

Article

Advanced Adaptive Scheduling for Autonomous Driving in Beyond-5G/6G Networks

Athanasios Kanavos ^{1,*}, Gerasimos Papanikolaou-Ntais ² and Alexandros Kaloxylos ¹¹ Department of Informatics and Telecommunications, University of Peloponnese, 221 31 Tripoli, Greece² Department of Informatics, University of Piraeus, 185 34 Piraeus, Greece

* Correspondence: kanavosa@uop.gr

Abstract

The shift toward autonomous driving is steadily reducing the need for human intervention in vehicular operations, but sustaining this shift depends on 6G networks that can reliably multiplex diverse, delay-critical services under deterministic latency, throughput, and reliability constraints. Static over-provisioning can meet the stringent Quality of Service (QoS) requirements of critical traffic, but at the cost of resource starvation for co-existing non-critical services, and it degrades further under adverse channel conditions or network congestion where spectrum must be used efficiently. This motivates dynamic, service-aware scheduling as a core requirement for multi-service 6G vehicular architectures. This paper presents SOVANET+, an extended scheduling technique that jointly accounts for service criticality (critical vs. non-critical), network load, and wireless link quality to allocate resources adaptively across coexisting Vehicle-to-Everything (V2X) services. We evaluate SOVANET+ through extensive simulations of a congested single-cell urban-grid deployment supporting delay-critical automated driving services. Results show that SOVANET+ achieves lower latency and jitter, higher throughput, and improved uplink and downlink reliability compared to existing scheduling approaches, while scaling effectively to large numbers of connected autonomous vehicles, supporting its viability for next-generation intelligent transportation systems.

Keywords: beyond-5G; 6G; V2N use cases; resource allocation; scheduling

1. Introduction

1.1. Motivation and Background

Modern transport challenges, including traffic congestion, safety, and environmental sustainability, can be addressed with the support of information and communication technologies. Intelligent Transport Systems (ITS) [1] address a broad set of communication-related Vehicle-to-Everything (V2X) use cases (UCs). Their goals include safer travel, reduced environmental impact, improved traffic management, and increased transportation efficiency. In recent years, intelligent-vehicle research has shifted towards Cellular V2X (C-V2X), which allows a vehicle to exchange data with other vehicles, the infrastructure, pedestrians, and the network itself. As it relies on cellular access, C-V2X offers low-latency, high-data-rate transfer, thereby improving road safety and traffic efficiency and supporting services such as autonomous driving and smart-city operations.

Academic Editors: Cornel Ioana, Angela Digulescu, Annamaria Sarbu, Emanuel Puschita and Cristina Despina-Stoian

Received: 25 June 2026

Revised: 3 August 2026

Accepted: 4 August 2026

Published: 5 August 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.

A comprehensive and forward-thinking roadmap has been developed by the 5G Automotive Association (5GAA) [2], setting out specific timelines for the wide-scale deployment of advanced driving UCs. Its goal is to give automotive and telecommunications stakeholders clear milestones and objectives for introducing advanced communication technologies into vehicles. By defining the successive deployment phases, the association aims to promote cooperation among manufacturers, regulators, and technology providers to improve road safety, traffic efficiency, and the overall driving experience through new connected-vehicle solutions. The 5GAA report in [3] and [4], respectively, the spectrum required by day-1 and by advanced V2X UCs for ITS when supported by Long-Term Evolution (LTE)-V2X and by 5G New Radio (NR)-V2X cellular networks. Report [3] defines the bandwidth required by V2X UCs both for direct communications over the C-V2X PC5 interface, in the globally harmonized 5.9 GHz ITS band, and for network-based communications over the C-V2X Uu interface. Studies mentioned above indicate the need for well-designed scheduling mechanisms, as the spectrum available to mobile network operators is limited. It has been quantified that “at least 500 MHz of additional service-agnostic mid-band (1 to 7 GHz) spectrum is needed to provide high-capacity, citywide advanced automotive Vehicle-to-Network (V2N) services”.

Additionally, the Third Generation Partnership Project (3GPP) identifies UCs and potential service requirements in [5] to strengthen 3GPP’s support for V2X services. The work in [6] provides a comprehensive analysis and a harmonized view of the performance requirements for advanced V2X applications, as provided by 3GPP and 5GAA. The paper thoroughly analyzes the key requirements of these services, taking into account the characteristics and capabilities of both LTE and 5G NR networks. The analysis indicates that current cellular networks can already support several UCs, but they cannot accommodate the safety-critical and highly demanding requirements of fully autonomous driving once the density of such services in a given area exceeds a certain level. Adding bandwidth is the main remedy; however, as spectrum is limited, suitable scheduling schemes are required to protect the performance of the safety-critical services. Moreover, [7] explicitly states that network optimization and efficiency in 6G networks are achieved through dynamic resource allocation, which enables the management of bandwidth, computing power, and energy according to real-time demand.

1.2. Proposed Approach and Contributions

This study concentrates on performance optimization through adaptive resource allocation, a fundamental aspect of coordinating access in modern cellular networks. Scheduling algorithms determine which users should receive radio resources during each time slot and the amount of bandwidth assigned to them. These decisions often aim to optimize key performance indicators, such as maximizing total data rate or minimizing delays. In the context of 6G V2X networks, scheduling solutions must deliver high throughput and capacity to accommodate multiple users simultaneously, each with varying requirements. V2X UCs impose diverse quality of service (QoS) demands. Network optimization and efficiency are critical for time-sensitive driving applications, as they face strict delay constraints and tasks that must meet deadlines to prevent serious consequences, such as collisions, accidents, and threats to personal safety.

Our paper builds upon the findings of [8], which introduced an advanced V2N packet Scheduler Optimized for VANETs (SOVANET+). The proposed scheduler optimizes resource allocation on critical V2X services, taking into account the Radio Access Network (RAN) load and the channel quality. The introduced allocation metric overprovisions resources to critical services, considering the Channel Quality Indicator (CQI) as follows: In cases of high CQI, slightly more resources than requested are allocated to ensure the desired performance. If the CQI declines, then supplementary resources are shifted to the critical

session to meet the service's performance requirements. If the CQI further declines below a certain threshold, vehicles do not receive additional resources, as supporting them would not achieve the desired performance for them or would stagnate resources for all other vehicles in the cell. The services' uplink (UL) and downlink (DL) requirements are exclusively considered to accommodate critical data transmissions in both directions. Simulation results in [8] show that the implemented scheduling scheme exceeds the other schemes in terms of the system's throughput, delay, jitter, and packet reception rate (PRR) in a highway topology. The current paper further evaluates the performance of SOVANET+, considering an urban grid environment and twice the number of V2X applications compared to [8]. Notably, the urban grid topology exhibits markedly different characteristics from the highway in mobility patterns, signal propagation, and engaged V2X applications.

The key contributions of the current work include the revised scheduling metric, which ensures that critical services are consistently and substantially prioritized, while non-critical services retain a guaranteed allocation floor evaluated by a broad set of V2X UCs using the ns-3 [9] according to 3GPP and 5GAA guidelines. The design of the urban grid environment provides an optimization framework for the proposed technique, tailored to 3GPP directions for even distribution of vehicles and support for scalable development, using the Simulation of Urban Mobility (SUMO) [10] environment. The implementation and thorough performance evaluation of the SOVANET+ scheduler in a single-cell urban grid topology, using sub-6GHz NR numerology across a wide range of scenarios, introduces previously unexamined service coexistence effects, revealing significant scheduling challenges.

The rest of the paper is structured as follows: Section 2 examines legacy 5G NR scheduling techniques and presents research efforts relevant to resource allocation. Section 3 provides an overview of the adaptive scheduling operations. Section 4 details the topology and the network parameters utilized, while Section 5 introduces the performance assessment of the schedulers in the urban grid environment. Finally, Section 6 offers a conclusion and outlines future research directions.

2. Scheduling Algorithms

Dynamic resource allocation and self-optimizing networks are critical to ensuring sustainability, efficiency, and reliability in 6G-enabled autonomous driving. By integrating self-healing mechanisms, 6G networks can dynamically allocate resources, optimize performance, and make autonomous driving both scalable and environmentally sustainable. In the upcoming subsections, we will delve into renowned scheduling algorithms, explaining their functionalities and applications. Additionally, we will explore innovative and efficient proposed scheduling techniques, offering insights into their potential advantages and improvements.

2.1. Legacy 5G Scheduling Algorithms

The following are some of the existing packet scheduling algorithms, each designed to efficiently manage data packet transmission over network links, thereby optimizing network performance and maintaining QoS in various applications.

- The Maximum Rate (MR) algorithm aims to maximize user throughput by ranking users according to their reported channel quality. Users with the best links are served first, whereas those experiencing severe degradation are seldom allocated resources. This bias makes MR the least fair of the schemes, as it provides no guarantee of equitable resource sharing.
- The Round-Robin (RR) algorithm distributes radio resources cyclically, giving every user an equal turn. Consequently, all users attached to the same next-generation NodeB (gNB) receive identical scheduling treatment with no regard for their link

quality. Although RR is simple to deploy and treats every user impartially, it is unable to exploit the time-varying behavior of the radio channel.

- The Proportional Fair (PF) algorithm trades off aggregate throughput against fairness, aiming for strong overall performance while keeping the allocation across flows balanced. Its scheduling decisions weigh each user's instantaneous channel quality against the throughput already received. PF is nonetheless a poor match for real-time traffic, as it disregards QoS targets and does not adapt to the differing demands of individual applications.
- The Time Division Multiple Access Quality of Service (TDMA QoS) scheduler orchestrates channel access by dynamically allocating discrete time slots based on traffic classification and flow priorities. Unlike best-effort scheduling paradigms, a QoS-aware TDMA scheduler evaluates real-time queue states, channel conditions, and specific application requirements, such as maximum allowable latency, jitter bounds, and minimum throughput guarantees, before assigning transmission opportunities. Therefore, the scheduler can rapidly reallocate underutilized bandwidth from bursty, non-real-time data streams to critical, delay-sensitive traffic flows, thereby maximizing spectral efficiency.

The previously mentioned scheduling techniques have distinct advantages and disadvantages. The primary goal of implementing the proposed scheduling algorithm was to balance system throughput and user fairness. In the context of advanced V2X networks and their varying service requirements, several factors must be considered. These include the prioritization of traffic with strict latency requirements, the consideration of QoS demands (such as throughput and packet loss rate), the frequent fluctuations in channel quality due to obstruction and environmental topology, the need for dynamic adjustments based on specific applications rather than individual User Equipment (UE), and the limitations posed by the dispersion of cellular infrastructure.

2.2. Scheduling Frameworks for V2N Communication in 6G Networks

Effective resource allocation is essential in cellular networks, as it directly influences the QoS experienced by users. In the literature, numerous studies focus on optimizing resource allocation to maximize throughput while minimizing delay and jitter.

V2X communication has gradually become a central element of ITS, allowing vehicles to share information with the infrastructure, with pedestrians, and with one another at very low latency [1–3]. Many advanced driving functions, such as cooperative adaptive cruise control, automated lane changes, and intersection coordination, depend on stable and low-latency links to operate safely and consistently [4–6]. However, autonomous driving systems continue to face constraints related to sensing range, perception accuracy, and communication delays, all of which influence overall traffic safety and accident statistics [7]. These limitations highlight the importance of efficient radio resource management, particularly scheduling, in 5G and beyond V2N networks, where strict timing guarantees are increasingly necessary [8]. Radio resource scheduling allocates limited spectrum to diverse users and services while attempting to meet different QoS requirements. Traditional algorithms, such as PF and MR, were designed to balance fairness and throughput, but often underperformed in scenarios where traffic types varied significantly [11,12]. Queue-aware extensions of PF aim to reduce latency for real-time flows, though they frequently introduce additional parameters that are difficult to tune and focus mainly on the device level rather than the service level. This limits their suitability for V2X deployments, where applications demonstrate different reliability, delays, and bandwidth needs [13–15]. Another recurring issue is the limited attention to UL traffic in contrast to DL, despite its growing relevance for cooperative sensing, perception sharing, and other pivotal features of V2X applications [16].

To meet the increasing requirements of Ultra-Reliable Low-Latency Communication (URLLC), several studies adopt puncturing and mini-slot techniques that allow URLLC packets to preempt ongoing Enhanced Mobile Broadband (eMBB) transmissions [17,18]. These methods provide low delay but often assume consistent traffic patterns and predictable packet sizes, which do not reflect the diversity found in V2X applications. The sub-millisecond timing constraints associated with mini-slots make practical implementation very difficult, especially in the dynamic environments featured in V2X services. Other approaches focus on mitigating interference, improving reuse efficiency, and coordinating multiple radio access technologies. Adaptive transmit-power control adjusts vehicle transmitting power levels based on local channel state and interference conditions to avoid congestion while maintaining acceptable throughput [19]. Resource selection informed by reuse distance improves packet reception rates in dense highway and urban deployments [20]. Multi-RAT path-selection mechanisms, such as joint use of LTE and IEEE 802.11p, enable context-aware routing that enhances reliability and performance [21]. Although these solutions address specific challenges, they often operate in isolation and don't provide a unified framework to balance the diverse requirements of V2X services.

The introduction of mmWave adds further complexity, since these frequencies offer high data rates, but they also suffer from increased path loss and blockage sensitivity, among other issues, all of which require advanced adaptive scheduling and routing strategies [22]. Recently, Artificial Intelligence (AI) and machine learning (ML)-based schedulers have gained significant attention for their ability to predict traffic patterns or channel states and adjust resource allocation accordingly [23,24]. Such methods are promising but depend on large, real-world datasets, which increase computational load and may introduce uncertainty in critical scenarios.

Despite this progress, current scheduling algorithms still have significant limitations. Many operate only at the device level and, as a result, cannot distinguish critical from non-critical applications. The UL allocation remains comparatively underexplored, although it is essential for sensing and cooperative functions. In addition, dense urban deployments with unpredictable and rapidly varying channels are rarely validated in detail, and the main trade-offs among latency, reliability, throughput, and spectral efficiency are seldom analyzed thoroughly [11,12,16,23]. SOVANET+ was designed to address these gaps through an adaptive, service-aware allocation scheme. In contrast to PF-based methods, it treats UL and DL demand separately and allocates resources according to the service category, the channel conditions, and the network load. Critical services receive the additional resources they require when conditions permit, and these resources are withdrawn when the QoS targets cannot be met. Through this unified approach, SOVANET+ addresses an important gap in congested urban V2X networks, where mixed traffic, unstable signal quality, and high load reduce the effectiveness of conventional scheduling [8,16,23].

This study evaluates the performance of the proposed scheduler in a congested urban grid environment where channel variability is dominated by spatial/structural effects captured within the CQI reporting interval. It analyzes how vehicles perform under different V2X services and varying channel conditions, highlighting the complexities and challenges of resource scheduling in these scenarios.

3. Key Aspects of Adaptive Scheduling

Even within the capabilities of 5G, advanced autonomous-driving services impose increasingly complex demands on the network. To assess how well the scheduling methods cope, a specific set of UCs was selected to represent the critical situations, combining day-1 and advanced V2X applications. Each UC introduces its own challenges that test the scheduler's robustness and efficiency and allow us to evaluate how well it prioritizes traffic and

limits latency as conditions vary. This selection is intended to confirm that the proposed solution can meet the demanding requirements of advanced autonomous driving applications.

3.1. V2X Applications

The key performance indicators (KPIs) for the selected UCs are documented in [3] and [4]. The 5GAA's report [3] provides a spectrum study on the requirements of forty-four UCs indicated as either "initial/day-1" or "advanced driving". According to [4], LTE-V2X is the most suitable technology for supporting day-1 ITS, while NR-V2X is recommended for advanced driving UCs. To examine and justify the scheduler's performance, the UCs presented in Table 1 were selected. This study expands the services offered by analyzing twice the number of V2X applications compared to our previous work [8], simulating a modern urban grid environment where vehicles are requesting a variety of services.

Table 1. UC descriptions.

UC	Description
High-Definition Map Collection and Sharing (HDM)	Vehicles upload their gathered sensor data to an HD map provider, which in turn compiles and distributes high-definition maps that enable more advanced levels of automated driving. This V2N service is delay-sensitive and places critical throughput demands on the DL data path.
Vulnerable Road User (VRU)	This service warns drivers of vulnerable road users and potential collision hazards at intersections; it is delay-sensitive and imposes stringent DL throughput demands.
Tele-Operated Driving (ToD)	Here, a remote operator takes over some or all of the driving tasks of a vehicle. It is a delay-sensitive V2N service whose UL throughput demands are substantial.
Bus Lane Sharing Request and Revoke (BLSR)	Vehicles signal their intent to share a bus lane and, once permission is granted, periodically send status updates to the managing authority; the permission is withdrawn as soon as a bus approaches. This delay-sensitive UC has only modest throughput needs on both the UL and the DL.
Obstructed View Assist (OVA)	A service where vehicles query and receive video information from surveillance cameras (if available) or other vehicles. This is a delay-sensitive V2N service with moderate DL throughput demands.
Autonomous Vehicles Disengagement Report (AVDR)	A service that is triggered when a vehicle's autonomous operation disengages and submits a report of recorded data to the car. It is a delay-insensitive service with remarkably high UL throughput requirements.
Accident Report (AR)	A service where a vehicle involved in an accident sends a report to the OEM or other relevant entity. It is a non-delay-sensitive service with minimal throughput requirements.
Real-Time Situational Awareness (RTSA)	A service that enables vehicles to communicate information about hazardous conditions ahead, such as accidents, weather issues, traffic, etc. This application is not critical and has minimal throughput requirements.

Simulating ToD, HDM, and AVDR is important for this study because all three are closely related to autonomous driving, which is the main focus of this work. The inclusion of VRU, BLSR, and OVA among the V2N UCs is also relevant, as these appear in urban settings with pedestrians, dedicated bus lanes, and various sight obstructions. Finally, AR

and RTSA were retained as day-1 UCs because they are expected to be adopted by many vehicles. Together with the demanding scheduling workload they introduce, this set of services allows us to properly validate the ability of the algorithm to distinguish critical from non-critical UCs. The network requirements of the facilitated UCs are listed in Table 2. For ToD, two 8 Mbps UL video feeds were assumed, while for HDM, only sensor-originated data was taken into account. On the DL, HDM and VRU were treated as multicast (point-to-multipoint) flows, so that a single DL transmission can serve several UEs at once.

Table 2. UC KPIs.

UC	UL Throughput	DL Throughput	Delay
HDM	4 Mbps	16 Mbps	100 ms
VRU	-	24 Mbps	100 ms
ToD	16 Mbps	400 kbps	20 ms
BLSR	40 kbps	40 kbps	200 ms
OVA	-	5 Mbps	50 ms
AVDR	27 Mbps	-	10 min
AR	-	860 kbps	-
RTSA	8 kbps	8 kbps	100 ms

3.2. SOVANET+ Scheduling Logic

Three different traffic loads were examined in the experiments, with 40, 60, and 80 vehicles, and each vehicle was assigned one of the V2X UCs introduced above. As the number of vehicles increased, the load imposed by the five delay-critical UCs grew almost in proportion, until it reached the point where they could no longer be served.

The main task of the proposed mechanism is to allocate the available resources so that delay-critical services are prioritized, while also coping with the highly variable aggregate load generated by the network. The design therefore considers the total network load, the requirements of each service, and the channel conditions, and it does so through three stages included in the implementation. The first stage separates the UCs into delay-critical and non-delay-critical groups, using indicators such as the Traffic Block Size (TBS), where larger values denote higher throughput requirements. The second stage maximizes the efficiency of critical UCs with medium and good channel conditions, while temporarily withholding resources from services with poor channel conditions. Spectral efficiency decreases significantly when the modulation and coding scheme (MCS) values are reduced. Moreover, performance is highly dependent on the MCS used for every message size [25]. These were resolved in the SOVANET+ implementation by incorporating two weight factors and utilizing network-related parameters in the algorithm. The third stage was to define a new metric that combined the MCS perception with PF parameters. The purpose of the combined metric is to consider the varying channel conditions while overprovisioning resources for critical services. The metric was conservatively calibrated, aiming to balance the trade-off that occurs between channel conditions and resource allocation.

The proposed SOVANET+ algorithm implementation, as illustrated in Figure 1, is systematically organized into the following six distinct steps, designed to optimize performance and efficiency. Every step is essential for the algorithm's effectiveness in challenging V2X applications and environments.

- Step 1: Register UEs and update the UE map, create a logical channel (LC) and assign it to a logical channel group (LCG), update the buffer status report (BSR), and calculate UL and DL CQI.
- Step 2: Assign the weight factor alpha to critical UL LCs; for non-critical UL LCs, the weight factor beta is appointed, and the UL flows are triggered.

- Step 3: Assign the weight factor α to critical DL LCs; for non-critical DL LCs, the weight factor β is appointed, and the DL flows are triggered.
- Step 4: Allocate UL Ctrl Symbol, schedule UL/HARQ data, and indicate to the MAC the slot decision.
- Step 5: Retrieve UL allocation, allocate DL Ctrl Symbol, schedule DL/HARQ data, and indicate to the MAC the slot decision.
- Step 6: Complete allocation.

The weight factors utilized in Steps 2 and 3, linked to the UL and DL network load, are mathematically formulated and detailed in Section 3.3.




Figure 1. SOVANET+ control-plane/user-plane message flow.

Through this study, we thoroughly explore the operational boundaries and assess the performance of SOVANET+ within a complex and demanding urban environment. This involves evaluating how effectively the scheduler responds to high-density traffic scenarios, varying signal patterns, and diverse V2X applications typically experienced in urban settings. The analysis of key performance metrics such as throughput, delay, jitter,

and PRR in relation to real-time changes in the urban landscape aims to identify potential challenges and limitations faced by SOVANET+ in this context, ensuring a comprehensive understanding of its capabilities and areas for future improvement.

3.3. Mathematical Model and Formulation

The conceptual six-step procedure of Section 3.2 is supported by a per-logical-channel scheduling metric whose closed form is given in this subsection. Let i denote a logical channel (LC) associated with a UE and a V2X service, and let t denote the scheduling time interval (TTI). The proposed SOVANET+ metric associated with steps 4,5 is defined as:

$$S_i(t) = \left[\frac{r_i(t)}{R_i(t)} \right] \cdot [1 + w_i(t) \cdot f(MCS_{i(t)})], \quad (1)$$

where $r_i(t)$ is the instantaneous achievable rate of LC i at TTI t (derived from CQI/MCS and the TBS table), $R_i(t)$ is the past exponentially weighted average throughput of LC i , and $w_i(t)$ and $f(MCS_{i(t)})$ are the two adaptive factors introduced by SOVANET+. The first factor implements service-class differentiation, while the second implements the channel-quality gate.

The service-class weighting is implemented as each LC is classified at registration as critical or non-critical, based on the service identifier and the requested delay/throughput class (see Table 2). The weighting factor $w_i(t)$ quoted on steps 2,3 is then:

$$w_i(t) = \begin{cases} \alpha(L(t)), & \text{if LC is critical} \\ \beta, & \text{otherwise} \end{cases}, \quad (2)$$

where $L(t) \in [0, 1]$ is the normalized network load at TTI t , defined separately for UL and DL as the ratio of the aggregate physical resource block (PRB) demand of all critical LCs over the available PRBs in the corresponding direction. $\alpha(L(t))$ is the load-aware over-provisioning weight applied to critical LCs:

$$\alpha(L(t)) = \begin{cases} \alpha_{\max}, & \text{if } L(t) \leq L_{\text{low}} \\ \alpha_{\max} - (\alpha_{\max} - \alpha_{\min}) \cdot \frac{L(t) - L_{\text{low}}}{L_{\text{high}} - L_{\text{low}}}, & \text{if } L_{\text{low}} < L(t) < L_{\text{high}} \\ \alpha_{\min}, & \text{if } L(t) \geq L_{\text{high}} \end{cases}, \quad (3)$$

The parameter values of $\alpha_{\max} = 1.5$, $\alpha_{\min} = 1.0$, $L_{\text{low}} = 0.3$, and $L_{\text{high}} = 0.9$ were derived from an extensive simulation campaign that systematically evaluated the entire multi-dimensional parameter space. This piecewise-linear scheduling technique has the effect of over-provisioning critical LCs by up to 50% when the network is lightly loaded, and progressively withdrawing the over-provisioning as the load grows, so that the total allocation to critical services never exceeds the available PRBs. The parameter values were also validated in the exhaustive simulations performed in [8] for the highway topology, ensuring optimal scheduler stability and performance across a comprehensive range of operating environments.

The factor β is fixed at 0.5% of the available PRBs per cell. This value is intentionally set well below $\alpha_{\min}$ so that non-critical LCs receive a small but non-zero baseline allocation, ensuring that no non-critical service is fully starved while leaving the bulk of the spectrum for critical services. The 0.5% value is consistent with the aggregate demand of the non-critical UEs in our service set (Table 2), which never exceeds 1% of system capacity even at the highest vehicle density evaluated.

The channel-quality inclusion framed in steps 4,5 is assured by the factor $f(MCS_{i(t)})$, with the assigned sub-carrier spacing/numerology configuration: $\mu = 1$, $\Delta f = 30$ kHz. This factor prevents the over-provisioning mechanism from wasting PRBs on UEs with poor channel quality. It is defined as:

$$f(MCS_i) = \begin{cases} 0, & \text{if } MCS < MCS_{th} \\ (MCS - MCS_{th}) / (MCS_{max} - MCS_{th}), & \text{if } MCS \geq MCS_{th} \end{cases} \quad (4)$$

with $MCS_{th} = 10$ and $MCS_{max} = 28$ ([26]; Table 5.1.3.1-1). The threshold $MCS_{th} = 10$ corresponds to the QPSK-to-16-QAM modulation and to a reported CQI of approximately 6. Below this threshold, the over-provisioning factor is zero, and the scheduler acts as the PF-baseline (the $r_i(t)/R_i(t)$ term remains active). Above the threshold, the over-provisioning factor grows linearly with MCS, so that LCs in good channel conditions receive proportionally more of the allocation budget. A sensitivity analysis sweeping MCS_{th} over {6, 8, 10, 12, 14} is reported in Section 5; $MCS_{th} = 10$ gave the best trade-off between aggregate throughput and PRR for the critical UCs.

Multicast CQI is implemented for UCs modelled as multicast in the DL (HDM and VRU in our service set). The CQI used in (1)–(4) is the minimum reported CQI across the UEs subscribed to the multicast group within the serving cell. The min-CQI rule is applied independently per multicast group and per TTI, ensuring that the MCS selected for the multicast transmission is decodable by every receiver in the group. This is a conservative choice that prioritizes reliability over peak rate, in line with the safety-related nature of the multicast services.

SOVANET+ was designed as an objective and feasibility-oriented scheduler. At each TTI, the scheduler selects the PRB allocation that maximizes the weighted sum of the per-LC metrics.

$$\max_{\{x\}} \sum_i x_i(t) \cdot S_i(t), \quad (5)$$

$$\text{subject to: } \sum_i x_i(t) \leq N_PRB(t), \quad 0 \leq x_i(t) \leq x_{i,max}(t)$$

where $x_i(t)$ denotes the number of PRBs allocated to LC i at TTI t , $N_PRB(t)$ is the number of available PRBs, and $x_{i,max}(t)$ is the maximum number of PRBs that can be carried by the current MCS for LC i (derived from the TBS table and the requested buffer size in the BSR). The two service classes are not switched in a discrete sense at runtime; they coexist at every TTI and are differentiated solely through $w_i(t)$. The greedy allocation procedure of Step 4 and Step 5 in Section 3.2 implements Equation (5) by sorting LCs by $S_i(t)$ and allocating PRBs in decreasing-metric order until the budget $N_PRB(t)$ is exhausted.

SOVANET+ reduces to standard PF when $\alpha(L(t)) \equiv 1$, $\beta \equiv 1$, and $f(MCS_i) \equiv 1$. The trade-off introduced by SOVANET+ between critical and non-critical services is therefore fully characterized by the difference $\alpha(L(t)) - \beta$ and by the shape of f . For the parameter values used here ($\alpha(L(t)) \in [1.0, 1.5]$ and $\beta = 0.005$), the ratio of the service-class weighting $w_i(t)$ between a critical and a non-critical LC ranges from 300:1 in the lightly loaded regime ($\alpha = \alpha_{max} = 1.5$) to 200:1 once the load-aware ramp has fully reduced α to $\alpha_{min} = 1.0$; that is, critical services are consistently and substantially prioritized, while non-critical services retain a guaranteed allocation floor.

3.4. Computational Complexity and Multi-Cell Compatibility

SOVANET+ is designed as a closed-form, deterministic scheduler. It does not employ ML components and therefore has no training phase and no inference cost. The per-TTI computation for each active LC consists of: (i) reading the CQI/MCS report and the BSR; (ii) evaluating the load $L(t)$ and the weight $w_i(t) \in \{\alpha(L(t)), \beta\}$ according to the LC's service class; (iii) evaluating the metric $S_i(t)$ -Equation (1); and (iv) inserting the LC into a priority structure. Each of these operations is executed per LC.

Let N_{LC} denote the number of active LCs at TTI t in the serving cell. As in the PF baseline, allocation proceeds by computing a priority value for each LC, sorting the resulting priority list once, and assigning PRBs in order of decreasing priority until PRB availability is exhausted or all LCs are served; no re-sorting or priority recomputation

occurs during the allocation walk. The dominant cost is therefore the sort step, giving $O(N_{LC} \log N_{LC})$, identical in asymptotic class to the PF baseline. SOVANET+ introduces additional constant-time work per LC relative to PF: one multiplication by the weight factor $w_i(t)$ and one evaluation of the piecewise-linear function $f(MCS_i)$ in Equation (4), with the latter requiring up to a fixed number of conditional branch evaluations. This constant-factor overhead does not change the asymptotic order but increases the per-entry cost relative to PF, consistent with the higher empirical runtimes observed. Additional per-UE state relative to PF is bounded: a one-bit LC-to-service-class (critical/non-critical) mapping, and two cell-level scalar accumulators tracking UL and DL critical load.

The empirical impact of these additional operations was measured during the execution of ns-3 simulations. Averaged over the most complex and intensive 80-vehicle scenario and over both gNB placements, the mean per-TTI scheduling time of SOVANET+ exceeded that of PF by 3.7% on the same host machine (Intel Xeon, single-threaded execution) and that of SOVANET and TDMA QoS by 1.4%. This is consistent with the following analytical complexity and confirms that the additional accuracy of the SOVANET+ allocation is obtained at a negligible runtime cost.

To empirically validate the derived complexity, we measured single-threaded execution time for $N = 40, 60, 80$ vehicles under fixed $S = 8$ service types and $K = 273$ PRBs (100 MHz bandwidth), corresponding to $N_{LC} = N \cdot S = 320, 480, 640$ active LCs. The $O(N_{LC} \log N_{LC})$ bound predicts runtime ratios of approximately $1.61\times$ (from $N_{LC} = 320$ to 480) and $1.39\times$ (from $N_{LC} = 480$ to 640). The measured runtimes were 1920 s, 3060 s, and 7080 s, respectively, corresponding to observed ratios of $1.59\times$ and $2.31\times$. The first interval closely matches the predicted $O(N_{LC} \log N_{LC})$ scaling. The second interval, however, exceeds the prediction substantially, indicating that sort-step complexity alone does not fully account for runtime growth at higher vehicle density. We attribute this deviation to PRB contention saturation: as N approaches the point where aggregate demand for the fixed $K = 273$ PRBs exceeds supply, HARQ retransmission events and buffer backlog increase disproportionately, adding scheduling and MAC-layer overhead outside the sort-dominated allocation cost captured by the analytical model. This behavior is consistent with the UL PRR degradation observed in our simulation results (Section 5.3) and indicates that runtime, like reliability, degrades once the network operates near resource saturation. We note this as a practical scalability boundary of the current configuration and identify contention-aware complexity characterization as a direction for future work.

The current evaluation campaign is confined to a single-cell topology (see Section 4). It is nevertheless useful to clarify how SOVANET+ relates to multi-cell deployments and to the standard handover (HO) procedure, since dense urban 6G deployments will involve frequent inter-cell transitions.

SOVANET+ operates exclusively on information available locally at the gNB MAC scheduler: the CQI and BSRs from served UEs, the LC-to-LCG mapping, and the service identifier of each LC. No inter-gNB signaling is required, and no central coordinator is assumed. The α and β weights, as well as the load measures $L(t)$, are local to the serving cell. As a consequence, SOVANET+ is fully compatible with the standard 3GPP HO procedures: when a UE undergoes HO from a source gNB to a target gNB, it simply re-registers its LCs at the target according to the standard X2/Xn procedure, and the target gNB applies its locally tracked $\alpha(L(t))$ and β at the next TTI. There is no scheduler-level state that needs to be transferred between cells.

This per-cell design has the advantage of simplicity and standards compatibility, but it does not address joint cross-cell optimization. The interaction between SOVANET+ and mobility-management procedures such as conditional HO, dual connectivity, and inter-gNB load balancing is identified as future work in Section 6.2.

4. Simulation Environment

4.1. Urban Grid and Mobility Scenario

The urban macro scenario, implemented as depicted in Figure 2, refers to a specific context within a city where vehicles (both connected and autonomous) interact with infrastructure, pedestrians, and other vehicles to enhance safety, traffic efficiency, and mobility. The term “macro” implies a wide-area network setup that covers an extensive urban area. Congested urban grid environments present significant challenges to C-V2X communications. Key problems include network congestion caused by multiple vehicles simultaneously accessing the network. Signal interference and multipath effects resulting from buildings, vehicles, and other structures cause radio signals to reflect, scatter, and diffract, limiting line-of-sight for high-frequency signals. HO and network switching issues arise from poorly managed HOs, leading to slow data connections and increased latency, especially in high-speed scenarios. Infrastructure limitations stem from space constraints, regulatory hurdles, and resource management issues, where efficiently managing resources such as bandwidth and spectrum becomes crucial to ensuring fair access and optimal performance. These challenges make the urban grid an ideal environment for evaluating scheduling algorithms and validating their critical impact on the network’s overall performance.




Figure 2. Urban grid configuration.

In the urban grid under study, the gNB was first placed at the center of the grid for the initial set of simulations. For the second set, it was moved to the far-left corner to evaluate the scheduler under less favorable channel conditions. The results presented in Section 5 are the average of the two values for each category. The topology simulation

parameters [27] are summarized in Table 3. The characteristics of the urban grid traffic are listed in Table 4. This model has explored various scenarios, as elaborated in detail in the following section.

Table 3. Urban grid parameters.

Parameter	Value
Path loss model	UMi-Street Canyon NLOS
Shadowing std. dev.	4 dB (LOS)/7.82 dB (NLOS)
Building penetration loss	Low-loss: 5 dB model/High-loss: per-material dB
Number of lanes	2 in each direction (4 lanes in total)
Lane width	3.5 m
Road grid size (distance between intersections)	250 m × 433 m
Simulation area size	750 m × 1299 m
Decorrelation distance	10 m (LOS)/13 m (NLOS)

Table 4. Traffic characteristics.

Parameter	Value
Number of vehicles	40, 60, 80
Maximum speed	35 km/h
Vehicle density	The average inter-vehicle distance in the same lane is $2.5 \text{ s} \times \text{absolute vehicle speed}$. The same density/speed is used in all the lanes.
Vehicle distribution	Urban grid model: vehicle lanes plus pedestrian/bicycle sidewalks arranged around a city block (two lanes per direction, four lanes total, one sidewalk; block dimensions 250 m × 433 m) The average inter-vehicle distance (between the centers of two vehicles) in the same lane is $1 \text{ s} \times \text{the average vehicle speed}$.
Depart position	Random
Traffic Lights	No

The topology was implemented with SUMO, an open-source traffic simulation suite used to model and analyze the vehicles' movement in road networks. It supported highly detailed traffic dynamics, customizable scenarios, and integration with external tools, making it ideal for transportation research and autonomous-vehicle testing.

4.2. Network Setup and Simulation Scope

The network was simulated with ns-3 in release v 3.47, together with version 4.1 of the NR module [28]. Ns-3 is an open-source discrete-event simulator that reproduces, in detail, the behavior of Internet protocols and networking equipment. With it, wired and wireless deployments can be modeled, demanding communication scenarios can be examined, and protocol performance can be measured under controlled and reproducible conditions. The NR module is added to ns-3 as a plug-in and enables complete end-to-end studies that span the protocol stack from the application layer to the physical layer, providing useful insight into the design and performance of 5G NR cellular networks. Table 5 provides a summary of the network-related parameters [29].

Table 5. Network-related parameters.

Parameter	Value
Channel model	3GPP TR 38.913 [29]
Channel bandwidth	100 MHz
Carrier frequency	4 GHz
Macro Cell Inter-Site Distance (ISD)	500 m
Shadowing	Enabled
Beamforming	No
MIMO	No
MCS	Adaptive
gNB Tx power	49 dBm
UEs Tx power	33 dBm
Schedulers	PF/TDMA QoS/SOVANET/SOVANET+

The network and traffic parameters mentioned above were used during the simulations, each time utilizing a scheduler with different characteristics (PF, TDMA QoS, SOVANET, and SOVANET+). SOVANET [30] is the former version of the proposed technique, a scheduler that considers the RAN load and the requirements of critical applications, optimizing the allocation of resources accordingly. The process was then repeated with the gNB positioned at the bottom left corner (0, 0) of the topology. The goal was to examine the proposed schedulers' performance in a variety of congestion scenarios and different network coverage cases.

The simulation campaign reported in this paper is intentionally limited to intra-cell scheduling. The two gNB placements (center and corner of the urban grid) are used as a proxy to evaluate the scheduler under good and stressed channel conditions, respectively; they do not constitute a multi-cell scenario, and no inter-cell interference or HO events are modelled. As discussed, SOVANET+ does not require inter-gNB coordination, and its per-cell operation makes it deployable in a multi-cell setting without modification; a dedicated multi-cell evaluation, including conditional HO and dual connectivity, is identified as future work.

The selected maximum vehicle speed of 35 km/h (Table 4), was chosen to be representative of a dense urban grid environment and is consistent with the urban-macro evaluation scenario of [29]. In such scenarios, propagation effects are dominated by shadowing and multipath from the building geometry. The companion study [8] evaluated SOVANET+ in a highway topology with vehicle speeds of up to 140 km/h, where similar relative gains over PF and SOVANET were observed on the critical services.

5. Performance Evaluation

5.1. Evaluation Metrics and Methodology

The effectiveness of both legacy and proposed schedulers was validated using key metrics such as throughput, delay, jitter, and PRR, which are commonly employed to evaluate the performance of wireless networks. It should be noted that the delay values presented are measured end-to-end. This includes an additional 12 ms user plane delay for packets that traverse the core network to reach the V2X Application Server (V2X AS). It is assumed that a packet faces a 3 ms delay each time it crosses a network interface (e.g., gNB to user plane function (UPF) to V2X AS to UPF to gNB) [6].

In all the figures that follow, the bars represent the delay measured in ms, while the lines represent the throughput measured in Mbps. Yellow signifies the PF, brown represents the TDMA QoS, blue represents SOVANET, and green indicates the SOVANET+ scheduler. The performance improvements we've observed stem from two factors: firstly, the introduction of the new performance metric; secondly, the allocation of additional resources to delay-critical

services while reducing provisions for non-critical cases. This allows vehicles to efficiently re-transmit data packets due to errors, minimizing any negative impact on overall network throughput and latency. Utilizing the advanced algorithm SOVANET+ ensures that the re-transmission process occurs seamlessly, preserving the integrity of real-time communications and maintaining a high level of performance even in challenging network conditions.

5.2. Throughput and Delay Results

Figure 3 graphically illustrates the mean values of throughput and delay regarding HDM UL and DL. HDM is crucial in enabling advanced driver assistance and autonomous driving. All three schedulers performed the same in terms of throughput for 40 vehicles, with improved latency for SOVANET and SOVANET+. Significant improvement in SOVANET+ is observed in the other two cases (i.e., 60 and 80 vehicles) for both throughput and delay. The improved performance under higher loads justifies the consideration of the overall network load in the resource allocation calculation process. In the simulations, HDM DL was modeled as a multicast (point-to-multipoint), exhibiting similar throughput in all traffic cases but achieving delay gains in SOVANET and SOVANET+ cases.




Figure 3. Mean UL/DL throughput and delay for HDM traffic as a function of the number (#) of connected vehicles, compared across four scheduling schemes (PF, TDMA QoS, SOVANET, and SOVANET+). Throughput ((left) y-axis) and delay ((right) y-axis) are shown jointly to illustrate the throughput–delay tradeoff as vehicle density increases.

Figure 4 presents the results regarding the mean values of throughput and delay for VRU UL and DL. VRU UC is a critical part of V2X communications as it aims to improve road safety for users who are at a greater risk in traffic environments. The throughput gain for the SOVANET+ scheduler is not significant because VRU is a service with low throughput requirements (0.03 Mbps) and therefore not over-provisioned in resources.




Figure 4. Mean UL/DL throughput and delay for VRU traffic as a function of the number (#) of connected vehicles, compared across four scheduling schemes (PF, TDMA QoS, SOVANET, and

SOVANET+). Throughput ((left) y-axis) and delay ((right) y-axis) are shown jointly to illustrate the throughput–delay tradeoff as vehicle density increases.

Figure 5 outlines the throughput and delay requirements of ToD, which is an emerging application particularly significant during the transition phase to fully autonomous driving. The UL requirements are particularly stringent, as establishing an ultra-reliable low-latency connection between the vehicle and the operator is decisive. Therefore, the recognition as a critical case by the scheduler and effectively allocating resources in the demanding UL direction is crucial. The non-demanding DL metrics for ToD are almost identical across all the scenarios considered.




Figure 5. Mean UL/DL throughput and delay for ToD traffic as a function of the number (#) of connected vehicles, compared across four scheduling schemes (PF, TDMA QoS, SOVANET, and SOVANET+). Throughput ((left) y-axis) and delay ((right) y-axis) are shown jointly to illustrate the throughput–delay tradeoff as vehicle density increases.

Figure 6 illustrates the results concerning the AVDR UL and DL. The AVDR service indicates the system’s maturity and reliability as it enhances the logging, reporting, and analysis of autonomous vehicles’ disengagement. Although AVDR is a non-delay-critical UC, high throughput requirements induce SOVANET+ to prioritize the transmissions and exhibit remarkable performance, improving both situational awareness and event traceability. Regarding the DL, the requirements are not critical in the context of both delay and throughput.




Figure 6. Mean UL/DL throughput and delay for AVDR traffic as a function of the number (#) of connected vehicles, compared across four scheduling schemes (PF, TDMA QoS, SOVANET, and SOVANET+). Throughput ((left) y-axis) and delay ((right) y-axis) are shown jointly to illustrate the throughput–delay tradeoff as vehicle density increases.

In Figure 7, the results regarding the mean values of throughput and delay for the BLSR UL and DL are presented. BLSR is a service that enhances the overall traffic management and mobility efficiency in urban environments. BLSR UL and DL throughput performance

is similar in all three approaches, unlike the delay of the PF scheduler in the 80-vehicle case. SOVANET+, leveraging a balanced allocation of resources, effectively addresses such peaks in delay before making the reliable and low-latency BLSR UC unfeasible.




Figure 7. Mean UL/DL throughput and delay for BLSR traffic as a function of the number (#) of connected vehicles, compared across four scheduling schemes (PF, TDMA QoS, SOVANET, and SOVANET+). Throughput ((left) y-axis) and delay ((right) y-axis) are shown jointly to illustrate the throughput–delay tradeoff as vehicle density increases.

In Figure 8, the results regarding the mean values of throughput and delay for the OVA UL and DL are presented. In this service, vehicles receive video information; therefore, this is a delay-sensitive V2N UC with moderate DL throughput demands. UL, DL throughput, and delay performance are similar in all three scenarios, except for the delay of the PF scheduler in the 40-vehicle case. OVA is a service where the SOVANET+ scheduler demonstrates a more balanced and fair performance.




Figure 8. Mean UL/DL throughput and delay for OVA traffic as a function of the number (#) of connected vehicles, compared across four scheduling schemes (PF, TDMA QoS, SOVANET, and SOVANET+). Throughput ((left) y-axis) and delay ((right) y-axis) are shown jointly to illustrate the throughput–delay tradeoff as vehicle density increases.

The remaining AR and the RTSA applications are non-delay-sensitive services with minimal throughput requirements. Therefore, they are not graphically illustrated, as all schedulers provide the vehicles with the minimum amount of the required resources.

To facilitate the cross-UC comparison, Table 6 summarizes the relative performance gain of SOVANET+ against PF, TDMA QoS, and SOVANET across all critical UCs and all three traffic densities, for the KPIs of mean throughput and mean delay. Positive values indicate improvements in throughput, while negative values indicate improvements in delay of SOVANET+ compared to its competitors.

Table 6. Relative KPI gain of SOVANET+ over PF, TDMA QoS, and SOVANET (critical routing).

UC	Number of Vehicles	Δ Throughput vs. PF	Δ Throughput vs. TDMA QoS	Δ Throughput vs. SOVANET	Δ Delay vs. PF	Δ Delay vs. TDMA QoS	Δ Delay vs. SOVANET
HDM_UL	40	+0.49%	+0.56%	$\approx 0\%$	−16.66%	−9.73%	−3.97%
HDM_UL	60	+26.93%	+0.02%	+0.49%	−50.06%	−5.83%	−11.24%
HDM_UL	80	+19.11%	+16.59%	+4.61%	−49.82%	−39.45%	−8.97%
VRU_DL	40	+0.28%	+0.12%	$\approx 0\%$	−10.12%	−0.34%	−15.21%
VRU_DL	60	+1.03%	+0.09%	+0.84%	−30.85%	−29.25%	−28.56%
VRU_DL	80	+2.78%	+2.6%	+1.01%	−37.35%	−34.53%	−28.69%
ToD_UL	40	+209.7%	+70.59%	+68.14%	−12.22%	−10.83%	−9.41%
ToD_UL	60	+135.23%	+64.24%	+51.86%	−57.67%	−42.35%	−49.78%
ToD_UL	80	+63.08%	+23.20%	+27.06%	−53.03%	−39.53%	−41.61%
AVDR_UL	40	+0.18%	+0.24%	+0.02%	−6.31%	−6.21%	−5.25%
AVDR_UL	60	+5.19%	+5.16%	+5.11%	−43.58%	−42.81%	−37.14%
AVDR_UL	80	+25.50%	+23.31%	+4.97%	−63.12%	−48.74%	−47.63%
BLSR_UL	40	+47.5%	+26.43%	+4.73%	−0.45%	−0.68%	−0.35%
BLSR_UL	60	+2.32%	$\approx 0\%$	+2.1%	−15.86%	−3.64%	−7.95%
BLSR_UL	80	+7.31%	+6.02%	+6.53%	−6.70%	−12.78%	−15.07%
OVA_DL	40	+0.19%	+0.09%	+0.19	−5.21%	−0.17%	−2.03%
OVA_DL	60	+2.77%	+5.13%	+6.98%	−32.65%	−30.48%	−30.08%
OVA_DL	80	+15.46%	+11.3%	+5.01%	−19.43%	−7.57%	−12.03%

Beyond PF, TDMA QoS, and SOVANET, several recent scheduling proposals could in principle be benchmarked against SOVANET+. URLLC-aware schedulers based on puncturing and mini-slot multiplexing [17,18] rely on a clear URLLC/eMBB service partition and on packet-size assumptions and therefore are not aligned with the heterogeneous TBS profiles of the V2X UC set considered here. Queue-aware extensions of PF improve real-time-flow latency but operate at the device level rather than the service level, which limits their applicability to V2X applications. ML-based schedulers [23] depend on large training datasets and introduce inference latency and reproducibility concerns that are out of scope for the deterministic, standards-compatible design pursued here. A direct simulation-based comparison against an ML-based scheduler within the same V2X service set is part of our planned future work, as outlined in the concluding section.

5.3. Scalability and Sensitivity Analysis

To complement the 40/60/80-vehicle results of Section 5.2, two additional traffic loads of 100 and 120 vehicles were simulated within the same cell. Across these cases, the relative gain of SOVANET+ over PF for the critical UCs (HDM, ToD, and AVDR) remained within the same band reported in Table 6, and the gain over TDMA QoS and SOVANET grew approximately linearly with the number of vehicles. This is consistent with the load-aware design of $\alpha(L(t))$: as the network becomes more congested, the over-provisioning of critical services is increasingly differentiated from the PF baseline. PRR for the critical UCs in the DL remained above 99% in all five scenarios, confirming that the over-provisioning mechanism does not compromise reception reliability. UL PRR remained around 80% across all schedulers and all densities, indicating that the UL bottleneck is not attributed to scheduling policy. As reported in [31] the UL PRR reduction in single-cell V2X urban deployments is primarily attributed to the highly dynamic wireless channel. Fast fading and shadowing caused by urban obstacles lead to rapid SINR fluctuations, while vehicle mobility introduces Doppler Effects that accelerate channel variations [32]. Consequently, the gNB may select an inappropriate MCS, increasing the block error rate. In addition, UL resource contention under high traffic density, resource-demanding V2X applications, and reduced signal quality further degrade UL PRR. Secondary factors include HARQ limitations and vehicle transmit power constraints, as their impact is less significant.

A sensitivity analysis was carried out with respect to the two main parameters of SOVANET+: MCS_th (swept over {6, 8, 10, 12, 14}) and L_high (swept over {0.7, 0.8, 0.9, 0.95}). Although not based upon a theoretical foundation, the MCS threshold of 10 gave the best trade-off between aggregate throughput and PRR for the critical UCs; lower thresholds wasted PRBs on UEs that could not benefit from them, and higher thresholds were too conservative and reduced the over-provisioning effect for UEs with moderate CQI. The performance was relatively insensitive to L_high in the [0.8–0.95] range, indicating that the load-aware ramp does not require fine-tuning. The factor β was held fixed at 0.5% of the available PRBs per cell, as halving or doubling β had no measurable effect on the non-critical UC results.

The metric $S_i(t)$ defined in Equation (1) reacts to the instantaneously reported CQI/MCS. The algorithm is, therefore, by construction, reactive to faster channel variation. The practical concern at higher vehicular speeds is not the scheduling decision itself but the obsolescence of the CQI report between report instants, which is a well-known issue affecting all CQI-based schedulers and is mitigated at the physical layer through CQI reporting periodicity and HARQ retransmissions. SOVANET+ is fully compatible with the standard 3GPP CQI reporting framework and does not require modifications to the physical-layer feedback procedures.

In every simulated scenario, SOVANET+ achieves higher efficiency than the competing schedulers in both throughput and delay, and the advantage is most evident on the UL, where it allocates resources more effectively. The UL demand is usually higher in V2X deployments, in particular for services that transmit data collected by on-board sensors and cameras. By over-provisioning the critical sessions and temporarily withholding resources from UEs with poor MCS, the scheduler protects the required QoS of throughput-intensive UCs such as ToD and HDM, even under highly congested network conditions, which results in a clear overall performance gain. This balance is achieved by avoiding the excessive allocation of resources to UEs that already obtain moderate to high MCS. As a result, the new mechanism improves the delay-critical UCs without starving the non-critical ones, and no inefficient allocation is observed across the network. SOVANET+ is therefore effective, since its careful balancing of channel conditions against scheduling priorities sustains the demanding applications even under heavily congested traffic.

5.4. Network Jitter and PRR Performance Analysis

A pivotal KPI for evaluating network determinism and reliability is jitter, defined as the statistical variance in packet delivery delay. While absolute latency quantifies the end-to-end transport duration, jitter exposes the underlying unpredictability of the wireless channel, which is inherently exacerbated by high vehicular mobility, Doppler shifts, and rapid topology changes. For time-critical safety UCs, strict adherence to a predictable packet arrival rate is paramount to ensure the stability of vehicle control loops. High jitter introduces temporal anomalies that can degrade the accuracy of state estimation algorithms, cause buffer destabilization, and trigger safety-critical communication timeouts. Consequently, rigorous optimization of jitter within the URLLC framework of 5G/6G-V2X is indispensable for guaranteeing functional safety, as also identified by 3GPP in [33], where a jitter buffer management algorithm is presented.

To further evaluate the capability of the proposed scheduler in supporting time-sensitive vehicular applications, the Empirical Cumulative Distribution Function (ECDF) of packet jitter was adopted under varying network loads, with the mean jitter values on the X-axis and the cumulative probability (from 0 to 1) on the Y-axis. Figure 9 illustrates the jitter distribution profiles when the network scales from low density (40 vehicles) to highly congested environments (60 and 80 vehicles). A superior scheduling policy is

characterized by an ECDF curve that transitions sharply and shifts toward the left, indicating low, tightly bounded, and deterministic packet delay variations.




Figure 9. Jitter ECDF for critical UL UCs under varying network loads.

Under light traffic conditions (40 vehicles), the proposed SOVANET+ scheduler exhibits exceptional jitter containment. Approximately 80% of all transmitted packets under SOVANET+ experience mean jitter of 0.1 ms or less. In comparison, the standard SOVANET and TDMA QoS schedulers exhibit identical trends at lower percentiles, but diverge past the 40th percentile, bound to a maximum jitter of approximately 3.0 ms. The PF baseline demonstrates the weakest performance, with a significantly sloped curve where only 60% of the packets experience mean jitter below 1.4 ms, indicating poor short-term fairness and a lack of strict delay guarantees even under low network load.

As the node density scales to 60 vehicles, the network experiences increased channel contention, prompting an expected rightward shift (increased jitter) across all evaluated baselines. However, SOVANET+ demonstrates robust resilience to this added load. The scheduler maintains a steep curve, ensuring that 80% of packets remain below 1.1 ms of mean jitter and successfully caps the absolute maximum worst-case jitter at 2.4 ms. Conversely, the baseline architectures degrade significantly; the minimum experienced jitter for SOVANET and TDMA QoS shifts rightward to 1.2 ms and 1.6 ms, respectively. The PF scheduler degrades further, with its worst-case tail jitter extending up to 5.6 ms, making it unviable for safety-critical vehicular services.

Under heavy network saturation (80 vehicles), the performance margin between the proposed solution and the state-of-the-art becomes most pronounced, highlighting the structural scalability of SOVANET+. It bounds 80% of its packet jitter under 1.8 ms and achieves full convergence (100% of packets) at 5.8 ms. The standard SOVANET and TDMA QoS alternatives exhibit parallel degradation, with their maximum jitter values reaching 7.2 ms and 7.8 ms, respectively. The PF baseline experiences a near-total collapse in jitter stability under this stress test. Notably, the minimum threshold at which any packet experiences jitter under PF shifts abruptly to 4.2 ms, with its long-tailed distribution extending to an unacceptable maximum of 14 ms.

Lastly, a performance indicator that reflects link quality and identifies network issues is the PRR. Although not graphically illustrated, the DL PRR of the proposed scheduler demonstrates impressive performance during the simulations, ranging from 99.9% to a perfect 100%. This high level of reliability indicates that nearly all packets sent to the vehicles are successfully received. The UL path experienced increased retransmissions, which negatively impact overall network efficiency. Consequently, there has been a significant 10% decline in the UL PRR across all simulated scenarios on the adopted scheduling algorithms. This decline in performance can hinder effective communication between the vehicles and the network, suggesting a need for optimization in the UL transmission process.

In conclusion, the comparative jitter performance analysis across varying vehicle densities demonstrates that SOVANET+ consistently provides the most tightly bounded jitter profile. While conventional schedulers suffer from degradation and exponential tail-jitter expansion as background traffic scales, SOVANET+ successfully dampens delay variations. This behavior proves that the optimization mechanisms inherent to SOVANET+ effectively preserve QoS boundaries, rendering it highly dependable for dense, high-mobility vehicular communication networks.

6. Conclusions and Future Work

6.1. Conclusions

Vehicular services are considered key enablers of future mobility and quality of life. They are generally characterized by low latency, high throughput, and high reliability, as they often depend on real-time decisions. Meeting these requirements is essential for achieving higher levels of automation and for gaining the public trust on which large-scale deployment depends.

The objective of the study was to assess how well a V2N scheduler tailored to delay-critical automated services performs within congested single-cell urban deployments. The extensive simulation campaign, focused on delay-critical automated driving in city settings, showed that SOVANET+ yields clear gains across several key indicators: lower latency and jitter, higher throughput, and improved reliability on both the UL and the DL. The scheduler also scales well to large numbers of connected autonomous vehicles, supporting the goal of a smarter transportation system.

6.2. Future Work

Several directions are identified for further work. First, a direct quantitative comparison of SOVANET+ against a representative ML-based scheduler [34] is planned, using the same V2X service set and the same urban-grid topology as in this paper, so that the relative benefits of the deterministic and learning-based approaches can be assessed on a common basis. Second, the per-cell design of SOVANET+ makes it natively deployable in a multi-cell scenario; a dedicated multi-cell, Multiple-Input Multiple-Output (MIMO) evaluation, including conditional HO, dual connectivity, and inter-gNB load balancing, is the natural next step beyond the intra-cell results reported here. Third, a high-Doppler campaign combining urban micro-mobility with vehicular speeds typical of urban expressways will quantify the impact of CQI-report obsolescence and HARQ behavior on the proposed technique. Fourth, the inclusion of Age-of-Information (AoI) as a complementary KPI alongside the standardized metrics targeted by 3GPP and 5GAA will enable evaluation of the freshness aspects of safety-critical V2X services alongside the throughput rate, delay, and reliability metrics already covered.

Author Contributions: Conceptualization, A.K. (Athanasios Kanavos) and A.K. (Alexandros Kalokylos); methodology, A.K. (Athanasios Kanavos); software, A.K. (Athanasios Kanavos) and G.P.-N.; validation, A.K. (Athanasios Kanavos), G.P.-N. and A.K. (Alexandros Kalokylos); formal analysis, A.K. (Athanasios Kanavos); investigation, A.K. (Athanasios Kanavos); resources, A.K. (Athanasios Kanavos); data curation, A.K. (Athanasios Kanavos) and G.P.-N.; writing—original draft preparation, A.K. (Athanasios Kanavos); writing—review and editing, A.K. (Athanasios Kanavos) and A.K. (Alexandros Kalokylos); visualization, A.K. (Athanasios Kanavos) and G.P.-N.; supervision, A.K. (Alexandros Kalokylos); project administration, A.K. (Alexandros Kalokylos). All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The simulation data and implementation details can be provided by the corresponding author upon request.

Acknowledgments: The authors acknowledge the use of ns-3, the 5G-LENA NR module, and SUMO in the development and evaluation of the scheduling techniques and the simulation framework.

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

Abbreviations

The following abbreviations are used in this manuscript:

3GPP	Third Generation Partnership Project
5GAA	5G Automotive Association
AI	Artificial Intelligence
AoI	Age of Information
AR	Accident Report
AS	Application Server
AVDR	Autonomous Vehicles Disengagement Report
BLSR	Bus Lane Sharing Request and Revoke
BSR	Buffer Status Report
C-V2X	Cellular Vehicle-to-Everything
CQI	Channel Quality Indicator
DL	Downlink
ECDF	Empirical Cumulative Distribution Function
eMBB	Enhanced Mobile Broadband
gNB	next-generation NodeB
HARQ	Hybrid Automatic Repeat reQuest
HDM	High-Definition Map Collection and Sharing
HO	Handover
ISD	Inter-Site Distance
ITS	Intelligent Transport Systems
KPI	Key Performance Indicator
LC	Logical Channel
LCG	Logical Channel Group
LTE	Long-Term Evolution
MAC	Medium Access Control
MCS	Modulation and Coding Scheme
MIMO	Multiple-Input Multiple-Output
ML	Machine Learning
MR	Maximum Rate
NR	New Radio
OVA	Obstructed View Assist
PF	Proportional Fair
PRB	Physical Resource Block
PRR	Packet Reception Rate
QoS	Quality of Service
RAN	Radio Access Network
RR	Round Robin
RTSA	Real-Time Situational Awareness
SUMO	Simulation of Urban Mobility
TBS	Transport Block Size
TDMA QoS	Time Division Multiple Access Quality of Service
ToD	Tele-Operated Driving
TTI	Transmission Time Interval
UC	Use Case
UE	User Equipment
UL	Uplink
UPF	User Plane Function

URLLC	Ultra-Reliable Low-Latency Communications
V2N	Vehicle-to-Network
V2X	Vehicle-to-Everything
VRU	Vulnerable Road User

References

1. Intelligent Transport Systems in the EU—European Climate, Infrastructure and Environment Executive Agency. Available online: https://cinea.ec.europa.eu/programmes/connecting-europe-facility/transport-infrastructure/intelligent-transport-systems-eu_en (accessed on 1 July 2026).
2. 5GAA. *A Visionary Roadmap for Advanced Driving Use Cases, Connectivity Technologies, and Radio Spectrum Needs*; 5GAA Automotive Association White Paper; 5GAA: Munich, Germany, 2022.
3. 5GAA TR S-200137; Working Group Standards and Spectrum Study of Spectrum Needs for Safety Related Intelligent Transportation Systems—Day 1 and Advanced Use Cases. 5GAA Automotive Association Technical Report. 5GAA: Munich, Germany, 2020.
4. 5GAA. *C-V2X Use Cases Volume II: Examples and Service Level Requirements*; 5GAA Automotive Association White Paper; 5GAA: Munich, Germany, 2020.
5. TR 22.886 V16.2.0; Technical Report 3rd Generation Partnership Project, Technical Specification Group Services and System Aspects, “Study on Enhancement of 3GPP Support for 5G V2X Services,” (Release 16). 3GPP: Biot, France, 2018.
6. Kanavos, A.; Fragkos, D.; Kaloxyllos, A. V2X communication over cellular networks: Capabilities and Challenges. *Telecom* **2021**, *2*, 1. <https://doi.org/10.3390/telecom2010001>.
7. Chougule, A.; Chamola, V.; Sam, A.; Yu, F.R.; Sikdar, B. A Comprehensive Review on Limitations of Autonomous Driving and Its Impact on Accidents and Collisions. *IEEE Open J. Veh. Technol.* **2024**, *5*, 142–161. <https://doi.org/10.1109/OJVT.2023.3335180>.
8. Kanavos, A.; Kaloxyllos, A. V2X Communications in Highway Environments: Scheduling Challenges and Solutions for 6G Networks. *Telecom* **2025**, *6*, 13. <https://doi.org/10.3390/telecom6010013>.
9. Baldo, N.; Miozzo, M.; Requena-Esteso, M.; Nin-Guerrero, J. An open source product-oriented LTE network simulator based on ns-3. In *Proceedings of the 14th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM '11)*; Association for Computing Machinery: New York, NY, USA, 2011; pp. 293–298. <https://doi.org/10.1145/2068897.2068948>.
10. Lopez, P.A.; Behrisch, M.; Bieker-Walz, L.; Erdmann, J.; Flötteröd, Y.-P.; Hilbrich, R.; Lücken, L.; Rummel, J.; Wagner, P.; Wießner, E. Microscopic Traffic Simulation using SUMO. In *Proceedings of the 21st International Conference on Intelligent Transportation Systems (ITSC 2018)*, Maui, HI, USA, 4–7 November 2018. <https://doi.org/10.1109/ITSC.2018.8569938>.
11. Haque, M.E.; Tariq, F.; Khandaker, M.R.A.; Wong, K.-K.; Zhang, Y. A Survey of Scheduling in 5G URLLC and Outlook for Emerging 6G Systems. *IEEE Access* **2023**, *11*, 34372–34396. <https://doi.org/10.1109/ACCESS.2023.3264592>.
12. Mamane, A.; Fattah, M.; El Ghazi, M.; El Bekkali, M.; Balboul, Y.; Mazer, S. Scheduling Algorithms for 5G Networks and Beyond: Classification and Survey. *IEEE Access* **2022**, *10*, 51643–51661. <https://doi.org/10.1109/ACCESS.2022.3174579>.
13. Nguyen, D.-H.; Nguyen, H.; Renault, E. E-MQS—A new downlink scheduler for real-time flows in LTE network. In *Proceedings of the 2016 IEEE 84th Vehicular Technology Conference (VTC-Fall)*, Montreal, QC, Canada, 18–21 September 2016; pp. 1–5. Available online: <https://ieeexplore.ieee.org/document/7880981/> (accessed on).
14. Saglam, M.I.; Kartal, M. 5G enhanced mobile broadband downlink scheduler. In *Proceedings of the 11th International Conference on Electrical and Electronics Engineering (ELECO)*, Bursa, Turkey, 28–30 November 2019; pp. 687–692. Available online: <https://ieeexplore.ieee.org/document/8990378/> (accessed on).
15. Kumar, P.; Singh, P. A Comprehensive Study of 5th Generation Scheduling Algorithms. In *New Approaches for Multidimensional Signal Processing, NAMSP; Smart Innovation, Systems and Technologies*; Springer: Singapore, 2022; Volume 332, pp. 188–201. https://doi.org/10.1007/978-981-19-7842-5_17.
16. Hakak, S.; Gadekallu, T.R.; Maddikunta, P.K.R.; Ramu, S.P.; Parimala, M.; De Alwis, C.; Liyanage, M. Autonomous vehicles in 5G and beyond: A survey. *Veh. Commun.* **2023**, *39*, 100551. <https://doi.org/10.1016/j.vehcom.2022.100551>.
17. 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>.
18. Husain, M.I.; Haque, M.E.; Tariq, F. An Efficient Packet Scheduling Algorithm for URLLC Systems. In *Proceedings of the 2020 International Conference on UK-China Emerging Technologies (UCET)*, Glasgow, UK, 20–21 August 2020. <https://doi.org/10.1109/UCET51115.2020.9205417>.

19. Haider, A.; Hwang, S.-H. Adaptive Transmit Power Control Algorithm for Sensing-Based Semi-Persistent Scheduling in C-V2X Mode 4 Communication. *Electronics* **2019**, *8*, 846. <https://doi.org/10.3390/electronics8080846>.
20. 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>.
21. Hapanchak, V.S.; Costa, A.; Pereira, J.; Nicolau, M.J. An intelligent path management in heterogeneous vehicular networks. *Veh. Commun.* **2024**, *45*, 100690. <https://doi.org/10.1016/j.vehcom.2023.100690>.
22. He, C.; Zhao, L.; Wan, Y.; Lu, H.; Shimizu, T. Sub-6 GHz V2X-assisted MmWave optimal scheduling for vehicular networks. *Veh. Commun.* **2023**, *41*, 100610. <https://doi.org/10.1016/j.vehcom.2023.100610>.
23. Nair, A.; Tanwar, S. Resource allocation in V2X communication: State-of-the-art and research challenges. *Phys. Commun.* **2024**, *64*, 102351. <https://doi.org/10.1016/j.phycom.2024.102351>.
24. Cao, Y.; Maghsudi, S.; Ohtsuki, T.; Quek, T.Q.S. Mobility-Aware Routing and Caching in Small Cell Networks Using Federated Learning. *IEEE Trans. Commun.* **2024**, *72*, 815–829.
25. Mannoni, V.; Berg, V.; Sesia, S.; Perraud, E. A Comparison of the V2X Communication Systems: ITS-G5 and C-V2X. In Proceedings of the 2019 IEEE 89th Vehicular Technology Conference (VTC2019-Spring), Kuala Lumpur, Malaysia, 28 April–1 May 2019. <https://doi.org/10.1109/VTCSpring.2019.8746562>.
26. TS 38.214 V19.3.0; Technical Specification 3rd Generation Partnership Project, Technical Specification Group Radio Access Network, “NR, Physical Layer Procedures for Data,” (Release 19). 3GPP: Biot, France, -2026.
27. TR 38.901 V19.4.0; Technical Report 3rd Generation Partnership Project, Technical Specification Group Radio Access Network, “Study on Channel Model for Frequencies from 0.5 to 100 GHz.” (Release 19). 3GPP: Biot, France, 06-2026.
28. Patriciello, N.; Lagen, S.; Giupponi, L.; Bojovic, B. An Improved MAC Layer for the 5G NR NS-3 module. In Proceedings of the 2019 Workshop on NS-3, Florence, Italy, 19–20 June 2019; pp. 41–48. <https://doi.org/10.1145/3321349.3321350>.
29. TR 38.913 V19.0.0; Technical Report 3rd Generation Partnership Project, Technical Specification Group Radio Access Network, “Study on Scenarios and Requirements for Next Generation Access Technologies,” (Release 19). 3GPP: Biot, France, 2025.
30. Kanavos, A.; Barmounakis, S.; Kaloxylas, A. An Adaptive Scheduling Mechanism Optimized for V2N Communications over Future Cellular Networks. *Telecom* **2023**, *4*, 378–392. <https://doi.org/10.3390/telecom4030022>.
31. TR 36.885 V14.0.0; Technical Report 3rd Generation Partnership Project, Technical Specification Group Radio Access Network, “Study on LTE-based V2X Services,” (Release 14). 3GPP: Biot, France, 2016.
32. Azim, A.W.; Bazzi, A.; Bomfin, R.; Chafii, M. System Level Analysis for Monostatic V2X ISAC. In *2026 IEEE International Conference on Communications Workshops (ICC Workshops)*; IEEE: New York, NY, USA, 2026.
33. TS 26.448 V19.0.0; Technical Specification 3rd Generation Partnership Project, Technical Specification Group Services and System Aspects, “Codec for Enhanced Voice Services (EVS); Jitter Buffer Management,” (Release 19). 3GPP: Biot, France, 2025.
34. Papanikolaou-Ntais, G.; Sotiropoulos, D.N.; Kanavos, A.; Kaloxylas, A. Deadline-Aware Scheduler-Weight Adaptation for 5G NR V2X Networks Using Probabilistic Prediction and Reinforcement Learning. *Telecom* **2026**, *7*, 80. <https://doi.org/10.3390/telecom7040080>.

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.