

Article

Memory-Based Predictive Resource Allocation for NR-V2X Traffic

Nurgüneş Yordanov *  and Bülent Çavuşoğlu 

Department of Electrical-Electronic Engineering, Atatürk University, 25240 Erzurum, Türkiye;
bcavusoglu@atauni.edu.tr

* Correspondence: nurgunesyordanov@gmail.com

Abstract

New Radio Vehicle-to-Everything (NR-V2X) safety scheduling is difficult because burst episodes increase urgent arrivals, lower transmission success, and create retransmissions that compete for future slots. A scheduler that waits for the visible queue can react late, whereas always reserving extra safety physical resource blocks (PRBs) consumes the best-effort (BE) capacity after the stress has passed. This study proposes Memory-Based Predictive Allocation (MPA), a finite-action PRB allocation rule for safety and BE coexistence. MPA combines the deadline queue and retry state with a decayed transient-deficit memory, online success calibration, and a recoverability-aware BE cost guard. At each slot, it tests feasible safety PRB increments and chooses the action that first limits urgent safety loss, then reduces next-slot carryover, and finally avoids unnecessary PRB use. The model uses an NR-V2X resource pool interpretation and a calibrated signal-to-interference-plus-noise-ratio (SINR)-to-success mapping with hybrid automatic repeat request (HARQ)-like combining. Monte Carlo results show that MPA lowers safety misses relative to queue-reactive scheduling while preserving more BE throughput than a maximum safety reservation. In dense non-line-of-sight (NLOS) stress, MPA keeps the 95th-percentile (p95) delivered packet delay within the three-millisecond budget and preserves 0.892 normalized BE throughput, versus 0.534 under fixed maximum reservation.

Keywords: NR-V2X; wireless resource allocation; safety traffic; best-effort traffic; physical resource block scheduling; burst traffic; low-latency communication; HARQ; predictive scheduling

1. Introduction

Safety-critical wireless traffic is usually discussed through stringent reliability and latency targets [1–3]; however, the scheduling problem is often shaped by a more ordinary constraint: the same radio resources must also serve non-safety traffic. In New Radio Vehicle-to-Everything (NR-V2X) and related low-latency systems [4,5], safety traffic denotes warning or control packets that must be delivered within a short deadline, whereas best-effort (BE) traffic denotes non-safety data that can tolerate delays but still requires usable capacity. This BE traffic may include map updates, telemetry, cooperative information exchange, user data, or other non-deadline services sharing the same radio pool; the sustained starvation of this traffic is therefore a real coexistence cost, not only a secondary throughput loss. A static safety reservation is easy to implement; however, it is rarely efficient in coexistence settings [6–8]. If the reservation is small, safety packets are exposed during stress. If it is large, the BE capacity is reduced even during normal operation.



Received: 11 June 2026

Revised: 30 June 2026

Accepted: 3 July 2026

Published: 8 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Difficult cases are often transient rather than simply high-load. In these cases, the channel and traffic change together. A burst episode can increase the arrival rate of safety packets, lower the success probability of each transmission attempt, and create a retry carryover that remains after the first failed attempts [5,9,10]. These effects interact over short horizons. The scheduler may still see a moderate queue at the beginning of the burst, but the system has already accumulated demand that the baseline allocation cannot serve. This gap between the visible backlog and unresolved demand is the main motivation of this study.

At the radio-protocol level, NR-V2X already includes mechanisms that shape usable sidelink resources, such as configured resource pools, sensing, reservation, reselection, feedback, and coordination. These mechanisms address resource availability and collision control before the scheduler-layer decision considered in this study. However, the question of scheduler-layer coexistence still remains: once a usable resource set and a packet success operating point are given, the system must decide how much capacity should be temporarily assigned to safety traffic and how much should remain available for the BE service. This study adopts this boundary [4,11–13].

Queue-reactive scheduling is a natural response and an important baseline in low-latency resource allocation [14–16]. The limitation of this method is its timing. By the time the queue is sufficiently large to trigger extra protection, some packets may already be close to their deadlines or retry limits [17,18]. A fixed high-safety reservation has the opposite behavior: it protects early, but it also keeps physical resource blocks (PRBs) away from BE traffic when no extra safety protection is needed [19]. Neither approach directly asks whether the current slot is part of a recoverable burst, where temporary safety protection is useful but full-time safety reservation is wasteful.

Within this boundary, the model maintains the mechanisms that affect the allocation decision: burst persistence, packet deadlines, retry states, finite PRB budgets, BE coupling, and a calibrated success probability that depends on the operating regime and hybrid automatic repeat request (HARQ)-like combining. Lower-layer physical-layer (PHY) and medium access control (MAC) details, such as waveform processing, sensing procedures, sensing-based resource selection, and control signaling, are represented through the effective PRB budget and success calibration, following the scheduler-layer abstraction used in NR-V2X system studies [11–13]. This ensures that the evaluation is centered on whether the scheduler-visible memory of recent underservice can improve safety protection without imposing a permanent high-safety reservation.

The proposed method is Memory-Based Predictive Allocation (MPA). MPA adds a decayed transient-deficit state to the usual queue, deadline, and retry information. The deficit is a compact memory of recent unresolved safety demands, updated from arrivals, failed attempts that remain eligible for retry, and delivered packets. At each slot, MPA enumerates the feasible safety PRB increments and compares the resulting short-horizon outcomes in a fixed order: avoid urgent safety losses first, reduce the near-term residual workload second, and use the smallest extra PRB allocation among comparable choices.

The finite-action rule is reproducible because every candidate action is evaluated explicitly, and each term in the allocation decision remains visible. The NR-V2X resource pool mapping specifies the PRB and timing assumptions, the channel/success calibration records the signal-to-interference-plus-noise-ratio (SINR)-to-success operating points, and the application-level validation checks the delivered packet delay and BE floor behavior across urban nominal, dense non-line-of-sight (NLOS) stress, and short-deadline heavy-load scenarios.

The main contributions are summarized as follows:

- A calibrated scheduler-layer NR-V2X safety-BE allocation model with explicit PRB coupling, deadline-constrained safety queues, retry states, burst persistence, and reported channel/success operating points;
- A transient-deficit memory state and finite-action one-step recourse rule that compare candidate allocations through urgent safety losses, residual workload, and PRB economy;
- A conservative online success calibration step that treats the calibrated channel curve as a prior and corrects the planning probability with realized transmission outcomes;
- A recoverability-aware PRB guard that uses baseline and maximum allocation shortfall ratios to reduce unnecessary safety expansion outside the recoverable burst stress region;
- A validation protocol with calibrated channel/success operating points, application-level delay checks, component removal ablations, calibration sweeps, and recoverable region coverage.

2. Related Work

NR-V2X provides the radio settings for this study. The Third-Generation Partnership Project (3GPP) specifications define the Vehicle-to-Everything (V2X) evaluation methodology, service architecture, sidelink resource pools, MAC operation, HARQ support, radio resource control (RRC) configuration, and physical-layer procedures used to interpret PRB-level allocation [4,5,20–23]. Release 17 lists sidelink enhancements among its radio projects, and Release 18 continues the 5G-Advanced evolution of sidelink, proximity, location, and positioning features [24,25]. Tutorial and standards-oriented studies provide a broader picture: New Radio (NR) sidelink extends earlier Long-Term Evolution V2X (LTE-V2X) operation toward lower latency, higher reliability, groupcast/unicast support, and advanced driving services [26–29]. Comparative V2X studies also show why the radio abstraction cannot be chosen casually; the Institute of Electrical and Electronics Engineers (IEEE) 802.11bd, LTE-V2X, and NR-V2X differ in access behavior, packet timing, coding, and physical-layer robustness [30–32].

Most NR-V2X resource allocation studies focus on sidelink resource selection. In Mode 2, user equipment (UE) devices select resources from a pool using sensing, semi-persistent scheduling (SPS), dynamic selection, reservation fields, sidelink control information (SCI), resource re-evaluation, pre-emption, and inter-UE coordination (IUC); therefore, the main concerns are collisions, reselection, hidden nodes, exposed nodes, and the effects of aperiodic traffic on packet reception [11–13,21,22]. These standardized procedures determine which sidelink resources are available or risky at the MAC/RRC layer. The scheduler-layer question studied here starts after this abstraction and asks how much of an admissible PRB pool should be temporarily assigned to safety. Recent studies have examined reuse distance-aided selection, Mode 1 dynamic resource adaptation, time-domain collision behavior, the coexistence of sidelink traffic types, and directional beamformed resource selection with SCI [33–37]. Link-level and physical-layer simulation studies connect numerology, HARQ, block error rate (BLER), range, latency, and packet reception ratio (PRR) behavior [9]. Learning-based resource allocation has also been studied for NR-V2X collision mitigation, joint communication and sensing, physical sidelink feedback channel (PSFCH)-assisted Mode 2 operation, graph attention-based cellular Vehicle-to-Everything (C-V2X) allocation, and state-aware V2X resource allocation [10,38–41]. Aperiodic event-triggered traffic is also an important part of the V2X picture, with Decentralized Environmental Notification Messages (DENMs) providing a standard example of hazard-triggered communication [42]. These studies ground the channel, resource pool, and traffic assumptions used in this study. The decision considered here is narrower at the radio-protocol level but different at the

scheduler level: for a shared PRB pool, it determines how much safety capacity should be assigned temporarily when burst stress appears.

A second line of work studies the coexistence of low-latency traffic with broadband or best-effort traffic. Ultra-reliable low-latency communication (URLLC) surveys and tutorials emphasize the joint pressure of latency, reliability, short packets, and uncertainty [1–3,43]. In enhanced mobile broadband (eMBB)/URLLC coexistence, the central issue is not only whether URLLC can be protected, but also the extent of broadband service loss while protecting it. Early joint scheduling work made this trade-off explicit through puncturing and rate loss models [6]. Subsequent studies examined radio access network (RAN) slicing, queue-based optimization, fair multiplexing, contract-based scheduling, deep reinforcement learning (DRL), spectrum partitioning, reconfigurable intelligent surface (RIS)-assisted resource optimization, highway V2X scheduling, and resource-adaptive client scheduling [7,8,16,19,44–49]. Recent highway work, including SOVANET+ (an adaptive extension of the Scheduler Optimized for Vehicular Ad Hoc Network), is especially relevant because it separates critical and non-critical V2X services under loads while using channel and RAN load information [48]. This literature motivates the comparator families used in the evaluation: fixed reservation, queue-triggered adaptation, stronger low-latency protection, and dynamic allocation under an explicit BE floor.

The proposed scheduler is also related to queue-, deadline-, and debt-aware scheduling. MaxWeight and stochastic network optimization provide the classical reason for treating the backlog as a useful control signal [14,15]. Real-time scheduling and wireless quality-of-service (QoS) theory add a second signal: a packet with little remaining time should not be treated as a packet with a loose deadline [17,18]. Lyapunov drift-plus-penalty and debt-based methods add another useful idea: a running service shortfall can guide resource decisions under long-term constraints [15,18,50]. Recent multi-hop wireless scheduling work uses a delay-deficit concept to reason about service guarantees through inter-scheduling times and route-level deadline feasibility [51]. The transient-deficit state follows the same broad service shortfall intuition, while its role is local to a single scheduler-layer PRB decision: it is a decayed burst-timing signal used inside finite PRB enumeration with deadline, retry, online success, and BE floor constraints. Therefore, the analysis focuses on the recoverable region and calibrated operating behavior rather than the throughput optimality theorem.

MPA uses a small amount of scheduler-side memory to address a specific timing issue. At the beginning of a burst, the visible queue can still appear manageable, even though new arrivals and failed attempts have already made the next few slots more difficult to manage. The one-step comparison is in the same broad family as finite-horizon dynamic programming and model-predictive control (MPC), where the current action is selected after a short predicted continuation is examined [52,53]. Recurrent neural networks (RNNs) and long short-term memory (LSTM) schedulers can learn richer temporal predictors, provided that representative training traces and a deployment observation interface are available [3,39,41,54]. The allocation rule studied here keeps the predictor explicit, considering the burst state, arrivals, visible queue, deadlines, retries, recent unresolved demand, and BE floor. Therefore, it occupies a middle ground between simple reactive rules and heavier MPC, RNN/LSTM, DRL, SOVANET+-style service scheduling, and long-horizon optimization approaches. Its target is the recoverable burst stress regime, where the baseline safety capacity is insufficient but temporary PRB headroom exists.

3. Materials and Methods

3.1. System Model and Scope

A slotted scheduler is considered, which shares a finite PRB pool between deadline-constrained safety traffic and BE traffic. The model is written at the scheduler decision layer, with both centralized and distributed NR-V2X interpretations. In a next-generation Node B (gNB)-assisted setting, the rule can be read as a supervisory PRB partition decision before lower-layer scheduling. In a distributed Mode 2 setting, sensing, SCI decoding, reservation, pre-emption, IUC, and reselection first shape the usable resource set and the planning success probability. The slot-level question is how much of this usable PRB pool should be assigned to safety traffic when burst stress creates an unresolved demand. Tables 1 and 2 summarize the main symbols used in the model, as well as the allocation rules, comparator definitions, and evaluation metrics.

In this ordering, Mode 2 collision avoidance and coordination procedures appear as resource availability inputs before PRB partitioning. Their effects enter through the finite action set, guard/control margin, and online-calibrated success probabilities.

The NR-V2X interpretation is kept explicit. The simulated resource pool corresponds to a 20 megahertz (MHz) bandwidth part (BWP) with 30 kilohertz (kHz) subcarrier spacing (SCS), giving 51 nominal PRBs. One PRB is reserved as a guard/control margin, and the scheduler allocates from the remaining $N_{\text{PRB}} = 50$ PRBs. One scheduler decision represents 1 millisecond (ms), i.e., two 0.5 ms NR slots at this numerology. A compact safety protocol data unit (PDU) is represented by $n_{\text{pkt}} = 2$ PRBs under the adopted modulation and coding scheme (MCS) setting. These assumptions are listed in Table 3.

Table 1. System, traffic, service, and PRB budget symbols used in the model.

Symbol	Meaning	Unit/Range
<i>Time and regime</i>		
t	Scheduler decision slot index	1 ms slot
B_t	Operating regime; normal if 0, burst stress if 1	{0, 1}
$p_{\text{NB}}, p_{\text{BN}}$	Normal-to-burst and burst-to-normal transition probabilities	Probability
<i>Traffic</i>		
A_t	Number of new safety arrivals in slot t	Packets
$\lambda_{\text{N}}, \lambda_{\text{B}}$	Per-user safety arrival probabilities in normal and burst regimes	Probability
λ_{BE}	BE arrival probability used in the traffic model	Probability
$N_{\text{UE}}^{\text{S}}, N_{\text{UE}}^{\text{BE}}$	Numbers of safety and BE users	Users
<i>Queue and retry</i>		
$Q_t, Q_t $	Safety queue and its visible length at beginning of slot t	Packets
$d_i(t)$	Remaining deadline of safety packet i	Slots/ms
$r_i(t), R_{\text{max}}$	Retry count of packet i and maximum retry count	Attempts
R_t^{pend}	Retry-labeled packets already present in the queue at the decision instant	Packets
R_t^{new}	Newly generated retry carryover from current-slot failed attempts	Packets
<i>Service</i>		
$p_{\text{succ}}(B_t, r_i), \bar{p}_i(t)$	Model success probability and online-calibrated planning probability	Probability
$p_{b,r}^0$	Prior success probability from the calibrated channel curve for regime b and retry count r	Probability
$n_{b,r}^+, n_{b,r}^-$	Online success and failure counts used by the calibration step	Attempts
$\alpha_{b,r}, \beta_{b,r}$	Beta posterior parameters for online success calibration	Unitless
S_t	Safety packets delivered in slot t	Packets
<i>PRB budget</i>		
N_{PRB}	Schedulable PRB pool	PRBs
$K_{\text{base}}, K_{\text{max}}$	Baseline and maximum admissible safety allocations	PRBs
$\Delta K_t, \Delta K_{\text{max}}$	Extra safety allocation in slot t and its maximum admissible value	PRBs
$K_t^{\text{S}}, K_t^{\text{BE}}$	Safety and BE PRB allocations in slot t	PRBs
$K_{\text{BE}}^{\text{min}}$	Minimum BE PRB reservation	PRBs
n_{pkt}	PRB demand of one safety packet transmission attempt	PRBs/attempt

Table 2. Allocation, comparator, regime analysis, and evaluation symbols.

Symbol	Meaning	Unit/Range
<i>Scheduler state and decision rule</i>		
$\mathcal{A}$	Finite set of admissible extra-safety PRB increments	PRBs
D_t, γ	Transient-deficit memory and decay factor	Packets, unitless
s_t	Scheduler-visible state used by MPA	–
<i>MPA planning</i>		
$C_t(\Delta K)$	Safety transmission attempts available under candidate increment ΔK	Attempts
$\hat{Q}_{t+1}$	Predicted next-slot safety queue	Weighted packets
$\hat{M}_t, \hat{W}_t$	Predicted urgent loss and residual workload terms	Packets
$\hat{D}_{t+1}$	Predicted next-slot deficit under a candidate action	Packets
$\hat{V}_{t+1}, \mathcal{J}_t$	Continuation value and current candidate comparison tuple	–
$\eta_{BE}(t), g_t$	Guarded BE opportunity cost weight and recoverability multiplier	Unitless
<i>Comparator policies</i>		
U_t, P_t	Urgency pressure and high-protection pressure score	Unitless
$m_B(B_t)$	Burst multiplier in the high-protection comparator	Unitless
$\alpha_Q, \alpha_A, \alpha_P$	Weights for queue, arrivals, and pressure-to-PRB mapping	Unitless
<i>Regime analysis</i>		
L_t^{guard}	Scheduler-visible demand used by the recoverability guard	Packets
P_t^{cap}	Mean online-calibrated planning probability used for the guard capacity estimate	Probability
$\rho_{\text{base}}, \rho_{\text{max}}$	Baseline insufficiency and recovery ratios	Unitless
<i>Evaluation</i>		
$P_{\text{miss}}, P_{\text{del}}$	Deadline miss ratio and delivery ratio	Ratio
$T_{BE}, T_{BE}^{\text{ref}}$	BE throughput and static-reference BE throughput	Mbps
BE_{norm}	BE throughput normalized by the static reference	Ratio

Table 3. NR-V2X settings used by the scheduler-layer model.

Item	Value	Role in the Simulation
<i>Radio configuration</i>		
Carrier and numerology	5.9 GHz, 30 kHz SCS	Carrier and slot-timing setting used for the NR-V2X mapping
Resource pool	20 MHz BWP, 51 nominal PRBs	Starting resource pool before the scheduler-level guard/control margin
<i>Scheduler abstraction</i>		
Decision interval	1 ms = 2 NR slots	Time unit for deadlines, queue updates, and PRB decisions
Schedulable PRBs	50 PRBs	PRB pool available to the safety-BE allocation rule
<i>Traffic and resource constraints</i>		
Safety packet demand	2 PRBs per attempt	Compact 96-byte safety PDU under the adopted MCS setting
Safety allocation range	6–30 PRBs	Range used in the dense NLOS stress and short-deadline heavy-load cases
BE floor	10 PRBs	Minimum BE reservation enforced in all evaluated cases
Delay budget	2–4 ms	Application-level safety deadlines used in the curated cases

3.1.1. Traffic and Burst State

Time is indexed by slots $t = 0, 1, \dots$. The operating condition is represented by a two-state process

$$B_t \in \{0, 1\}, \quad (1)$$

where $B_t = 0$ denotes the normal regime and $B_t = 1$ denotes burst stress. The regime follows a Markov transition model:

$$\Pr(B_{t+1} = 1 \mid B_t = 0) = p_{NB}, \quad \Pr(B_{t+1} = 0 \mid B_t = 1) = p_{BN}. \quad (2)$$

This is a compact way to represent persistence. A burst is not treated as an isolated bad slot; it may last long enough for failed attempts and new arrivals to overlap.

Safety arrivals are regime-dependent. With N_{UE}^{S} safety users,

$$A_t \sim \begin{cases} \text{Binomial}(N_{\text{UE}}^{\text{S}}, \lambda_{\text{N}}), & B_t = 0, \\ \text{Binomial}(N_{\text{UE}}^{\text{S}}, \lambda_{\text{B}}), & B_t = 1, \end{cases} \quad (3)$$

where $\lambda_{\text{B}} > \lambda_{\text{N}}$. The burst state therefore affects both the load and channel quality. BE arrivals are modeled separately and stored in a finite BE queue. The exact BE arrival law is not the central issue; what matters is that BE service uses the same PRB pool as safety protection.

Each Monte Carlo replicate starts from $B_0 = 0$, an empty safety queue, $D_0 = 0$, zero retry carryover, and an empty BE queue. A warm-up interval is discarded before the metrics are collected; therefore, the reported values are not driven by these startup conditions.

3.1.2. Safety Queue and Retry-Aware Service

Each safety packet carries two scheduler-visible attributes: its remaining deadline $d_i(t)$ and retry count $r_i(t)$. A packet leaves the queue when it is delivered, when its deadline expires, or when its retry limit is exceeded. If packet i is attempted in slot t , its success probability is

$$p_i(t) = p_{\text{succ}}(B_t, r_i(t)). \quad (4)$$

The function $p_{\text{succ}}(\cdot)$ is obtained from a calibrated scheduler-layer channel model. It uses the regime-dependent SINR, an MCS-based BLER curve, PRB diversity, and HARQ-like combining. SINR values are reported in decibels (dB). The deterministic mean values are reported in Table 4; for example, the nominal burst first-attempt success probability is 0.582 and rises to 0.907 by retry count 3 under the HARQ gain cap. In the stress burst setting, the first-attempt value is 0.275 and rises to 0.727 under the retry cap. The simulation still applies per-slot SINR jitter for realized transmissions; the reported values are the planning means used by the scheduler.

Table 4. Channel/success operating points used in the evaluated scenarios. Values are deterministic mean planning probabilities; stochastic simulation additionally applies per-slot SINR jitter.

Scenario	Normal SINR	Burst SINR	Retry Cap	Burst $p_{\text{succ}}(r = 0)$	Burst $p_{\text{succ}}(r = R_{\text{max}})$
Urban nominal	14.0 dB	5.0 dB	3	0.582	0.907
Dense NLOS stress	14.0 dB	3.0 dB	4	0.275	0.727
Short-deadline heavy load	14.0 dB	3.0 dB	4	0.275	0.727

Let S_t be the number of safety packets delivered in slot t , and let R_t^{new} be the newly generated retry carryover from failed attempts in the same slot. Both are endogenous: they depend on the current queue composition, burst state, retry counts, and PRB allocation selected in that slot. The term R_t^{new} counts only current-slot failed attempts that remain eligible for retry; retry-labeled packets already present in Q_t are represented in the queue state and are not added again in the physical queue recursion.

3.1.3. PRB Partition

The scheduler reserves a baseline safety budget K_{base} and may add a nonnegative increment ΔK_t . The safety and BE allocations are

$$K_t^{\text{S}} = K_{\text{base}} + \Delta K_t, \quad K_t^{\text{BE}} = N_{\text{PRB}} - K_t^{\text{S}}. \quad (5)$$

The decision must respect a safety ceiling and a BE floor:

$$0 \leq \Delta K_t \leq K_{\max} - K_{\text{base}}, \quad K_t^{\text{BE}} \geq K_{\text{BE}}^{\min}. \quad (6)$$

Equivalently, the largest admissible increment is

$$\Delta K_{\max} = \min(K_{\max} - K_{\text{base}}, N_{\text{PRB}} - K_{\text{base}} - K_{\text{BE}}^{\min}). \quad (7)$$

Since safety packets are served in units of n_{pkt} PRBs, the feasible action set is finite:

$$\mathcal{A} = \left\{ 0, n_{\text{pkt}}, 2n_{\text{pkt}}, \dots, \left\lfloor \frac{\Delta K_{\max}}{n_{\text{pkt}}} \right\rfloor n_{\text{pkt}} \right\}. \quad (8)$$

This finite action set is important for the interpretation of MPA. The online decision is solved by enumeration over feasible PRB increments, not by adjusting a continuous tuning knob.

3.1.4. Transient-Deficit Memory

The model includes a decayed transient-deficit state D_t :

$$D_{t+1} = \gamma[D_t + A_t + R_t^{\text{new}} - S_t]^+, \quad 0 < \gamma < 1. \quad (9)$$

The deficit is a memory of recent unresolved safety demand. It grows when arrivals and newly generated retry carryover exceed the delivered safety service, and it decays after recovery. The queue records packets that are still waiting; the deficit records the recent service shortfall that produced the pressure. This distinction is useful during burst onset, when the visible queue may still be moderate but the combination of fresh arrivals, failures, and retries already signals that the baseline allocation is falling behind.

3.1.5. Observed State and Modeling Boundary

At the beginning of a slot, the scheduler observes

$$s_t = (B_t, A_t, D_t, Q_t, \{d_i(t), r_i(t)\}_{i \in Q_t}), \quad (10)$$

where Q_t is the safety queue. This state contains the information used by MPA and by the structural baselines.

The model boundary follows the scheduler-layer question. It retains the elements that change the safety-BE allocation decision: burst persistence, arrivals, deadlines, retries, success probabilities, finite PRB budgets, and BE coupling. The lower-layer procedures enter through the effective resource pool and calibrated success map, keeping the scheduler question explicit. The evaluation is a calibrated scheduler-layer study with NR-V2X parameter settings.

In the simulations, $n_{\text{pkt}} = 2$ means that one compact safety PDU occupies two PRBs under the adopted MCS. This value fixes the packet resource granularity used in the main operating points. With another MCS or a lower-power IoT-oriented waveform, both the packet-to-PRB mapping and the BLER/success probability curve would change. In the scheduler abstraction, these changes enter through n_{pkt} and $p_{\text{succ}}(\cdot)$. Recent IoT modulation studies, including energy-efficient CPM and index-aided DCSK designs, provide examples of how modulation choices can shift link-layer robustness and spectral costs [55,56]. A later sensitivity check varies the packet-resource granularity around this setting.

Table 5 records how lower-layer NR-V2X mechanisms enter the scheduler-layer abstraction. PHY/MAC effects first shape the usable resource set or the planning success

probability; the allocation rule then decides how much of this usable set is assigned to safety traffic.

Table 5. Scheduler-layer representation of lower-layer NR-V2X mechanisms.

NR-V2X Mechanism	Scheduler-Layer Representation	Role in the Evaluation
Mode 1 scheduling	gNB-side supervisory PRB partition	Direct interpretation of the safety-BE PRB decision
Mode 2 resource selection and sensing-based SPS	Usable resource set determined before the scheduler-layer partition	The allocation rule partitions the admissible PRB set after sensing, reservation, and reselection constraints
SCI and PSFCH signaling	Effective PRB budget or guard/control margin	Reduces schedulable resources available to safety and BE traffic
Pre-emption, IUC, and re-evaluation	Candidate-set filtering before allocation	Removes risky or unavailable PRB increments from the admissible action set
Sensing errors, hidden nodes, and collisions	Conservative planning success probability and online correction	Moves the operating point toward higher safety pressure when observed outcomes are worse than the calibrated curve
Geometry, distance, and interference	Scenario-dependent success calibration	Reflected through the nominal and dense NLOS channel/success operating points

3.2. Memory-Based Predictive Allocation

MPA chooses the extra safety allocation ΔK_t at the beginning of each slot. The input is the observed state s_t in Section 3.1.5, and the output is one element of the feasible set $\mathcal{A}$. The rule uses a short horizon by design. It takes the current queue as observed, propagates the likely one-slot consequences of each candidate allocation, and then checks how much recourse is needed in the next slot. This keeps the online decision sufficiently small for enumeration while still accounting for retry carryover and transient underservice.

3.2.1. Current-Slot Propagation

For a candidate action $\Delta K \in \mathcal{A}$, the number of safety transmission attempts available in the current slot is

$$C_t(\Delta K) = \left\lfloor \frac{K_{\text{base}} + \Delta K}{n_{\text{pkt}}} \right\rfloor. \quad (11)$$

Packets are considered in queue order, with their own deadlines and retry counts. For attempted packet i , the model success probability is first obtained from the calibrated link abstraction, $p_{\text{succ}}(B_t, r_i(t))$. The scheduler then applies a lightweight online calibration. For each burst regime $b \in \{0, 1\}$ and retry count r , let $p_{b,r}^0 = p_{\text{succ}}(b, r)$ be the calibrated model value, and let $n_{b,r}^+(t)$ and $n_{b,r}^-(t)$ be the realized success and failure counts observed before the current decision. With prior strength c_{cal} , the Beta posterior is

$$\alpha_{b,r}(t) = c_{\text{cal}} p_{b,r}^0 + n_{b,r}^+(t), \quad \beta_{b,r}(t) = c_{\text{cal}} (1 - p_{b,r}^0) + n_{b,r}^-(t). \quad (12)$$

The posterior mean and standard deviation proxy are

$$\mu_{b,r}(t) = \frac{\alpha_{b,r}(t)}{\alpha_{b,r}(t) + \beta_{b,r}(t)}, \quad \sigma_{b,r}(t) = \sqrt{\frac{\mu_{b,r}(t)(1 - \mu_{b,r}(t))}{\alpha_{b,r}(t) + \beta_{b,r}(t) + 1}}. \quad (13)$$

The lower planning value from this counter is

$$\underline{p}_t(b, r) = \max\{p_{\text{min}}, \mu_{b,r}(t) - z_{\text{cal}} \sigma_{b,r}(t)\}, \quad (14)$$

where $p_{\min}$ prevents a zero-probability collapse and z_{cal} controls the conservative margin. MPA uses the following planning success probability:

$$\bar{p}_i(t) = \min \left\{ p_{\text{succ}}(B_t, r_i(t)), (1 - \omega) p_{\text{succ}}(B_t, r_i(t)) + \omega p_{\text{t}}(B_t, r_i(t)) \right\}, \quad (15)$$

where $0 \leq \omega \leq 1$ is a fixed blending factor. The minimum keeps the planner conservative: online observations can reduce the planning probability under model mismatch, but they do not make the scheduler more optimistic than the calibrated channel model. A successful packet leaves the queue; a failed packet either becomes retry mass in the next slot or becomes irrecoverable if its deadline or retry budget is exhausted. Thus, the calibrated channel model acts as the prior planning curve, while the online posterior correction provides a conservative adjustment from realized transmission outcomes.

The propagation gives three quantities for each candidate action:

- $\hat{M}_t(\Delta K)$, the expected urgent loss in the current slot;
- $\hat{S}_t(\Delta K)$ and $\hat{R}_t^{\text{new}}(\Delta K)$, the expected delivered packets and newly generated retry carryover;
- $\hat{Q}_{t+1}(\Delta K)$, a weighted next-slot queue containing unserved packets and failed packets that remain eligible for retry.

The next deficit used by the planner is then

$$\hat{D}_{t+1}(\Delta K) = \gamma \left[D_t + A_t + \hat{R}_t^{\text{new}}(\Delta K) - \hat{S}_t(\Delta K) \right]^+. \quad (16)$$

3.2.2. One-Step Recourse

A current-slot action may reduce the visible queue while leaving a difficult next slot. MPA therefore evaluates a one-step continuation problem. From the predicted state

$$\hat{s}_{t+1}(\Delta K) = \left(\hat{B}_{t+1}, \hat{D}_{t+1}(\Delta K), \hat{Q}_{t+1}(\Delta K) \right), \quad (17)$$

the method again enumerates feasible safety increments for the next slot. The next burst condition is represented by the Markov transition probability from the observed B_t , so the future success probability is the corresponding normal–burst mixture.

For a next-slot action $\Delta K'$, let $\hat{M}_{t+1}(\Delta K'; \hat{s}_{t+1})$ denote the expected urgent loss after the predicted queue has aged by one slot, and let $\hat{W}_{t+2}(\Delta K'; \hat{s}_{t+1})$ denote the residual workload after this action. These two quantities are obtained by applying the same service, retry, and deficit propagation used in (16) to the predicted state $\hat{s}_{t+1}(\Delta K)$. The residual workload is defined as

$$\hat{W}_{t+2} = \sum_{i \in \hat{Q}_{t+2}} w_i + \hat{D}_{t+2}, \quad (18)$$

where w_i is the probability mass of a predicted packet state and $\hat{Q}_{t+2}$ is the predicted queue after the candidate next-slot action. The tuple minimum used below is read in a fixed left-to-right order. For two tuples $x = (x_1, \dots, x_L)$ and $y = (y_1, \dots, y_L)$, x is preferred to y if there exists an index ℓ such that $x_j = y_j$ for every $j < \ell$ and $x_\ell < y_\ell$. If all entries are equal, the two tuples are tied. Thus, the minimum symbol in the next equation is not a scalar weighted sum:

$$\hat{V}_{t+1}(\hat{s}_{t+1}) = \min_{\Delta K' \in \mathcal{A}} \left(\hat{M}_{t+1}(\Delta K'), \hat{W}_{t+2}(\Delta K'), \eta_{\text{BE}}(t) \frac{\Delta K'}{n_{\text{pkt}}}, \frac{\Delta K'}{n_{\text{pkt}}} \right). \quad (19)$$

Equation (19) assigns the predicted state its best next-slot continuation tuple. It is obtained by taking the predicted state in (17), applying the same one-slot service and

retry propagation to each feasible continuation action $\Delta K'$, using (18) for the remaining workload, and then selecting the best tuple under the fixed-ordered comparison. The entries are the expected urgent loss, residual workload, guarded BE opportunity cost, and PRB increment. A candidate with a lower urgent loss is preferred before the residual workload or BE cost is considered. The BE cost term is considered only after the current urgent loss term in $\mathcal{J}_t$ and the urgent loss component embedded in the recourse tuple $\hat{V}_{t+1}$ have been compared. The coefficient $\eta_{BE}(t)$ is the guarded BE opportunity cost used by the scheduler; it is distinct from the BE arrival probability used in the traffic model. The BE term is written in packet-sized PRB units because actions are enumerated in this unit.

3.2.3. Recoverability-Aware PRB Guard

The allocation rule also includes a recoverability-aware PRB guard. The guard translates the recoverable region idea into a cost multiplier inside the existing candidate comparison. The ratio test keeps feasible actions in the comparison and changes only the BE cost entry, after the current urgent loss and one-step safety carryover have been evaluated. The demand signal used by the guard is

$$L_t^{\text{guard}} = \max \left\{ D_t + A_t + R_t^{\text{pend}}, \sum_{i \in Q_t} \mathbf{1}\{d_i(t) \leq 2\} \right\}. \quad (20)$$

Here, L_t^{guard} is the scheduler-visible demand used only by the guard. It takes the larger of two signals: the recent unresolved load $D_t + A_t + R_t^{\text{pend}}$ and the number of packets whose remaining deadline is at most two slots. The term R_t^{pend} denotes retry-labeled packets already present in the queue at the decision instant. This is different from R_t^{new} in (9), which is generated only after current-slot failed attempts. The planning capacity term is

$$p_t^{\text{cap}} = \begin{cases} |Q_t|^{-1} \sum_{i \in Q_t} \bar{p}_i(t), & |Q_t| > 0, \\ \bar{p}_t(B_t, 0), & |Q_t| = 0, \end{cases} \quad (21)$$

where p_t^{cap} is the packet success probability, used only for the guard capacity estimate. When the queue is nonempty, it is the mean of the online-calibrated planning probabilities $\bar{p}_i(t)$ defined in (15); when the queue is empty, it uses the current regime and first-attempt retry class. The two ratio tests are

$$\rho_{\text{base}}(t) = \frac{L_t^{\text{guard}}}{\max\{C_{\text{base}} p_t^{\text{cap}}, \epsilon\}}, \quad \rho_{\text{max}}(t) = \frac{L_t^{\text{guard}}}{\max\{C_{\text{max}} p_t^{\text{cap}}, \epsilon\}}, \quad (22)$$

where $C_{\text{base}} = \lfloor K_{\text{base}}/n_{\text{pkt}} \rfloor$ is the number of packet attempts supported by the baseline safety allocation, $C_{\text{max}} = \lfloor K_{\text{max}}^{\text{adm}}/n_{\text{pkt}} \rfloor$ is the number supported by the largest admissible safety allocation after the BE floor is enforced, $K_{\text{max}}^{\text{adm}}$ is the admissible safety ceiling, and ϵ avoids division by zero. The products $C_{\text{base}} p_t^{\text{cap}}$ and $C_{\text{max}} p_t^{\text{cap}}$ are the expected planning capacities in packets; they are the guard-level counterparts of the slot service capacity $C_t(\Delta K)$ in the current-slot propagation. The guarded BE cost is

$$\eta_{BE}(t) = g_t \eta_{BE}, \quad (23)$$

where

$$g_t = \begin{cases} g_{\text{suf}}, & \rho_{\text{base}}(t) \leq 1 \text{ and } \sum_{i \in Q_t} \mathbf{1}\{d_i(t) \leq 2\} = 0, \\ g_{\text{inf}}, & \rho_{\text{max}}(t) \geq 1, \\ g_{\text{rec}}, & \text{otherwise.} \end{cases} \quad (24)$$

The constants g_{suf} , g_{inf} , and g_{rec} are fixed nonnegative multipliers. The implementation uses $g_{\text{suf}} > g_{\text{rec}}$ to reduce unnecessary extra safety allocation when the baseline appears sufficient and no packet is immediately urgent. It uses $g_{\text{inf}} > g_{\text{rec}}$ near the resource ceiling, where additional PRBs may not recover the safety backlog and can mainly consume BE capacity. In the recoverable middle case, $g_{\text{rec}} = 1$, and the standard MPA ordering is left unchanged. Thus, the guard affects only the BE cost entry of the ordered comparison. Safety-improving actions remain available whenever they can reduce the current urgent loss, even if one of the ratio thresholds has been crossed. In the near-ceiling case, if no packet is currently urgent, the implementation also limits the current-slot candidate list while retaining a configurable fraction of the near-urgent coverage for packets with $d_i(t) \leq 2$. This prevents workload-only safety expansion when the ratio test indicates that the backlog is not recoverable with the available PRB ceiling. The cap is not applied in the recoverable middle case or when the current urgent loss can still be reduced.

3.2.4. Selection Rule

The current action is selected by comparing candidate tuples

$$\mathcal{J}_t(\Delta K) = \left(\hat{M}_t(\Delta K), \hat{V}_{t+1}(\hat{s}_{t+1}(\Delta K)), \hat{W}_{t+1}(\Delta K) + \eta_{\text{BE}}(t) \frac{\Delta K}{n_{\text{pkt}}}, \frac{\Delta K}{n_{\text{pkt}}} \right), \quad (25)$$

where

$$\hat{W}_{t+1}(\Delta K) = \sum_{i \in \hat{\mathcal{Q}}_{t+1}(\Delta K)} w_i + \hat{D}_{t+1}(\Delta K). \quad (26)$$

The selected allocation is

$$\Delta K_t^* = \arg \min_{\Delta K \in \mathcal{A}} \mathcal{J}_t(\Delta K). \quad (27)$$

The minimum in (27) uses the same left-to-right tuple comparison. The current urgent loss entry is compared first. If candidate actions are tied on this entry, the recourse tuple is compared. If they remain tied, the residual workload/BE cost term and then the PRB increment are compared. Consequently, the BE cost affects the decision only when the earlier safety-priority entries do not separate the candidates. The final entry selects the smaller safety increment when the two candidates are otherwise equivalent. This provides a safety-first rule while avoiding the permanent assignment of all spare PRBs to safety.

This ordered rule selects a safety-first operating point on the safety-BE frontier. It does not trace the full Pareto set: a candidate that saves many BE PRBs by accepting a larger current urgent loss falls behind a candidate with lower urgent loss. The evaluation therefore reports the miss ratio and BE throughput together, placing this operating point against the more BE-preserving Reactive policy and the more safety-heavy Look-Ahead, LSTM, and Overprovision references.

3.2.5. Implementation Notes

The implementation follows the above equations by direct enumeration. For each feasible ΔK , it propagates the current queue, computes the predicted deficit, solves the one-step recourse enumeration, and stores the resulting tuple. The tuple with the best left-to-right priority order determines the action. The actual transmission outcomes in the simulator are sampled using the stochastic channel model, and planning uses the online-calibrated probabilities in (15). The success counters are reset at the beginning of each Monte Carlo replicate and updated from the realized attempts, including the warm-up interval.

The number of evaluated actions is bounded by the PRB headroom and the BE floor:

$$|\mathcal{A}| = 1 + \left\lfloor \frac{\Delta K_{\max}}{n_{\text{pkt}}} \right\rfloor. \quad (28)$$

The recourse step repeats the same finite enumeration once. Listing 1 summarizes the online sequence in one scheduler slot. Table 6 reports the resulting decision structure, together with the heavier two-step look-ahead reference used as a robustness check.

Listing 1. One-slot MPA allocation procedure.

1	Observe $s_t = (B_t, A_t, D_t, Q_t, \{d_i(t), r_i(t)\}_{i \in Q_t})$ at the beginning of slot t .
2	Form the feasible PRB-increment set $\mathcal{A}$ from $\Delta K_{\max}$, n_{pkt} , and the BE floor constraint.
3	Update the online success estimate from realized attempts in earlier slots and obtain the conservative planning probability $\bar{p}_i(t)$.
4	Compute $\rho_{\text{base}}(t)$, $\rho_{\max}(t)$, and the guarded BE cost coefficient $\eta_{\text{BE}}(t)$; in the nonurgent near-ceiling case, apply the candidate cap described in Section 3.2.3.
5	For each $\Delta K \in \mathcal{A}$, predict the current-slot service, urgent loss $\hat{M}_t(\Delta K)$, retry carryover, next queue $\hat{Q}_{t+1}(\Delta K)$, and next deficit $\hat{D}_{t+1}(\Delta K)$.
6	For the predicted next state, enumerate each feasible continuation action $\Delta K'$ and keep the best continuation tuple $\hat{V}_{t+1}$ under the fixed left-to-right ordered comparison.
7	Store the current candidate tuple $\mathcal{J}_t(\Delta K)$, containing the current urgent loss, continuation tuple, residual workload with guarded BE cost, and PRB increment.
8	Select ΔK_t^* as the candidate with the best-ordered tuple; if all entries are tied, choose the smaller PRB increment.
9	Apply $K_t^S = K_{\text{base}} + \Delta K_t^*$, sample realized transmissions, and update the safety queue, retry labels, deficit state, online success counters, and BE queue.

Table 6. Scheduler policies used in the safety–BE comparison.

Policy	Comparison Role	Scheduler-Visible Information	Online Decision Form
Reactive	BE-preserving baseline	Visible queue and deadlines	Queue-triggered PRB increment
High-Protection Urgency+Q+R+A MPA	Safety-protective comparator Proposed allocator	Urgency, queue, retry, and arrivals Burst state, arrivals, deficit, queue, deadlines, retries, online success estimate	Tuned score-based pressure rule Exact finite-action one-step recourse selection with recoverability-aware BE cost
Look-Ahead (H = 2) Overprovision	Heavier reference Safety ceiling/BE cost reference	Same state as MPA Fixed maximum safety allocation	Two-step finite-action reference Maximum safety PRBs subject to BE floor

3.3. Analytical Insight: Recoverable Burst Stress Regime

The deficit state in (9) is useful only if it has a clear interpretation. This section provides an interpretation of this through a simple mean service argument. The argument locates the operating region where memory should help, and the numerical claims remain with the simulation and calibration results.

3.3.1. Base-Only Deficit Reference

Consider first the reference process that uses only the baseline safety allocation. Let

$$X_t^{\text{base}} = A_t + R_t^{\text{new}} - S_t^{\text{base}} \quad (29)$$

denote the one-slot safety shortfall under the baseline allocation, where S_t^{base} is the safety service obtained from K_{base} PRBs. The corresponding deficit recursion is

$$D_{t+1}^{\text{base}} = \gamma [D_t^{\text{base}} + X_t^{\text{base}}]^+. \quad (30)$$

Because $0 < \gamma < 1$, old shortfalls are discounted. A short burst can still create a large transient deficit; however, the memory does not grow without bound once the service recovers. The state D_t therefore represents recent pressure rather than a permanent backlog.

Taking a regime-average view, the steady-state base-only deficit is bounded by

$$\frac{\gamma \mathbb{E}[X_t^{\text{base}}]^+}{1 - \gamma} \leq \mathbb{E}D_{\infty}^{\text{base}} \leq \frac{\gamma \mathbb{E}[(X_t^{\text{base}})^+]}{1 - \gamma}. \quad (31)$$

The lower bound reflects the average drift of baseline service. The upper bound reflects the average positive part of the shortfall, so it remains sensitive to burst periods even when normal periods are well served.

3.3.2. Insufficiency and Recovery Ratios

The ratio test follows directly from the expected one-slot shortfall. Under a baseline allocation, the expected drift of the shortfall term is

$$\mathbb{E}[X_t^{\text{base}}] = \mathbb{E}[A_t + R_t^{\text{new}}] - \mathbb{E}[S_t^{\text{base}}]. \quad (32)$$

Thus, the baseline tends to fall behind when $\mathbb{E}[X_t^{\text{base}}] > 0$. Equivalently, it is insufficient when the incoming safety work exceeds the expected baseline service. The same calculation can be performed at the largest admissible safety allocation. The expected service terms are estimated from the packet attempt capacity and the calibrated planning success probability,

$$\mathbb{E}[S_t^{\text{base}}] \approx \mathbb{E}[C_{\text{base}} p_t^{\text{cap}}], \quad \mathbb{E}[S_t^{\text{max}}] \approx \mathbb{E}[C_{\text{max}} p_t^{\text{cap}}], \quad (33)$$

where C_{base} and C_{max} are the packet attempt capacities at the baseline and maximum admissible safety allocations. This leads to two ratios,

$$\rho_{\text{base}} = \frac{\mathbb{E}[A_t + R_t^{\text{new}}]}{\mathbb{E}[S_t^{\text{base}}]} \quad (34)$$

and

$$\rho_{\text{max}} = \frac{\mathbb{E}[A_t + R_t^{\text{new}}]}{\mathbb{E}[S_t^{\text{max}}]}, \quad (35)$$

where S_t^{max} is the service available under the maximum admissible safety allocation. These ratios separate two different questions. If $\rho_{\text{base}} > 1$, the baseline allocation is insufficient on average and some dynamic protection is justified. If $\rho_{\text{max}} < 1$, the system still has enough PRB headroom for an adaptive policy to recover without reserving maximum safety resources all the time.

The averages in (34) and (35) are finite-window load–service summaries for a scenario or sweep point, not steady-state claims about the first few burst slots. At the controller level, the guard in (22) uses the same logic as the current observed demand L_t^{guard} and the current planning probability p_t^{cap} . The regime-level values identify where recovery is possible, whereas the slot-level guard decides whether the current burst sample is pressing against the baseline or the ceiling.

The calculation is a load–service diagnostic with the same practical logic as Lyapunov drift and debt arguments: recovery requires the service opportunity to exceed incoming

work over the interval that matters. If the shortfall process is written as a decayed debt, $D_{t+1} = \gamma[D_t + X_t]^+$, the one-step drift is governed by the sign and size of the service shortfall X_t . A negative or small positive shortfall causes the debt to decay, whereas a sustained positive shortfall keeps it elevated. Equations (34) and (35) apply this idea to the two PRB limits that matter for the scheduler: the baseline allocation and the maximum admissible safety allocation. The resulting region identifies where a short scheduler memory has useful timing information and where the same memory adds little because the baseline is already sufficient or the PRB ceiling is itself too small.

A simple drift bound also explains why the memory remains transient. For any bounded one-slot shortfall X_t ,

$$\mathbb{E}[D_{t+1} - D_t \mid D_t] \leq -(1 - \gamma)D_t + \gamma \mathbb{E}[X_t^+ \mid D_t]. \quad (36)$$

Because arrivals, retry carryover, and PRB service are bounded in the scheduler model, the positive shortfall term is finite. The bound shows that a large deficit is pulled down by the decay factor unless new positive shortfalls keep arriving. This is the property needed for a burst-timing signal: the recent unresolved demand remains visible for a few slots, but old stress does not dominate later decisions.

This means that the most relevant region for MPA is

$$\rho_{\text{base}} > 1 \quad \text{and} \quad \rho_{\text{max}} < 1. \quad (37)$$

In this region, permanently reserving the maximum safety allocation is not necessary. The scheduler needs a temporary protection signal. The transient deficit provides such a signal because it responds to arrivals, retry carryover, and realized service before the visible queue length alone fully reflects the stress.

3.3.3. Interpretation for Scheduling

The analysis also clarifies where the proposed memory is less effective. If $\rho_{\text{base}} \leq 1$, the baseline allocation is already adequate in the mean; therefore, a memory signal has little potential to improve the decision. If $\rho_{\text{max}} \geq 1$, even the largest admissible safety allocation cannot fully absorb the average demand; in this case, the resource budget, traffic load, deadline, or link conditions dominate the results. Between these extremes, memory is useful because the system is stressed but remains recoverable.

This interpretation matches the design of MPA. The deficit enters the finite-action comparison together with the deadline urgency, retry state, predicted carryover, online-calibrated success probability, and guarded BE cost. As a result, extra safety PRBs are favored when unresolved demand is building, and the allocation relaxes when the predicted workload falls or the current ratio test indicates a weak safety return for additional PRBs. This gives the deficit state a practical role: the queue length can be late under burst stress, while a decayed memory of recent shortfall gives the scheduler an earlier signal.

3.4. Evaluation Design and Validation Protocol

The evaluation is organized around the safety-BE trade-off created by burst stress. All policies run on the same slot process, PRB budget, traffic traces, retry rules, deadline model, and calibrated success probability map. The reported comparisons change the allocation rule while keeping the operating point fixed. Each Monte Carlo replicate resets the random number generator with the deterministic seed protocol exported with the results package, and all summary metrics are computed after the warm-up interval. Unless otherwise stated, the main figures, application validation results, calibration summaries, and LSTM reference results use 25 independent runs per policy/scenario. Tables report replicate means together

with 95% confidence interval half-widths when run-level replication is available. The summary values supporting the reported tables and figures are provided as Dataset S1.

The number of runs was chosen to keep the confidence intervals visible without allowing the heavier look-ahead and LSTM references to dominate the evaluation time. Increasing the number of independent runs would mainly narrow the confidence intervals at the usual $1/\sqrt{N}$ rate; it would not change the policy definitions, PRB constraints, or deterministic seed protocol. The statistical tables report standard deviations and 95% confidence interval half-widths, making the sampling uncertainty explicit.

All simulations, post-processing steps, and figure/table generation were implemented in MATLAB. MATLAB was used as a controlled simulation environment for deterministic Monte Carlo scripting, explicit finite-action enumeration, and direct reproduction of the reported tables and figures. The scheduler decisions do not rely on MATLAB-specific optimization solvers; each policy is evaluated through the same scripted slot process and finite PRB action set. Therefore, timing entries are interpreted as prototype execution measurements from one software environment. For portability, the evaluation also reports candidate evaluation counts and recourse depths, which describe the size of the online decision independently of a specific MATLAB installation.

3.4.1. Operating Regimes

Three curated regimes are used as the main application-level cases. Urban nominal uses a 4 ms deadline, moderate burst persistence, and a nominal SINR profile. Dense NLOS stress uses a 3 ms deadline, a lower burst SINR, longer retry allowance, and smaller baseline safety allocation. A short-deadline heavy load reduces the safety deadline to 2 ms, whereas the BE demand remains high. MPA uses the same controller parameters across these cases; only the operating environment changes.

The curated regimes are complemented by controlled sweeps. The calibration sweep changes the burst channel severity and HARQ-like combining assumptions. The recoverable region sweep varies the burst arrival rate, baseline safety allocation, burst persistence, and deadline. Together, these experiments separate the main performance comparison from the link abstraction sensitivity and regime coverage. Table 7 summarizes the scenarios used for application-level validation.

Table 7. Operating scenarios used in the main validation. Shared settings are $N_{\text{PRB}} = 50$, $K_{\text{max}} = 30$, $K_{\text{BE}}^{\text{min}} = 10$, $n_{\text{pkt}} = 2$, $N_{\text{UE}}^{\text{S}} = 10$, $T_{\text{sim}} = 3000$ slots, and $T_{\text{warmup}} = 200$ slots.

Parameter	Urban Nominal	Dense NLOS Stress	Short-Deadline Heavy Load
Safety deadline D_{max}	4 ms	3 ms	2 ms
Baseline safety PRBs K_{base}	8	6	6
Burst arrival probability λ_{B}	0.72	0.72	0.82
Burst recovery probability p_{BN}	0.10	0.04	0.04
Burst SINR SINR_{B}	5.0 dB	3.0 dB	3.0 dB
Retry cap R_{max}	3	4	4
BE users $N_{\text{UE}}^{\text{BE}}$	30	26	30
BE arrival probability λ_{BE}	0.80	0.72	0.80
Main summary runs	25	25	25
Application validation runs	25	25	25

3.4.2. Policies Compared

The comparison set was organized based on scheduling principles. Each policy uses the same PRB budget, BE floor, packet deadlines, retry limit, and channel success map; only the allocation rule is changed for each policy.

The comparison separates the scheduler-layer mechanisms under a single observation interface. *Static* and *Reactive* provide low-complexity reference points; *Threshold-Guard*,

Urgency + Q + R, and the high-protection variants add stronger reactive protection; *Look-Ahead* and the LSTM reference add predictive comparators. Modern SPS, SOVANET+, MPC, and DRL approaches remain important in the literature; however, their input assumptions differ, i.e., the sensing window history, channel quality reports, service-level classifications, distributed feedback, or trained predictors [12,39,48]. Maintaining the same PRB budget, BE floor, channel success map, and scheduler-visible state makes the differences attributable to the allocation rules.

- *Static* keeps the safety allocation fixed at its baseline value and therefore provides a no-adaptation reference. It represents the fixed K_{base} safety reservation used for within-scenario normalization. Fixed or semi-static resource partitioning is commonly used as a reference point when evaluating dynamic RAN slicing and URLLC/eMBB coexistence mechanisms [7]. In the NR-V2X context, this is also consistent with configured or persistent reservation viewpoints in sidelink resource pool studies [11,12,21].
- *Reactive* is the queue-reactive baseline. This increases safety protection from visible backlogs and retry-dependent service information. It uses B_t , Q_t , and retry states through the channel/retry service model, but excludes current arrivals, deadline urgency, transient deficit memory, and recourse. This policy is grounded in queue-aware low-latency scheduling, where the queuing state is used to monitor the delay and control the service [14–16].
- *Threshold-Guard* is a quantized dynamic protection controller. It selects among a small number of safety guard levels using visible queue pressure and retry-aware service prediction while excluding current arrivals, deadline urgency, transient deficit memory, and recourse. This is motivated by dynamic RAN slicing and scheduling approaches that adjust low-latency resources according to service conditions rather than maintaining a fixed partition [8].
- *Urgency + Q + R* is the main strengthened reactive comparator. It uses the burst state, visible queue length, packet deadline urgency, and retry labels but excludes current arrivals, transient deficit memory, and one-step recourse. It tests whether deadline-, queue-, and retry-aware reactivity is sufficient without memory or predictive carryover. This comparator is grounded in fair URLLC/eMBB multiplexing and deadline-sensitive scheduling formulations [17,18,44].
- *High-Protection Urgency + Q + R* and *High-Protection Urgency + Q + R + A* are deliberately strengthened urgency-aware comparators. Both use the burst state, visible queue length, packet deadline urgency, and retry labels, but neither uses transient-deficit memory or recourse. The second variant additionally uses the current arrival count A_t , which tests whether the MPA gain can be reproduced by simply exposing arrivals to a high-protection urgency-aware rule.

For these high-protection comparators, the urgency pressure is

$$U_t = \sum_{i \in Q_t} w_d(d_i(t)) w_r(r_i(t)), \quad (38)$$

where $w_d(\cdot)$ increases as the residual deadline decreases, and $w_r(\cdot)$ increases with the retry count. The generic pressure score is

$$P_t = m_B(B_t) (U_t + \alpha_Q |Q_t| + \alpha_A A_t), \quad (39)$$

where $|Q_t|$ is the visible queue length, $m_B(B_t)$ is a burst multiplier, α_Q weights visible backlog, $\alpha_A = 0$ for *High-Protection Urgency* + $Q + R$, and $\alpha_A > 0$ for the proposed *High-Protection Urgency* + $Q + R + A$ method. The selected allocation is

$$\Delta K_t = \min \left\{ \Delta K_{\max}, n_{\text{pkt}} \lceil \alpha_P P_t \rceil \right\}, \quad (40)$$

where α_P maps the pressure score to a quantized extra safety allocation with the BE floor enforced.

- *Overprovision* persistently applies the maximum admissible safety allocation. It represents the persistent prioritization extreme: low-latency traffic receives strong protection, whereas the BE service pays the corresponding opportunity cost. This is motivated by reservation- or contract-style URLLC protection mechanisms in coexistence settings [19].
- *Look-Ahead* ($H = 2$) is an internal validation reference for the one-step design of MPA. It uses the same scheduler-facing state family as MPA, including the burst state, visible queue, current arrivals, packet urgency, retry labels, and transient-deficit memory, but evaluates a deeper finite-horizon approximation. Its role is to test whether the one-step MPA design captures the main qualitative behavior of a heavier predictive reference, following the limited look-ahead dynamic programming view of sequential decision-making.
- *LSTM predictive* is a learned short-history comparator based on the standard LSTM sequence model [54]. It uses the same scheduler-visible variables as the other predictive policies, arranged over an eight-slot history, and maps them to a five-slot safety pressure target expressed in packet-sized extra safety allocations. The training target is the clipped mean excess safety pressure over the next five slots: future arrivals, pending retries, the urgency risk count, and the transient deficit are averaged, compared with the baseline packet service capacity, and clipped to the admissible extra allocation range. During evaluation, these future terms are not exposed to the policy; the trained LSTM predicts the target from the preceding eight scheduler-visible slots. One LSTM reference is trained on five Reactive stress traces, and a second uses mixed Reactive, MPA, and Overprovision traces. Both learned references are evaluated on 25 dense NLOS stress traces generated from disjoint random seeds; none of the evaluation traces are used during training.

The same admissible action set and PRB constraints are used for MPA and the comparators. The key difference is the state used in the allocation comparison: MPA includes the decayed transient-deficit state, online success calibration, and a recoverability-aware BE cost guard while evaluating the immediate action, together with a one-step carryover. The comparison set therefore covers fixed reservation, queue-triggered reaction, quantized guard control, urgency/retry scoring, stronger safety protection, finite-horizon look-ahead, and learned short-history prediction.

3.4.3. Ablation Setup

The ablation study is an internal component removal test of the MPA. It is evaluated over the first $W = 5$ slots after each detected burst onset, where timing matters most: the safety pressure has started to rise, while the visible queue may not yet fully reflect the new operating condition.

The compared variants are as follows.

- *MPA full*: Uses burst state, current arrivals, transient-deficit memory, visible queue length, packet deadline labels, retry labels, online success calibration, the recoverability-aware BE cost guard, and the one-step continuation check.

- *No online calibration*: Keeps the calibrated channel curve, transient-deficit memory, packet urgency, retry state, recoverability guard, and one-step continuation check, but disables the posterior correction from realized transmission outcomes.
- *No recoverability guard*: Keeps online success calibration and the full MPA state, but disables the recoverability multiplier and the near-ceiling candidate cap; $\rho_{\text{base}}(t)$ and $\rho_{\text{max}}(t)$ are therefore excluded from the PRB comparison.
- *No deficit*: Removes D_t while keeping arrivals, visible queue length, deadline urgency, retry information, and the one-step continuation check.
- *No deficit + no online calibration*: Removes both the transient-deficit memory and the posterior success correction, while leaving the remaining queue, deadline, retry, guard, and continuation logic unchanged.
- *No recourse*: Keeps the full scheduler-facing state, including D_t , but replaces the one-step continuation check with a myopic finite-action decision.
- *A + Q + R*: Uses arrivals, burst state, visible queue length, and retry labels, but removes packet deadline urgency, transient-deficit memory, and the one-step continuation check.
- *Deficit only*: Uses only the scalar transient-deficit memory D_t as the dynamic demand-side signal and removes arrivals, visible queue length, packet deadline urgency, retry labels, and the one-step continuation check. It serves as a diagnostic low-information variant.

The sensitivity checks also vary γ , n_{pkt} , $K_{\text{BE}}^{\text{min}}$, and the online calibration blending factor ω around the dense NLOS stress setting.

3.4.4. Metrics

The primary safety metrics are the deadline miss ratio, delivery ratio, average delivered packet delay, and 95th-percentile (p95) delivered packet delay. For a run with N_{gen} generated safety packets, N_{miss} deadline misses, and N_{del} delivered packets,

$$P_{\text{miss}} = \frac{N_{\text{miss}}}{\max(N_{\text{gen}}, 1)}, \quad P_{\text{del}} = \frac{N_{\text{del}}}{\max(N_{\text{gen}}, 1)}. \quad (41)$$

Delay is measured only for delivered safety packets, in scheduler milliseconds. The p95 delay is compared with the deadline budget in the application-level table, because a low average delay can hide a tail latency problem.

The BE metric is the normalized throughput,

$$\text{BE}_{\text{norm}} = \frac{T_{\text{BE}}}{T_{\text{BE}}^{\text{ref}}}, \quad (42)$$

where $T_{\text{BE}}^{\text{ref}}$ is the BE throughput of *Static* under the same regime, i.e., the policy that keeps the safety allocation fixed at K_{base} . This within-scenario normalization makes the BE loss comparable across the cases. The BE metric is not treated as a safety reliability metric; it measures the non-safety service capacity sacrificed to obtain a given safety outcome. The evaluation also records the average BE PRB allocation and checks the configured BE floor. To capture burst timing, onset miss rates are reported over the first 5 and 10 slots after a normal-to-burst transition. This is the interval where a queue-only response is most likely to be late.

3.4.5. Reported Artifacts

Table 3 lists the shared NR-V2X resource pool settings, Table 4 reports the success probability operating points used in the curated cases, and Table 6 identifies the policies

needed to interpret the safety–BE trade-off. The application-level delivery and delay checks are reported in the corresponding Results section.

4. Results

The results below use the parameterization given in Tables 3, 4, 6 and 7.

4.1. Safety–BE Trade-Off

The stress regime operating point is shown in Figure 1. Reactive scheduling preserves almost all BE throughput, but it leaves a large fraction of safety packets unserved before their deadline in dense NLOS stress. A fixed maximum safety allocation and high-protection urgency-aware allocation move to the opposite side of the trade-off; they reduce the miss ratio substantially but reserve far more PRBs for safety. MPA occupies the middle of this curve. In this scenario, the miss ratio is reduced from about 0.665 under *Reactive* to about 0.288 under MPA, while the normalized BE throughput remains about 0.892. The fixed maximum safety allocation reaches a lower miss ratio of approximately 0.15, but its normalized BE throughput decreases to 0.534, meaning that almost half of the static-reference non-safety service capacity is removed at the same operating point.




Figure 1. Stress regime safety–BE operating point. Error bars denote 95% confidence intervals over 25 Monte Carlo runs.

The broader metric comparison in Figure 2 shows the same reading from several angles. MPA uses average extra safety allocation of approximately 10.6 PRBs in dense NLOS stress, compared with 24 PRBs for *Overprovision* and approximately 15 PRBs for the high-protection urgency-aware comparators. Therefore, the extra safety protection is temporary and selective. In the urban nominal case, where the baseline protection is already close to adequate, MPA and *Reactive* have nearly the same BE throughput and very similar safety delivery. This behavior is useful because the memory state should change the allocation mainly when burst stress creates an unresolved demand.

4.2. Response to Burst Onset

The onset-aligned results in Figure 3 show the source of the trade-off. During the first few slots after a normal-to-burst transition, reactive scheduling still has limited evidence in the visible queue. MPA has an additional signal: recent underservice stored in the transient

deficit. This allows it to raise safety protection earlier, before the queue alone fully reflects the burst.

In the policy comparison run set for dense NLOS stress, the five-slot onset miss rate is 0.263 for *Reactive* and 0.138 for MPA. Those for the *High-Protection Urgency + Q + R* and *Overprovision* references are lower, at approximately 0.05, but they spend more safety PRBs immediately. The comparison separates the two behaviors: high-protection references minimize early misses by spending resources early, whereas MPA closes a large part of the onset gap while preserving a larger BE share than full-time protection.




Figure 2. Stress regime metric comparison across safety reliability, BE throughput, extra PRB allocation, and queue behavior. Error bars denote 95% confidence intervals over 25 Monte Carlo runs.




Figure 3. Burst onset response over the first slots after a normal-to-burst transition. Panels (a–d) show the early protection increment, visible queue ramp, early BE preservation, and first-five-slot miss rate, respectively. Panel-level error bars denote 95% confidence intervals over 25 Monte Carlo runs where applicable.

4.3. Role of the MPA Components

Figure 4 and Table 8 show which parts of MPA account for burst onset behavior in a separate matched dense NLOS stress ablation protocol. The absolute MPA entry in this table is used as the within-protocol reference, rather than as a replacement for the policy comparison value in Figure 3. Removing online success calibration slightly raises the five-slot onset miss rate from 0.138 to 0.140 and lowers the onset BE throughput from 199.3 Mbps to 197.9 Mbps. Removing the recoverability guard gives a lower onset miss rate, 0.127, but also consumes more safety capacity and lowers the onset BE throughput to 196.8 Mbps. This matched-channel onset test separates the burst memory effect from the two safeguards.

The larger onset effect comes from the timing state and the continuation check. Removing the deficit memory raises the five-slot onset miss rate from 0.138 to 0.243, and removing both the deficit memory and online calibration yields 0.247. Removing the one-step continuation raises the value to 0.228. Using arrivals, queue, and retry information without the memory state yields a higher onset miss rate of 0.326, whereas using the deficit alone yields 0.354. The full MPA also keeps the post-onset queue lower, at approximately 11.4 packets, compared with about 14.5 without deficit memory.

The full method spends more safety PRBs during the first burst slot: its mean early increment is approximately 10.8 PRBs, compared with approximately 4.3 PRBs when the memory or continuation step is removed. This lowers the BE throughput during the onset window, but it is a temporary cost tied to the stress interval. The comparison supports the intended interpretation of MPA: the memory state helps the scheduler to act before the visible queue alone has fully grown, and the continuation check keeps this action tied to near-term carryover.

Table 8. Component removal check over the first five slots after burst onset in the dense NLOS stress setting. Values are replicate means over 25 runs.

Variant	Onset Miss	Early ΔK_t	Post-Onset Queue	BE Throughput
MPA full	0.138	10.78	11.42	199.3 Mbps
No online calibration	0.140	11.02	11.60	197.9 Mbps
No recoverability guard	0.127	11.20	11.26	196.8 Mbps
No deficit memory	0.243	4.34	14.51	225.1 Mbps
No deficit + no online calibration	0.247	4.31	14.50	225.3 Mbps
No recourse	0.228	4.30	13.99	226.6 Mbps
A + Q + R	0.326	0.03	15.26	225.2 Mbps
Deficit only	0.354	0.00	15.56	224.8 Mbps

Table 9 reports two targeted checks for online calibration and the recoverability-aware guard. The first uses a success model mismatch: the planner starts from an optimistic burst success curve, while realized transmissions follow a lower-success channel. Online posterior correction gives a lower miss ratio point estimate, from 0.539 to 0.524, and an approximately 4.8 Mbps higher BE throughput point estimate. The second case places the scheduler near the recovery boundary. Maintaining the guard gives a lower miss ratio point estimate, from 0.211 to 0.204, and an approximately 4.0 Mbps higher BE throughput point estimate by avoiding workload-only safety expansion near the resource ceiling.

Table 9. Targeted checks for online success calibration and the recoverability-aware guard. Values are replicate means over 25 runs.

Case	Policy	Miss	BE Throughput	Mean ΔK_t
Success model mis-match	MPA full	0.524	190.4 Mbps	12.27
Success model mis-match	No online calibration	0.539	185.6 Mbps	13.06
Near recovery boundary	MPA full	0.204	205.2 Mbps	9.79
Near recovery boundary	No recoverability guard	0.211	201.3 Mbps	10.42




Figure 4. Component removal check during burst onset. Panels (a–d) show onset miss rate, mean early protection, mean post-onset queue, and mean onset BE throughput, respectively. Variants remove online success calibration, the recoverability-aware guard, the deficit memory, the one-step continuation, or selected scheduler-visible information. Error bars denote 95% confidence intervals over 25 Monte Carlo runs.

Table 10 reports the joint removal of online calibration and deficit memory. In the calibrated dense NLOS setting, removing online calibration alone changes the point estimate only slightly, because the realized outcomes remain close to the calibrated prior during the first few burst slots. Removing the deficit memory has a much larger effect. Removing both components gives nearly the same onset behavior as removing the deficit alone, which indicates that the early burst response in this calibrated case is driven mainly by the memory signal.

Table 10. Additional component removal check in dense NLOS stress. Values are mean $\pm$ 95% confidence-interval half-widths over 25 runs.

Variant	Onset Miss	Full-Window Miss	BE Throughput
MPA full	0.140 $\pm$ 0.002	0.292 $\pm$ 0.005	199.1 $\pm$ 2.0 Mbps
No online calibration	0.140 $\pm$ 0.004	0.294 $\pm$ 0.005	197.9 $\pm$ 1.9 Mbps
No deficit memory	0.249 $\pm$ 0.004	0.533 $\pm$ 0.008	224.6 $\pm$ 0.2 Mbps
No deficit + no online calibration	0.245 $\pm$ 0.004	0.536 $\pm$ 0.009	224.6 $\pm$ 0.3 Mbps

4.4. LSTM Learned Reference

The LSTM reference provides a learned short-history comparator under the same scheduler-layer inputs and PRB constraints. It receives eight slots of scheduler-visible history and estimates the five-slot safety pressure target defined in Section 3.4.2. The first LSTM is trained on five stress regime traces generated under the Reactive policy, giving 2650 sequence samples. A second LSTM uses mixed Reactive, MPA, and Overprovision traces, namely five from each policy, giving 7950 sequence samples. Both learned references are evaluated on 25 dense NLOS stress runs with disjoint seeds.

Table 11 presents the learned reference comparison with 95% confidence-interval half-widths for the run-level metrics. The mixed-trace LSTM gives a lower safety miss estimate than the Reactive-trained LSTM and moves close to *Overprovision* in onset behavior. This lower miss estimate comes with larger safety allocation and a smaller BE share than MPA. The comparison keeps MPA as the explicit finite-action operating point, with no training phase and a direct link to the PRB, deadline, retry, success calibration, and BE floor constraints.

Table 11. LSTM learned references in dense NLOS stress. Run-level metrics are reported as mean $\pm$ 95% confidence-interval half-widths over 25 evaluation runs.

Policy	Miss	BE Norm	Mean ΔK_t	Onset Miss
Reactive	0.661 $\pm$ 0.009	1.000 $\pm$ 0.001	0.78 $\pm$ 0.02	0.262 $\pm$ 0.003
MPA	0.294 $\pm$ 0.004	0.879 $\pm$ 0.007	11.06 $\pm$ 0.25	0.140 $\pm$ 0.003
LSTM predictive, Reactive-trained	0.178 $\pm$ 0.003	0.792 $\pm$ 0.006	13.97 $\pm$ 0.22	0.058 $\pm$ 0.002
LSTM predictive, mixed-trace	0.153 $\pm$ 0.004	0.734 $\pm$ 0.010	16.52 $\pm$ 0.36	0.055 $\pm$ 0.002
Overprovision	0.147 $\pm$ 0.003	0.534 $\pm$ 0.000	24.00 $\pm$ 0.00	0.056 $\pm$ 0.002

4.5. Application-Level Validation

Table 12 lists the delivered packet delay and BE floor behavior for the curated application cases.

This complements the miss ratio by indicating whether the delivered safety packets remain within the intended latency budget. In the urban nominal case, MPA reaches 0.999 delivery with a p95 delay of 3 ms under a 4 ms budget, and the normalized BE throughput remains at 1.000. In dense NLOS stress, MPA maintains the p95 delivered packet delay at a 3 ms budget and preserves 0.892 normalized BE throughput. The fixed maximum safety allocation reaches a lower miss ratio in the same case, but its BE throughput falls to 0.534.

Short-deadline heavy loads are closer to resource ceilings. The safety miss ratios of MPA, High-Protection Urgency + Q + R + A, Look-Ahead ($H = 2$), and Overprovision are all close, whereas the BE throughput separates them. MPA preserves 0.715 normalized BE throughput, compared with 0.460 under Overprovision. This is consistent with the analytical view: as the system moves closer to the ceiling, the safety advantage of memory narrows, but selective protection can still reduce unnecessary BE losses.

Table 12. Application-level validation for the three curated scenarios. Values are replicate means over 25 runs. BE throughput is normalized to the static allocation reference in the same scenario.

Scenario	Policy	Delivery	Miss	p95/Budget	BE Norm	BE Floor
Urban nominal	High-Protection Urgency + Q + R + A	0.999	0.001	3.0/4.0 ms	0.948	Yes
	MPA	0.999	0.001	3.0/4.0 ms	1.000	Yes
	Look-Ahead (H = 2)	0.999	0.001	3.0/4.0 ms	0.994	Yes
	Overprovision	0.999	0.001	3.0/4.0 ms	0.544	Yes
Dense NLOS stress	High-Protection Urgency + Q + R + A	0.852	0.148	3.0/3.0 ms	0.776	Yes
	MPA	0.711	0.288	3.0/3.0 ms	0.892	Yes
	Look-Ahead (H = 2)	0.823	0.177	3.0/3.0 ms	0.827	Yes
	Overprovision	0.852	0.148	3.0/3.0 ms	0.534	Yes
Short-deadline heavy load	High-Protection Urgency + Q + R + A	0.678	0.322	2.0/2.0 ms	0.647	Yes
	MPA	0.674	0.326	2.0/2.0 ms	0.715	Yes
	Look-Ahead (H = 2)	0.675	0.325	2.0/2.0 ms	0.715	Yes
	Overprovision	0.677	0.323	2.0/2.0 ms	0.460	Yes

4.6. Calibration and Regime Coverage

The calibration sweep, whose results are depicted in Figure 5, tests whether the result depends on one exact success probability curve. Under the nominal, pessimistic burst, and reduced retry gain calibrations, MPA lowers the miss ratio relative to Reactive while keeping the BE throughput below the static BE reference but well above the fixed maximum safety allocation. In the reduced retry gain case, Reactive receives less benefit from retry combining, and its curve stays close to the nominal stress case until the visible queue builds enough pressure for a stronger response. Under the mild mismatch calibration, Reactive is sufficiently effective, and the two miss ratios are indistinguishable at the reported precision. This boundary case indicates that the method is most relevant when the burst stress is sufficiently strong for memory to matter. Table 13 lists the corresponding Reactive–MPA point estimates.

Table 13. Channel/success calibration check for Reactive and MPA. Values are replicate means over 25 runs. BE throughput is normalized to the static allocation reference.

Calibration Case	Reactive Miss	MPA Miss	Reactive BE Throughput	MPA BE Throughput	Observation
Nominal	0.665	0.288	1.000	0.892	MPA lower miss
Pessimistic burst	0.719	0.535	1.000	0.852	MPA lower miss
Reduced retry gain	0.662	0.332	1.000	0.890	MPA lower miss
Mild mismatch	0.027	0.027	0.886	0.885	Reactive sufficient

Table 14 reports the actual online calibration correction during the first five slots after detected dense NLOS burst onset. The posterior correction is small in this nominal calibrated case: the mean conservative adjustment stays below 10^{-4} because the realized outcomes are consistent with the calibrated prior and only a few burst samples have been observed. Consistently, the component removal table shows only a small difference between MPA and *No online calibration* under matched calibration. The online posterior mainly protects against success model mismatch, while the transient-deficit memory carries the early timing response.




Figure 5. Calibration robustness across nominal, pessimistic burst, reduced retry gain, and mild mismatch channel/success settings. Panels (a,b) show miss ratio and normalized BE throughput across calibration settings, respectively. Error bars denote 95% confidence intervals over 25 Monte Carlo runs.

Table 14. Online calibration magnitude during the first five slots after dense NLOS burst onset. Values are replicate means over 25 runs.

Slot after Onset	Model Planning Probability	Online Planning Probability	Conservative Correction
0	0.2752	0.2752	2.3×10^{-6}
1	0.3020	0.3020	2.2×10^{-5}
2	0.3258	0.3258	3.3×10^{-5}
3	0.3443	0.3443	4.9×10^{-5}
4	0.3629	0.3628	6.1×10^{-5}

The recoverable region sweep in Figure 6 connects the results back to the ratios in Section 3.3. The sweep axes are the burst load and baseline safety allocation. The two ratio tests are baseline insufficiency, $\rho_{\text{base}} > 1$, and recovery under the admissible safety ceiling, $\rho_{\text{max}} < 1$. Panel (a) uses these tests to form the plotted classes, with a separate near-ceiling band for points close to $\rho_{\text{max}} = 1$. Panel (d) marks the exact recoverable condition, $\rho_{\text{base}} > 1$ and $\rho_{\text{max}} < 1$. The largest gains appear mainly inside this condition. When $\rho_{\text{base}} \leq 1$, the baseline service is already sufficient, and memory adds little. When $\rho_{\text{max}} \geq 1$, all policies are constrained by the same resource limit.

4.7. Robustness, Sensitivity, and Complexity Analysis

Table 15 adds standard deviations to the main dense NLOS stress comparison. The separation between *Reactive* and MPA is much larger than the confidence intervals. The same table also shows the intended trade-off: *Overprovision* reaches the lowest miss ratio, while MPA preserves a larger BE share.

Table 16 checks two asymmetric traffic settings. When BE traffic is heavier than safety traffic, MPA reduces the safety miss ratio relative to *Reactive* while preserving substantially more BE throughput than *Overprovision*. When safety traffic is heavier than BE traffic, the BE queue is not the bottleneck and all BE-normalized values stay close to one; the safety ordering is still consistent with the dense NLOS interpretation.




Figure 6. Recoverable region sweep over burst load, baseline safety allocation, burst persistence, and deadline. Each sweep point is averaged over 25 runs. Panel (a) maps each operating point to the plotted ratio test classes, including a near-ceiling band around the maximum allocation boundary. Panel (d) marks the exact recoverable condition, $\rho_{\text{base}} > 1$ and $\rho_{\text{max}} < 1$. Panels (b,c) show the corresponding safety gains over *Reactive* and BE preservation over *Overprovision*.

Table 15. Dense NLOS stress replication statistics. Values are means, standard deviations (SD), and 95% confidence-interval half-widths (CI) over 25 runs.

Policy	Miss Mean	Miss SD	Miss CI	BE Norm Mean	BE Norm SD	BE Norm CI
Static	0.712	0.014	0.006	0.999	0.002	0.001
Reactive	0.665	0.018	0.007	1.000	0.003	0.001
MPA	0.288	0.010	0.004	0.892	0.018	0.008
Overprovision	0.149	0.007	0.003	0.534	0.000	0.000

Table 16. Asymmetric load checks. Values are mean $\pm$ 95% confidence-interval half-widths over 25 runs.

Case	λ_B	λ_{BE}	Policy	Miss	BE Norm
BE-heavy asymmetric	0.35	0.90	Static	0.482 ± 0.010	1.000 ± 0.000
	0.35	0.90	Reactive	0.458 ± 0.009	0.989 ± 0.000
	0.35	0.90	MPA	0.233 ± 0.007	0.919 ± 0.003
	0.35	0.90	Overprovision	0.091 ± 0.003	0.455 ± 0.000
Safety-heavy asymmetric	0.90	0.45	Static	0.757 ± 0.007	1.000 ± 0.002
	0.90	0.45	Reactive	0.702 ± 0.011	1.000 ± 0.002
	0.90	0.45	MPA	0.327 ± 0.005	1.000 ± 0.002
	0.90	0.45	Overprovision	0.278 ± 0.005	1.001 ± 0.002

Table 17 varies the memory decay, packet PRB demand, BE floor, and online calibration blending factor around the dense NLOS stress setting. A smaller decay factor gives a more safety-heavy response in this setting; a larger packet demand increases the PRB cost of each attempt; and the BE floor and blending factor sweeps do not change the qualitative ordering.

Table 17. Parameter sensitivity around dense NLOS stress. Values are mean $\pm$ 95% confidence-interval half-widths over 25 runs.

Parameter	Value	Miss	BE Norm	Mean ΔK_t
γ	0.50	0.148 ± 0.003	0.818 ± 0.010	13.39
	0.70	0.292 ± 0.005	0.883 ± 0.008	10.92
	0.90	0.290 ± 0.005	0.885 ± 0.008	10.84
n_{pkt}	1	0.365 ± 0.007	1.000 ± 0.001	3.95
	2	0.291 ± 0.005	0.885 ± 0.010	10.87
	3	0.310 ± 0.008	0.753 ± 0.012	15.68
$K_{\text{BE}}^{\text{min}}$	6	0.294 ± 0.004	0.882 ± 0.007	10.98
	10	0.293 ± 0.004	0.882 ± 0.007	10.96
	14	0.292 ± 0.005	0.883 ± 0.008	10.92
ω	0.00	0.288 ± 0.003	0.893 ± 0.005	10.55
	0.35	0.293 ± 0.004	0.881 ± 0.007	11.01
	0.70	0.287 ± 0.004	0.891 ± 0.008	10.63

Table 18 reports prototype execution measurements together with finite enumeration counts. The timing columns compare policies within the same scripted implementation environment; the evaluation counts and recourse depth provide a more portable description of the online decision size. MPA remains between a simple reactive rule and the two-step look-ahead reference.

Table 18. Prototype execution measurements and finite enumeration counts for one scheduler decision. Timing values are measured in the same MATLAB implementation, while candidate evaluation counts and recourse depths describe the implementation-independent size of the online decision. Values are mean $\pm$ 95% confidence-interval half-widths over 25 runs.

Policy	Mean Decision Time (ms)	p95 Decision Time (ms)	Mean Candidate Evaluations	Recourse Depth
Reactive	0.119 ± 0.003	0.153 ± 0.010	1 ± 0	0
MPA	0.862 ± 0.013	1.274 ± 0.018	1776 ± 46	1
Look-Ahead (H = 2)	8.704 ± 0.168	13.163 ± 0.369	$13,264 \pm 276$	2
Overprovision	0.003 ± 0.000	0.004 ± 0.000	1 ± 0	0

5. Discussion

The results point to a clear interpretation: burst stress is a timing problem as much as a load problem. At the beginning of a burst, the visible queue can still appear moderate, even though new arrivals and failed attempts have already created work for the next few slots. This explains why Reactive behaves well in urban nominal traffic but falls behind in dense NLOS stress. In the dense NLOS run set, Reactive has a miss ratio of 0.665, whereas MPA reduces it to 0.288 while keeping the normalized BE throughput at 0.892. The gain comes from acting before the visible queue alone becomes a strong signal.

The same dense NLOS case also shows the safety-BE frontier. *High-Protection Urgency* + *Q* + *R* + *A* and *Overprovision* reduce the miss ratio to approximately 0.148, and the mixed-trace LSTM reference provides similar safety-heavy behavior with a miss ratio of 0.153. These points spend more safety PRBs. *Overprovision* leaves 0.534 normalized BE throughput, and the mixed-trace LSTM leaves 0.734. MPA sits between *Reactive* and these safety-heavy references: it does not minimize misses at any BE cost, but it closes a large part of the Reactive safety gap while preserving a larger BE share. This operating point is relevant because BE traffic represents non-safety data services that still share the same radio

resources, such as infotainment, map updates, sensor sharing, and background application traffic.

The recoverable region sweep explains when this trade-off appears. If the baseline allocation is already sufficient, there is little unresolved demand for memory to expose. If the maximum admissible safety allocation is also insufficient, all schedulers are constrained by the same PRB ceiling. Between these cases, the baseline service is not sufficient, but temporary PRB headroom remains available. In this region, the deficit memory encourages earlier safety protection, the online success correction adapts the planning probability to observed outcomes, and the recoverability guard reduces expansion when extra PRBs have weak expected returns. The asymmetric load checks support the same reading: under a BE-heavy load, MPA preserves much more BE throughput than *Overprovision*; under a safety-heavy load, BE is no longer the bottleneck, and the main separation is safety delivery.

The transient-deficit state is similar to virtual queues, debt variables, and delay-deficit ideas in stochastic scheduling. The difference is in how the state is used. Classical Lyapunov drift-plus-penalty methods usually build a control rule around long-term stability or constraint satisfaction. Here, the deficit is a short-lived scheduler signal inside a bounded PRB enumeration, combined with packet deadlines, retry labels, online-calibrated success probabilities, a BE floor, and a one-step continuation check. The drift bound in Section 3.3 explains why the memory decays after the burst; the recoverable region ratios identify the operating points where this short memory is likely to help. This gives the method a queuing-theoretic basis without turning the paper into a throughput-optimality proof.

The NR-V2X mapping provides the radio interpretation of these findings. In a Mode 1-style setting, MPA can be read as a gNB-side PRB partition rule. In Mode 2, autonomous resource selection first determines the usable resource set through sensing, SCI decoding, pre-emption, IUC, re-evaluation, and local implementation rules; the scheduler-layer rule then partitions the usable set between safety and BE traffic. Hidden-node effects, sensing errors, SCI/PSFCH overhead, directional transmission, and geometry-driven propagation therefore enter the scheduler as a smaller admissible PRB budget or as more conservative packet success probabilities. Both effects increase the safety pressure, which is exactly the situation in which recent service shortfall becomes informative.

The comparison set also clarifies the position of the method relative to recent work. SOVANET+ addresses critical and non-critical V2X scheduling under RAN load and channel quality information, while SPS, re-evaluation, pre-emption, IUC, and Mode 2 learning methods address resource selection and collision mitigation at the sidelink procedure level [12,13,38,39,48]. MPC, RNN/LSTM, and DRL approaches can use richer prediction or training, but they also introduce model, data, and deployment interface requirements. MPA keeps the scheduler-layer rule explicit, with a finite action set, direct PRB constraints, online success correction, a recoverability guard, and no offline training phase. The complexity table reflects this position; the prototype MPA decision is heavier than in Reactive but much lighter than in the two-step look-ahead reference.

Limitations and Future Work

The results are most informative in recoverable operating points, where $\rho_{\text{base}} > 1$ and $\rho_{\text{max}} < 1$. In this region, the baseline safety allocation falls behind, but the admissible safety ceiling still provides sufficient PRB headroom for temporary recovery. When $\rho_{\text{max}} \geq 1$, the operating point is different: the traffic load, deadline pressure, channel quality, and PRB ceiling dominate the outcome. A scheduler can still choose a better allocation among the available actions, but the main bottleneck is then the resource budget itself. The higher miss ratios in dense NLOS stress should therefore be interpreted as constrained-regime results, not as a system-level URLLC reliability claim.

The Mode 2 interpretation depends on how accurately the usable resource set is represented. Sensing errors, hidden nodes, SCI decoding errors, directional-beam mismatch, and feedback-channel limits can reduce the resources that remain available during burst periods. In this scheduler-layer model, such effects enter through the effective PRB budget and planning success probability. A lower-layer sidelink implementation would need to make this link explicit by coupling MPA with the resource selection window, reselection timing, SCI/PSFCH overhead, Release 17/Release 18 pre-emption and IUC behavior, and geometry-dependent PRR behavior.

Scalability is mainly tied to the number of packet states propagated at each slot and to the finite action set. The action set remains bounded by the PRB headroom and the packet size granularity; however, much larger UE populations would require aggregation or priority grouping before the slot-level decision. A practical next step is a multi-level scheduler in which one layer groups vehicles or service classes under sustained overload, while MPA handles the slot-level safety-BE split during recoverable bursts. This direction is related to medium-term scheduling for resource thrashing, where an intermediate control layer smooths overload before short-term decisions become unstable [57]. Release 18 sidelink coordination, multi-beam operation, and joint safety-BE QoS slicing are also suitable directions for extending the present scheduler-layer model.

6. Conclusions

This study examined scheduler-layer PRB allocation for safety and BE coexistence under burst stress. The main challenge is timing. At burst onset, the queue may still appear manageable, while arrivals and failed attempts have already created a demand for the next decisions. A queue-only response can therefore arrive late, whereas a fixed maximum safety allocation protects early at a high BE cost.

MPA addresses this timing problem with a short memory of recent unresolved safety demands, an online conservative success update, and a recoverability-aware BE cost guard. When arrivals, failed attempts, and delivered packets indicate that the baseline safety allocation is falling behind, MPA can raise the safety PRB allocation before the queue alone provides a strong signal. When the pressure falls, the allocation relaxes and returns the capacity to the BE traffic.

The results show that this mechanism is most useful in recoverable burst stress regimes. In dense NLOS stress, MPA lowers the miss ratio relative to Reactive while preserving substantially more BE throughput than a fixed maximum safety allocation. Safety-heavy, learned, and deeper look-ahead references can reduce misses further, but they require more safety capacity. Application-level validation also shows that delivered safety packets remain within the p95 delay budget in urban nominal, dense NLOS stress, and short-deadline heavy-load scenarios.

The main conclusion is that recent service shortfall is a useful scheduler-layer signal for deciding when temporary safety protection is worth its BE cost. Under calibrated NR-V2X resource pool and success probability assumptions, this signal provides a practical way to act in the recoverable region between queue-reactive delays and permanent safety reservation. Outside this region, the resource ceiling or the lower-layer sidelink resource set becomes the dominant limitation, which is why the method is best read as a scheduler-layer allocation rule rather than a complete sidelink reliability solution.

Supplementary Materials: The following supporting information can be downloaded at: <https://www.mdpi.com/article/10.3390/electronics15142995/s1>, Dataset S1: minimal dataset containing the simulation summary values used for the reported tables and figures.

Author Contributions: Conceptualization, N.Y. and B.Ç.; methodology, N.Y.; software, N.Y.; validation, N.Y. and B.Ç.; formal analysis, N.Y.; investigation, N.Y.; data curation, N.Y.; writing—original draft preparation, N.Y.; writing—review and editing, N.Y. and B.Ç.; visualization, N.Y.; supervision, B.Ç.; project administration, B.C.; funding acquisition, B.Ç. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: No external datasets were used in this study. The original contributions presented in this study are included in the article and in Supplementary Dataset S1. Further inquiries can be directed to the corresponding author.

Acknowledgments: This study was carried out within the scope of Atatürk University Scientific Research Projects Coordination Unit (BAP) project FDK-2023-11959.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript.

3GPP	Third-Generation Partnership Project
BE	Best effort
BLER	Block error rate
BWP	Bandwidth part
C-V2X	Cellular Vehicle-to-Everything
dB	Decibel
DENM	Decentralized Environmental Notification Message
DRL	Deep reinforcement learning
eMBB	Enhanced mobile broadband
GHz	Gigahertz
gNB	Next-generation Node B
HARQ	Hybrid automatic repeat request
IEEE	Institute of Electrical and Electronics Engineers
IUC	Inter-UE coordination
kHz	Kilohertz
LTE-V2X	Long-Term Evolution Vehicle-to-Everything
LSTM	Long short-term memory
MAC	Medium access control
MCS	Modulation and coding scheme
MHz	Megahertz
MPC	Model predictive control
ms	Millisecond
MPA	Memory-Based Predictive Allocation
NLOS	Non-line-of-sight
NR	New Radio
NR-V2X	New Radio Vehicle-to-Everything
p95	95th percentile
PDU	Protocol data unit
PHY	Physical layer
PRB	Physical resource block
PRR	Packet reception ratio

PSFCH	Physical sidelink feedback channel
QoS	Quality of service
RAN	Radio access network
RIS	Reconfigurable intelligent surface
RNN	Recurrent neural network
RRC	Radio resource control
SCS	Subcarrier spacing
SCI	Sidelink control information
SINR	Signal-to-interference-plus-noise ratio
SOVANET+	Adaptive extension of Scheduler Optimized for Vehicular Ad Hoc Network
SPS	Semi-persistent scheduling
UE	User equipment
URLLC	Ultra-reliable low-latency communication
V2X	Vehicle-to-Everything

References

1. Popovski, P.; Nielsen, J.J.; Stefanovic, C.; de Carvalho, E.; Strom, E.; Trillingsgaard, K.F.; Bana, A.-S.; Kim, D.M.; Kotaba, R.; Park, J.; Sorensen, R.B. Wireless Access for Ultra-Reliable Low-Latency Communication: Principles and Building Blocks. *IEEE Netw.* **2018**, *32*, 16–23. <https://doi.org/10.1109/MNET.2018.1700258>.
2. Bennis, M.; Debbah, M.; Poor, H.V. Ultrareliable and Low-Latency Wireless Communication: Tail, Risk, and Scale. *Proc. IEEE* **2018**, *106*, 1834–1853. <https://doi.org/10.1109/JPROC.2018.2867029>.
3. She, C.; Sun, C.; Gu, Z.; Li, Y.; Yang, C.; Poor, H.V.; Vucetic, B. A Tutorial on Ultrareliable and Low-Latency Communications in 6G: Integrating Domain Knowledge into Deep Learning. *Proc. IEEE* **2021**, *109*, 204–246. <https://doi.org/10.1109/JPROC.2021.3053601>.
4. 3GPP. *Study on NR Vehicle-to-Everything (V2X)*; 3GPP TR 38.885, V16.0.0, Release 16; 3rd Generation Partnership Project: Valbonne, France, 2019.
5. 3GPP. *NR; Physical Layer Procedures for Data*; 3GPP TS 38.214, V15.9.0, Release 15; ETSI TS 138 214; European Telecommunications Standards Institute: Sophia Antipolis, France, 2020.
6. Anand, A.; De Veciana, G.; Shakkottai, S. Joint Scheduling of URLLC and eMBB Traffic in 5G Wireless Networks. In Proceedings of the IEEE INFOCOM 2018—IEEE Conference on Computer Communications, Honolulu, HI, USA, 15–19 April 2018; pp. 1970–1978. <https://doi.org/10.1109/INFOCOM.2018.8486430>.
7. Korrai, P.; Lagunas, E.; Sharma, S.K.; Chatzinotas, S.; Bandi, A.; Ottersten, B. A RAN Resource Slicing Mechanism for Multiplexing of eMBB and URLLC Services in OFDMA Based 5G Wireless Networks. *IEEE Access* **2020**, *8*, 45674–45688. <https://doi.org/10.1109/ACCESS.2020.2977773>.
8. Feng, L.; Zi, Y.; Li, W.; Zhou, F.; Yu, P.; Kadoch, M. Dynamic Resource Allocation with RAN Slicing and Scheduling for uRLLC and eMBB Hybrid Services. *IEEE Access* **2020**, *8*, 34538–34551. <https://doi.org/10.1109/ACCESS.2020.2974812>.
9. Liu, P.; Shen, C.; Liu, C.; Cintrón, F.J.; Zhang, L.; Cao, L.; Rouil, R.; Roy, S. Towards 5G New Radio Sidelink Communications: A Versatile Link-Level Simulator and Performance Evaluation. *Comput. Commun.* **2023**, *208*, 231–243. <https://doi.org/10.1016/j.comcom.2023.06.005>.
10. Aslam, S.; Khan, D.; Chang, K. The Shared Experience Actor–Critic (SEAC) Approach for Allocating Radio Resources and Mitigating Resource Collisions in 5G-NR-V2X Mode 2 under Aperiodic Traffic Conditions. *Sensors* **2024**, *24*, 6769. <https://doi.org/10.3390/s24206769>.
11. Ali, Z.; Lagén, S.; Giupponi, L.; Rouil, R. 3GPP NR V2X Mode 2: Overview, Models and System-Level Evaluation. *IEEE Access* **2021**, *9*, 89554–89579. <https://doi.org/10.1109/ACCESS.2021.3090855>.
12. Lusvarghi, L.; Molina-Galan, A.; Coll-Perales, B.; Gozalvez, J.; Merani, M.L. A Comparative Analysis of the Semi-Persistent and Dynamic Scheduling Schemes in NR-V2X Mode 2. *Veh. Commun.* **2023**, *42*, 100628. <https://doi.org/10.1016/j.vehcom.2023.100628>.
13. Molina-Galan, A.; Lusvarghi, L.; Coll-Perales, B.; Gozalvez, J.; Merani, M.L. On the Impact of Re-Evaluation in 5G NR V2X Mode 2. *IEEE Trans. Veh. Technol.* **2024**, *73*, 2669–2683. <https://doi.org/10.1109/TVT.2023.3318235>.
14. Tassioulas, L.; Ephremides, A. Stability Properties of Constrained Queueing Systems and Scheduling Policies for Maximum Throughput in Multihop Radio Networks. *IEEE Trans. Autom. Control* **1992**, *37*, 1936–1948. <https://doi.org/10.1109/9.182479>.
15. Neely, M.J. *Stochastic Network Optimization with Application to Communication and Queueing Systems*; Morgan & Claypool Publishers: San Rafael, CA, USA, 2010.
16. Zhang, W.; Derakhshani, M.; Lambbotharan, S. Stochastic Optimization of URLLC-eMBB Joint Scheduling with Queueing Mechanism. *IEEE Wirel. Commun. Lett.* **2021**, *10*, 844–848. <https://doi.org/10.1109/LWC.2020.3046628>.

17. Liu, C.L.; Layland, J.W. Scheduling Algorithms for Multiprogramming in a Hard-Real-Time Environment. *J. ACM* **1973**, *20*, 46–61. <https://doi.org/10.1145/321738.321743>.
18. Hou, I.-H.; Borkar, V.; Kumar, P.R. A Theory of QoS for Wireless. In Proceedings of the IEEE INFOCOM 2009, Rio de Janeiro, Brazil, 19–25 April 2009; pp. 486–494. <https://doi.org/10.1109/INFCOM.2009.5061954>.
19. Manzoor, A.; Kazmi, S.M.A.; Pandey, S.R.; Hong, C.S. Contract-Based Scheduling of URLLC Packets in Incumbent eMBB Traffic. *IEEE Access* **2020**, *8*, 167516–167526. <https://doi.org/10.1109/ACCESS.2020.3023128>.
20. 3GPP. *Study on Evaluation Methodology of New Vehicle-to-Everything (V2X) Use Cases for LTE and NR*; 3GPP TR 37.885, V15.3.0, Release 15; 3rd Generation Partnership Project: Valbonne, France, 2019.
21. 3GPP. *NR; Medium Access Control (MAC) Protocol Specification*; 3GPP TS 38.321, V19.2.0, Release 19; 3rd Generation Partnership Project: Valbonne, France, 2026.
22. 3GPP. *NR; Radio Resource Control (RRC) Protocol Specification*; 3GPP TS 38.331, V19.2.0, Release 19; 3rd Generation Partnership Project: Valbonne, France, 2026.
23. 3GPP. *Architecture Enhancements for 5G System (5GS) to Support Vehicle-to-Everything (V2X) Services*; 3GPP TS 23.287, V19.1.0, Release 19; 3rd Generation Partnership Project: Valbonne, France, 2026.
24. 3GPP. *Release 17*. Available online: <https://www.3gpp.org/specifications-technologies/releases/release-17> (accessed on 5 June 2026).
25. 3GPP. *Release 18*. Available online: <https://www.3gpp.org/specifications-technologies/releases/release-18> (accessed on 5 June 2026).
26. Molina-Masegosa, R.; Gozalvez, J. LTE-V for Sidelink 5G V2X Vehicular Communications: A New 5G Technology for Short-Range Vehicle-to-Everything Communications. *IEEE Veh. Technol. Mag.* **2017**, *12*, 30–39. <https://doi.org/10.1109/MVT.2017.2752798>.
27. Garcia, M.H.C.; Molina-Galan, A.; Boban, M.; Gozalvez, J.; Coll-Perales, B.; Şahin, T.; Kousaridas, A. A Tutorial on 5G NR V2X Communications. *IEEE Commun. Surv. Tutor.* **2021**, *23*, 1972–2026. <https://doi.org/10.1109/COMST.2021.3057017>.
28. Lien, S.-Y.; Deng, D.-J.; Lin, C.-C.; Tsai, H.-L.; Chen, T.; Guo, C.; Cheng, S.-M. 3GPP NR Sidelink Transmissions Toward 5G V2X. *IEEE Access* **2020**, *8*, 35368–35382. <https://doi.org/10.1109/ACCESS.2020.2973706>.
29. Bagheri, H.; Noor-A-Rahim, M.; Liu, Z.; Lee, H.; Pesch, D.; Moessner, K.; Xiao, P. 5G NR-V2X: Toward Connected and Cooperative Autonomous Driving. *IEEE Commun. Stand. Mag.* **2021**, *5*, 48–54. <https://doi.org/10.1109/MCOMSTD.001.2000069>.
30. Naik, G.; Choudhury, B.; Park, J.-M. IEEE 802.11bd & 5G NR V2X: Evolution of Radio Access Technologies for V2X Communications. *IEEE Access* **2019**, *7*, 70169–70184. <https://doi.org/10.1109/ACCESS.2019.2919489>.
31. Anwar, W.; Franchi, N.; Fettweis, G. Physical Layer Evaluation of V2X Communications Technologies: 5G NR-V2X, LTE-V2X, IEEE 802.11bd, and IEEE 802.11p. In Proceedings of the 2019 IEEE 90th Vehicular Technology Conference (VTC2019-Fall), Honolulu, HI, USA, 22–25 September 2019; pp. 1–7. <https://doi.org/10.1109/VTCFall.2019.8891313>.
32. Wu, Z.; Bartoletti, S.; Martinez, V.; Bazzi, A. A Methodology for Abstracting the Physical Layer of Direct V2X Communications Technologies. *Sensors* **2022**, *22*, 9330. <https://doi.org/10.3390/s22239330>.
33. Yin, J.; Hwang, S.-H. Reuse Distance-Aided Resource Selection Mechanisms for NR-V2X Sidelink Communication. *Sensors* **2024**, *24*, 253. <https://doi.org/10.3390/s24010253>.
34. Sabeeh, S.; Wołosowski, K. On Adaptation of Resources in New Radio Vehicle-to-Everything Mode 1 Dynamic Resource Allocation. *Electronics* **2025**, *14*, 77. <https://doi.org/10.3390/electronics14010077>.
35. Tabassum, M.; Oliveira, A. Collision Probabilities Between User Equipment Using 5G NR Sidelink Time-Domain-Based Resource Allocation in C-V2X. *Electronics* **2025**, *14*, 751. <https://doi.org/10.3390/electronics14040751>.
36. Guerra, J.; Luís, M.; Rito, P. Performance Evaluation of 5G New Radio V2X Sidelink for Coexisting Traffic. *IEEE Access* **2025**, *13*, 131400–131410. <https://doi.org/10.1109/ACCESS.2025.3590868>.
37. Traspadini, A.; Deshpande, A.A.; Giordani, M.; Mahabal, C.; Shimizu, T.; Zorzi, M. Sensing-Based Beamformed Resource Allocation in Standalone Millimeter-Wave Vehicular Networks. In Proceedings of the ICC 2025–IEEE International Conference on Communications, Montreal, QC, Canada, 8–12 June 2025; pp. 3088–3094. <https://doi.org/10.1109/ICC52391.2025.11161005>.
38. Li, Z.; Wang, P.; Shen, Y.; Li, S. Reinforcement Learning-Based Resource Allocation Scheme of NR-V2X Sidelink for Joint Communication and Sensing. *Sensors* **2025**, *25*, 302. <https://doi.org/10.3390/s25020302>.
39. Saad, M.M.; Tariq, M.A.; Ajmal, M.; Kim, D.; Srivastava, G. Federated Multiagent Reinforcement Learning for Resource Allocation in NR-V2X Mode 2. *IEEE Internet Things J.* **2025**, *12*, 23402–23417. <https://doi.org/10.1109/JIOT.2025.3555195>.
40. Li, Z.; Li, G.; Wu, Z.; Zhang, W.; Bazzi, A. Dynamic Allocation of C-V2X Communication Resources Based on Graph Attention Network and Deep Reinforcement Learning. *Sensors* **2025**, *25*, 5209. <https://doi.org/10.3390/s25165209>.
41. Sun, M.; Xu, J.; Wang, J. State-Aware Resource Allocation for V2X Communications. *Sensors* **2026**, *26*, 344. <https://doi.org/10.3390/s26010344>.
42. ETSI. *Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Part 3: Specifications of Decentralized Environmental Notification Basic Service*; ETSI EN 302 637-3, V1.3.1; European Telecommunications Standards Institute: Sophia Antipolis, France, 2019.

43. Parvez, I.; Rahmati, A.; Guvenc, I.; Sarwat, A.I.; Dai, H. A Survey on Low Latency Towards 5G: RAN, Core Network and Caching Solutions. *IEEE Commun. Surv. Tutor.* **2018**, *20*, 3098–3130. <https://doi.org/10.1109/COMST.2018.2841349>.
44. Yin, H.; Zhang, L.; Roy, S. Multiplexing URLLC Traffic within eMBB Services in 5G NR: Fair Scheduling. *IEEE Trans. Commun.* **2021**, *69*, 1080–1093. <https://doi.org/10.1109/TCOMM.2020.3035582>.
45. Huang, Y.; Li, S.; Li, C.; Hou, Y.T.; Lou, W. A Deep-Reinforcement-Learning-Based Approach to Dynamic eMBB/URLLC Multiplexing in 5G NR. *IEEE Internet Things J.* **2020**, *7*, 6439–6456. <https://doi.org/10.1109/JIOT.2020.2978692>.
46. Peng, H.; Wang, L.-C.; Jian, Z. Data-Driven Spectrum Partition for Multiplexing URLLC and eMBB. *IEEE Trans. Cogn. Commun. Netw.* **2023**, *9*, 386–397. <https://doi.org/10.1109/TCCN.2022.3231690>.
47. Shen, X.; Zeng, Z.; Liu, X. RIS-Assisted Network Slicing Resource Optimization Algorithm for Coexistence of eMBB and URLLC. *Electronics* **2022**, *11*, 2575. <https://doi.org/10.3390/electronics11162575>.
48. Kanavos, A.; Kaloxylos, A. V2X Communications in Highway Environments: Scheduling Challenges and Solutions for 6G Networks. *Telecom* **2025**, *6*, 13. <https://doi.org/10.3390/telecom6010013>.
49. Liu, J.; Chen, Y.; Wu, W.; Tian, F. Resource-Adaptive Semantic Transmission and Client Scheduling for OFDM-Based V2X Communications. *Sensors* **2026**, *26*, 2615. <https://doi.org/10.3390/s26092615>.
50. Jung, S.; Kim, J.-H.; Kim, J. Intelligent Extra Resource Allocation for Cooperative Awareness Message Broadcasting in Cellular-V2X Networks. *IEEE Trans. Netw. Serv. Manag.* **2025**, *22*, 1677–1689. <https://doi.org/10.1109/TNSM.2024.3496394>.
51. Jones, N.; Modiano, E. Optimal Slicing and Scheduling with Service Guarantees in Multi-Hop Wireless Networks. In Proceedings of the Twenty-Fifth International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing, Athens, Greece, 14–17 October 2024; pp. 181–190. <https://doi.org/10.1145/3641512.3686385>.
52. Bertsekas, D.P. *Dynamic Programming and Optimal Control, Volume I*, 4th ed.; Athena Scientific: Belmont, MA, USA, 2012.
53. Rawlings, J.B.; Mayne, D.Q.; Diehl, M.M. *Model Predictive Control: Theory, Computation, and Design*, 2nd ed.; Nob Hill Publishing: Madison, WI, USA, 2020.
54. Hochreiter, S.; Schmidhuber, J. Long Short-Term Memory. *Neural Comput.* **1997**, *9*, 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>.
55. Yuan, L.; Wen, W.; Wu, P.; Xia, M. Energy-Efficient CPM Modulation for IoT: A Joint Index and Spread Spectrum Approach. In Proceedings of the 2025 IEEE/CIC International Conference on Communications in China Workshops (ICCC Workshops), Shanghai, China, 10–13 August 2025; pp. 1–6. <https://doi.org/10.1109/ICCCWorkshops67136.2025.11148219>.
56. Ma, H.; Tao, Y.; Fang, Y.; Chen, P.; Li, Y. Multi-Carrier Initial-Condition-Index-Aided DCSK Scheme: An Efficient Solution for Multipath Fading Channel. *IEEE Trans. Veh. Technol.* **2025**, *74*, 15743–15757. <https://doi.org/10.1109/TVT.2025.3567450>.
57. Reuven, M.; Wiseman, Y. Medium-Term Scheduler as a Solution for the Thrashing Effect. *Comput. J.* **2006**, *49*, 297–309. <https://doi.org/10.1093/comjnl/bxl001>.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.