}\left(\dfrac{d_{j,i}}{d_{\text{BP}}}\right) & d_{j,i} \geq d_{\text{BP}} \\ \quad + 7n_w + X_\sigma, \end{cases} \tag{1}$$

with $PL_0 = 40.05 + 20\log_{10}(f_c/2.4)$ dB the frequency-dependent reference loss, break point $d_{\text{BP}} = 10$ m, wall penetration loss $7n_w$ dB (n_w walls), and log-normal shadowing $X_\sigma \sim \mathcal{N}(0, \sigma_{\text{sh}}^2)$. Interference at STA j arises from all concurrently active APs $k \neq i$, each contributing through the channel coefficient $h_{j,k}$, representing the path from AP k to the receiver STA j.

The received instantaneous signal-to-noise ratio (SINR) for STA j served by AP i on stream ℓ at epoch t is

$$\gamma_{j,i,\ell}(t) = \frac{P_{i,\ell}(t)\,|h_{j,i}(t)|^2}{\sigma^2 + \sum_{k \neq i} P_{k,\ell}(t)\,|h_{j,k}(t)|^2}, \tag{2}$$

where numerator is the desired signal power from AP i and the denominator accounts for thermal noise σ^2 and

aggregate co-channel interference. The achievable data rate $R_j(t) = f(\gamma_{j,i,\ell}(t))$ is determined by the MCS selected from instantaneous SINR, where $f(\cdot)$ is the step-wise rate function defined by the IEEE 802.11bn MCS table; a higher SINR allows selection of a higher MCS index, yielding a greater number of successfully decoded bits per transmission.

Let $\gamma_{\min}$ denote the minimum SINR threshold of the lowest supported MCS index. Packet error rate (PER) for STA j at epoch t is

$$\mathrm{PER}_j(t) = \begin{cases} 1 & \gamma_{j,i,\ell}(t) < \gamma_{\min}, \\ 10^{-2} & \gamma_{j,i,\ell}(t) \geq \gamma_{\min}. \end{cases} \tag{3}$$

When $\gamma_{j,i,\ell}(t) \geq \gamma_{\min}$, each packet is successfully decoded independently with probability $1 - \mathrm{PER}_j(t)$, so the number of successfully decoded packets S at epoch t follows $S_j(t) \sim \mathrm{Binomial}(M_j(t), 1 - \mathrm{PER}_j(t))$, where $M_j(t) = \min\{Q_j(t), A_j^{\max}(t)\}$, $Q_j(t)$ is the length of queue at epoch t and $A_j^{\max}(t)$ is the maximum A-MPDU aggregation size determined by the selected MCS and TXOP duration. For the stability analysis in Section III, we restrict attention to the regime $\gamma_{j,i,\ell}(t) \geq \gamma_{\min}$, under which $P_e = \mathrm{PER}_j(t) \leq 10^{-2}$ uniformly across all STAs and epochs.

III. Proposed Scheme

The proposed EDF-Lyapunov-Robbins-Monro (ELR) framework integrates three mechanisms: a Lyapunov drift-based group selection for provable queue stability, an EDF urgency term for per-packet delay enforcement, and a Robbins-Monro loop that adaptively balances both objectives.

A. Traffic Model and Queue Dynamics

Each STA j generates traffic via a two-state MMPP with state space $\{\mathrm{ON}, \mathrm{OFF}\}$ and exponentially distributed durations $\mathbb{E}[T_{\mathrm{ON}}]$, $\mathbb{E}[T_{\mathrm{OFF}}]$, with long-run average arrival rate

$$\bar{\lambda}_j = \lambda_j \frac{\mathbb{E}[T_{\mathrm{ON}}]}{\mathbb{E}[T_{\mathrm{ON}}] + \mathbb{E}[T_{\mathrm{OFF}}]} \in (0, \lambda_j). \tag{4}$$

Let $A_j(t)$ and $S_j(t)$ denote arriving and decoded packet counts for STA j at epoch t. Under a work-conserving discipline, the queue length evolves as (Algorithm 1, line 12):

$$Q_j(t+1) = \max\{Q_j(t) - S_j(t),\ 0\} + A_j(t). \tag{5}$$

Since parameters are finite, $A_j^{\max}(t) \leq A_{\mathrm{phy}}^{\max} < \infty$ a.s. Each packet fails independently with probability P_e, giving $S_j(t) \mid Q_j(t) \sim \mathrm{Binomial}(M_j(t), 1 - P_e)$, with conditional mean

$$\mathbb{E}[S_j(t) \mid Q_j(t)] = \min\{Q_j(t),\ A_j^{\max}(t)\}\,(1 - P_e) \triangleq \mu_j^{\mathrm{inst}}(t), \tag{6}$$

where $\mu_j^{\mathrm{inst}}(t)$ is the instantaneous expected service rate, approximated by the exponential moving average (EMA) estimate $\hat{\mu}_j(t)$ in (7). Since $A_j^{\max}(t)$ varies with channel conditions, an EMA with smoothing factor $\beta \in (0, 1)$ is updated only when STA j is scheduled (Algorithm 1, lines 14-16), estimating the time-varying service rate:

$$\hat{\mu}_j(t+1) = \begin{cases} (1-\beta)\,\hat{\mu}_j(t) + \beta\,\frac{S_j(t)}{T_{\mathrm{tx}}} & \text{if scheduled at } t, \\ \hat{\mu}_j(t) & \text{otherwise,} \end{cases} \tag{7}$$

where $T_{\mathrm{tx}} = 5$ ms is the TXOP duration and $S_j(t)/T_{\mathrm{tx}}$ is the instantaneous observed throughput.

B. Group Scheduling Metric

The coordinating AP selects $g^*(t) \in \mathcal{G}$, where $\mathcal{G}$ is the set of all candidate STA subsets. The set $\mathcal{G}$ is formed by exhaustive search over only those combinations where STAs gain from coordinated (rather than individual) transmission, evaluated from semi-static RSSI tables exchanged over the Distribution System. Under low mobility, grouping is updated periodically and not per TXOP. To schedule among these candidates, the aggregate backlog is $Q_g(t) = \sum_{j \in g} Q_j(t)$ with mean estimated service rate $\hat{\mu}_g(t) = \frac{1}{|g|} \sum_{j \in g} \hat{\mu}_j(t)$. Let $T_{j,\mathrm{first}}(t)$ denote the arrival epoch of the oldest undelivered packet at STA j. The group head of line (HOL) delay and deadline slack are computed as:

$$D_g(t) = \max_{j \in g} \left(t - T_{j,\mathrm{first}}(t)\right), \tag{8}$$

$$\Delta_g(t) = \max\{D_{\text{target}} - D_g(t),\ \epsilon\}, \tag{9}$$

where $D_{\text{target}} > 0$ is the prescribed delay target and $\epsilon > 0$ prevents division by zero. The urgency ratio $D_g(t)/\Delta_g(t)$ increases monotonically as $D_g(t) \to D_{\text{target}}$, and is bounded above by $D_g(t)/\epsilon$ owing to the floor $\epsilon > 0$ in (9), ensuring the scheduling metric remains finite at all times. The joint metric:

$$\Phi_g(t) = \Big(\sum_{j \in g} Q_j(t)\,\hat{\mu}_j(t)\Big)\left(1 + \alpha_t\,\frac{D_g(t)}{\Delta_g(t)}\right), \tag{10}$$

where $\alpha_t \in [0, 1]$ is adapted by Robbins-Monro (12); $\alpha_t = 0$ recovers the max-weight metric. The scheduled group is

$$g^*(t) = \arg\max_{g \in \mathcal{G}}\ \Phi_g(t). \tag{11}$$

C. Adaptive Tuning via Robbins-Monro

To obtain the optimal balance between queue backlog and deadline urgency, the projected Robbins-Monro scheme drives α_t toward α^*, at which $\bar{D}(t) = \frac{t}{t+1}\bar{D}(t-1) + \frac{1}{t+1}D_{g^*}(t)$ converges to D_{target}:

$$\alpha_{t+1} = \Pi_{[0,1]}\left(\alpha_t + \frac{\eta_0}{1+t}\big(\bar{D}(t) - D_{\mathrm{target}}\big)\right), \tag{12}$$

where $\Pi_{[0,1]}(x) = \min\{1, \max\{0, x\}\}$ projects the updated iteration onto $[0, 1]$, $\eta_0 > 0$ is the initial step size, and $\eta_t = \eta_0/(1+t)$ satisfies $\sum \eta_t = +\infty$ and $\sum \eta_t^2 = \eta_0^2 \pi^2/6 < +\infty$. When $\bar{D}(t) > D^{\text{target}}$, α_t increases to embed stronger deadline urgency into (10); when $\bar{D}(t) < D^{\text{target}}$, it decreases to relax the urgency weight and restore throughput-optimal group selection. The Robbins-Monro update drives α_t continuously toward the value that satisfies $\bar{D}(t) = D^{\text{target}}$.

D. Lyapunov Stability Analysis

Let $L(t) = \frac{1}{2}\sum_{j=1}^{K} Q_j(t)^2$ and $\Delta(t) = \mathbb{E}[L(t+1) - L(t) \mid \mathbf{Q}(t)]$. Since $\Delta_g(t) \geq \epsilon$ and $D_g(t) \leq D_{\max}$, the urgency factor in (10) obeys $1 \leq 1 + \alpha_t D_g/\Delta_g \leq 1 + \Gamma$ with $\Gamma \triangleq D_{\max}/\epsilon$.

Algorithm 1 ELR: Stochastic Lyapunov Scheduler

Require: $\beta \in (0,1)$, $\eta_0 > 0$, $\epsilon > 0$, $D_{\text{target}} > 0$

Init: $Q_j(0) \leftarrow 0$, $\hat{\mu}_j(0) > 0$, $\alpha_0 \in [0,1]$, $\bar{D}(0) \leftarrow 0$, $\forall j \in K$

for $t = 0, 1, 2, \ldots$ **do**

  for each $g \in \mathcal{G}$ **do**

    $Q_g(t) \leftarrow \sum_{j \in g} Q_j(t)$; $\hat{\mu}_g(t) \leftarrow \frac{1}{|g|} \sum_{j \in g} \hat{\mu}_j(t)$

    $D_g(t) \leftarrow \max_{j \in g}(t - T_{j,\text{first}}(t))$

    $\Delta_g(t) \leftarrow \max\{D_{\text{target}} - D_g(t), \epsilon\}$

    $\Phi_g(t) \leftarrow \Big(\sum_{j \in g} Q_j(t)\, \hat{\mu}_j(t)\Big)\Big(1 + \alpha_t \frac{D_g(t)}{\Delta_g(t)}\Big)$

  end for

  $g^*(t) \leftarrow \arg\max_{g \in \mathcal{G}} \Phi_g(t)$; $D^*(t) \leftarrow D_{g^*}(t)$

  Transmit g^*; observe $S_j(t)$, $\forall j \in g^*$

  for each $j \in K$ **do**

    $Q_j(t+1) \leftarrow \max\{Q_j(t) - S_j(t), 0\} + A_j(t)$

  end for

  for each $j \in g^*$ **do**

    $\hat{\mu}_j(t+1) \leftarrow (1-\beta)\hat{\mu}_j(t) + \beta S_j(t)/T_{\text{tx}}$

  end for

  $\bar{D}(t) \leftarrow \frac{t}{t+1}\bar{D}(t-1) + \frac{1}{t+1} D_{g^*}(t)$

  $\alpha_{t+1} \leftarrow \Pi_{[0,1]}\Big(\alpha_t + \frac{\eta_0}{1+t}(\bar{D}(t) - D_{\text{target}})\Big)$

end for

Applying $(\max\{a,0\})^2 \le a^2$ to (5), summing over j, and taking $\mathbb{E}[\cdot \mid \mathbf{Q}(t)]$ yields the conditional drift

$$\Delta(t) \le B + \sum_{j=1}^{K} Q_j(t)\bar{\lambda}_j - \sum_{j=1}^{K} Q_j(t)\,\mathbb{E}[S_j(t) \mid \mathbf{Q}(t)], \quad (13)$$

where the PHY bounds $A_j \le A_{\text{phy}}^{\max}$, $S_j \le S_{\text{phy}}^{\max}$ give $B = \frac{K}{2}[(A_{\text{phy}}^{\max})^2 + (S_{\text{phy}}^{\max})^2] < \infty$ and $\bar{\lambda}_j = \mathbb{E}[A_j(t)]$ by stationarity. Only $g^\star(t)$ is served, so by (6)-(7) the served weight is $\sum_j Q_j \mathbb{E}[S_j \mid \mathbf{Q}] = \sum_{j \in g^\star} Q_j \hat{\mu}_j \triangleq W_{g^\star}(t)$. With $g^{\text{MW}} = \arg\max_g W_g$ and the urgency factor in $[1, 1+\Gamma]$, maximizing Φ_g gives $W_{g^\star} \ge (1+\Gamma)^{-1} W_{g^{\text{MW}}}$. Since stability concerns large backlogs ($Q_j \ge A_j^{\max}$), the rate $\bar{\mu}_j \triangleq A_j^{\max}(1 - P_e)$ is queue-independent and $\sum_j Q_j \bar{\mu}_j$ is linear; max-weight optimality [12] gives $W_{g^{\text{MW}}} \ge \sum_j Q_j \bar{\mu}_j$ for interior arrivals. Substituting into (13) gives the drift bound

$$\Delta(t) \le B + \sum_{j=1}^{K} Q_j(t)\big(\bar{\lambda}_j - \tfrac{\bar{\mu}_j}{1+\Gamma}\big). \quad (14)$$

If $\bar{\lambda}_j < \bar{\mu}_j/(1+\Gamma)$ for all j, then $\Delta(t) \le B - \tilde{\varepsilon}_{\min} \sum_j Q_j(t)$ with $\tilde{\varepsilon}_{\min} = \min_j(\bar{\mu}_j/(1+\Gamma) - \bar{\lambda}_j) > 0$. Summing over $t = 1, \ldots, T$, telescoping, and using $\mathbb{E}[L(T+1)] \ge 0$ establishes strong stability:

$$\limsup_{T \to \infty} \frac{1}{T} \sum_{t=1}^{T} \sum_{j=1}^{K} \mathbb{E}[Q_j(t)] \le \frac{B}{\tilde{\varepsilon}_{\min}} < \infty. \quad (15)$$

The factor $(1+\Gamma)$ quantifies the stability-region reduction from the urgency term; $\Gamma \to 0$ recovers the max-weight margin.

IV. Simulation Results

The simulation study is carried out using MATLAB software to model an IEEE 802.11bn grid enterprise WLAN under bursty MMPP traffic across medium and high load traffic, with packets classified under the best-effort (BE) access category. The physical layer is configured with a clear channel assessment (CCA) threshold of -82 dBm, $N_{ss} = 2$ spatial streams, and a frame length of 12 kbits. All results are averaged over 5×10^5 scheduling epochs. Three scheduling schemes, namely MNP, OP, and TAT as considered in [5], and the proposed ELR are evaluated under two deployments where 4 APs have a fixed position, with 8 and 16 randomly placed STAs. Performance is evaluated using the 99th percentile delay, capturing tail latency behavior critical to real-time enterprise applications.

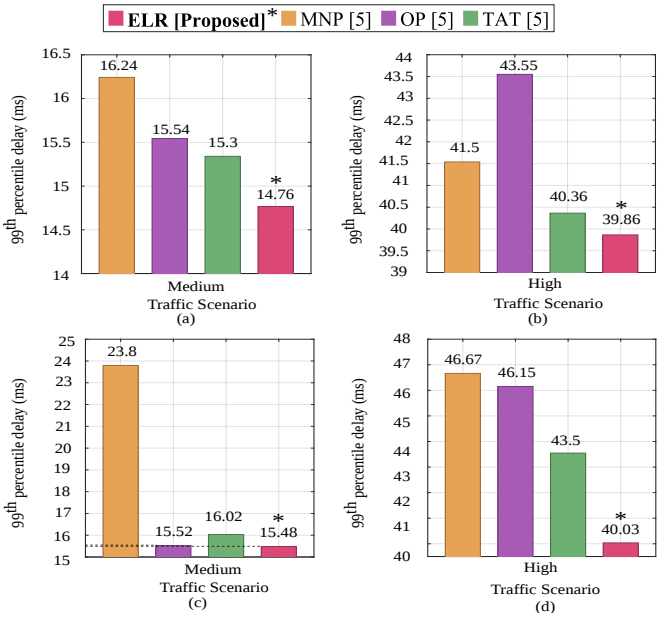

Fig. 3. Comparison of delay performance under varying bursty traffic scenarios (a) 8 STAs, Medium; (b) 8 STAs, High; (c) 16 STAs, Medium; (d) 16 STAs, High.

In Fig. 3a, ELR achieves the lowest tail delay of 14.76 ms. Under high traffic, ELR maintains the lowest tail delay at 39.86 ms against TAT (40.36 ms), MNP (41.5 ms), and OP (43.55 ms). On doubling the number of stations as shown in Fig. 3c, under medium traffic, ELR achieves a tail delay of 15.48 ms against OP (15.52 ms), TAT (16.02 ms), and MNP (23.8 ms), demonstrating converging tail performance. In case of ELR, queue backlogs are uniformly small across all groups, reducing the differentiation signal of the max-weight term $\sum_{j \in g} Q_j(t)\, \hat{\mu}_j(t)$. Under high traffic, where backlog pressure is large, ELR achieves a significant margin by reducing the 99th percentile delay by 14.23% over MNP, 13.26% over OP, and 7.97% over TAT. This aggregation motivates an analysis per-iteration, where the dynamic adaptation of ELR becomes observable.

Fig. 4a presents the 99th percentile delay for MNP, OP, TAT, and the proposed ELR over 25 iterations under 16 STAs bursty-medium traffic. MNP demonstrates several sharp transitions, while OP and TAT remain comparatively stable in the range of 12-18 ms. The adaptive rate control of ELR effectively balances delay performance where α updates at

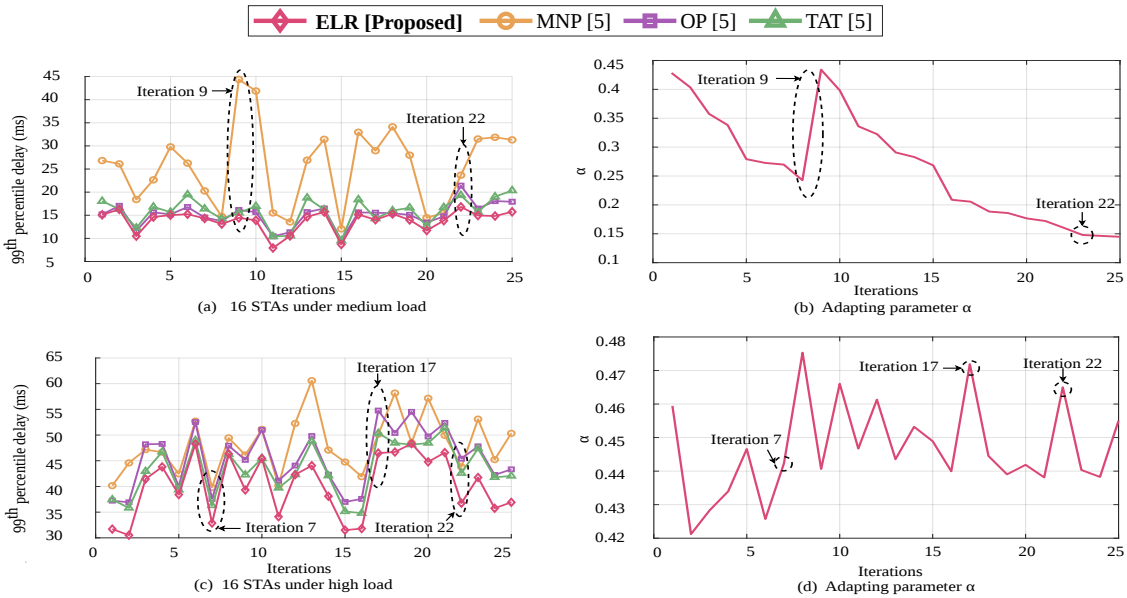

Fig. 4. 99th percentile queuing delay for MNP, OP, TAT, and ELR over 25 iterations for 16 STAs under (a) medium and (c) high loads and the evolution of adaptive parameter α in (b) and (d) respectively illustrating the dynamic congestion response unique to ELR.

iteration 9 as shown in Fig. 4b, and self-attenuates at iteration 22 once congestion is resolved. In Fig. 4c, under 16 STAs bursty-high traffic, the three mechanisms experience elevated delay across the range 34-62 ms. However, ELR consistently adapts to the bursts that occur at iterations 7 (transitioning to iteration 8), 17, and 22. Fig. 4d accounts for this behavior where α oscillates within a narrow elevated band ($\approx$ 0.42-0.48) unlike convergence in the case of medium load. Each upward update in α corresponds to a delay constraint in Fig. 4c. The continuous adaptation is event-specific, implying that under higher STA density, next set of congestion events arrive before ELR relaxes, enabling continuous online adaptation.

V. Conclusion

We present a novel adaptive scheduling framework for dense IEEE 802.11bn multi-AP networks under C-SR, jointly optimizing queue stability and deadline-aware packet delivery to minimize tail latency. An EDF urgency metric with Lyapunov-based adaptive slack weighting Robbins-Monro approximation address time-varying channels and bursty MMPP traffic. TGax simulations demonstrate reduction in the 99th percentile delay up to 14.23%, 13.26%, and 7.97% over MNP, OP, and TAT respectively confirming the effectiveness of adaptive scheduling for low-latency Wi-Fi 8 enterprise deployment.

Acknowledgment

The author sincerely thanks the anonymous reviewers for their valuable feedback and suggestions. The author expresses her heartfelt gratitude to Professor Dhaval K. Patel, MICxN Research Laboratory, Ahmedabad University, for his dedicated mentorship, continuous guidance, and unwavering support throughout this research journey, which were instrumental in achieving the outcomes presented in this work. The author also acknowledges the Stepwell High-Performance Computing Facility for providing the computational resources necessary for this study.

References

[1] S. He, C. Peng, W. Huang, Z. An, Y. Qian, and L. Liu, "Toward Wi-Fi 8 standard: A survey of state-of-the-art technologies," *IEEE Open Journal of the Communications Society*, vol. 6, pp. 10 150–10 170, 2025.

[2] J.-C. Chen and K.-W. Cheng, "EDCA/CA: Enhancement of IEEE 802.11e EDCA by contention adaption for energy efficiency," *IEEE Transactions on Wireless Communications*, vol. 7, no. 8, pp. 2866–2870, 2008.

[3] K. Chemrov, D. Bankov, A. Lyakhov, and E. Khorov, "A scheduler for real-time service in Wi-Fi 8 multi-AP networks with parameterized spatial reuse," *IEEE Communications Letters*, vol. 28, no. 7, pp. 1654–1657, 2024.

[4] D. Nunez, F. Wilhelmi, L. Galati-Giordano, G. Geraci, and B. Bellalta, "Spatial reuse in IEEE 802.11bn coordinated multi-AP WLANs: A throughput analysis," *arXiv:2407.16390*, 2024.

[5] D. Nunez, P. Imputato, S. Avallone, M. Smith, and B. Bellalta, "Enabling reliable latency in Wi-Fi 8 through multi-AP joint scheduling," *IEEE Open Journal of the Communications Society*, 2025.

[6] M. Wojnar, W. Ciezobka, K. Kosek-Szott, K. Rusek, S. Szott, D. Nunez, and B. Bellalta, "IEEE 802.11bn multi-AP coordinated spatial reuse with hierarchical multi-armed bandits," *IEEE Communications Letters*, 2024.

[7] F. Wilhelmi, B. Bellalta, S. Szott, K. Kosek-Szott, and S. Barrachina-Muñoz, "Coordinated multi-armed bandits for improved spatial reuse in Wi-Fi," *arXiv:2412.03076*, 2025.

[8] 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.

[9] S. Liu, K. Ma, and S. M. Goetz, "A Robbins–Monro sequence that can exploit prior information for faster convergence," *arXiv:2401.03206*, 2024.

[10] M. J. Neely, "Fast learning for renewal optimization in online task scheduling," *arXiv:2007.09532*, 2021.

[11] P. Bianchi, "Ergodic convergence of a stochastic proximal point algorithm," *arXiv:1504.05400*, 2016.

[12] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Trans. Autom. Control*, vol. 37, no. 12, pp. 1936–1948, 1992.