Università degli Studi di Modena e Reggio Emilia

Dipartimento di Ingegneria "Enzo Ferrari"

International Doctorate in
Information and Communication Technologies

XXXV Cycle

Vehicular Connectivity in 5G and Beyond

Supervisor:
Prof. Maria Luisa Merani

Candidate:
Luca Lusvarghi

PhD School Coordinator:
Prof. Sonia Bergamaschi

A Giorgio, ciao Nonno

"In the beginner's mind there are many possibilities, in the expert's mind there are few."
初心

Shunryu Suzuki
Zen Mind, Beginner's Mind

"Heaven is not a place, and is not a time. Heaven is being perfect.
You will begin to touch heaven in the moment that you touch perfect speed.
And that isn't flying a thousand miles an hour, or a million, or flying at the speed of light. Because any number is a limit, and perfection doesn't have limits.
Perfect speed is being there."

Richard Bach
Jonathan Livingston Seagull

Acknowledgments

First, I would like to thank my supervisor, Prof. Maria Luisa Merani, for giving me the opportunity to pursue the PhD and for supporting me with valuable advice. Next, I want to thank Javier Gozalvez, Baldomero Coll-Perales, and the rest of the UWICORE group at the Universidad Miguel Hernandez (UMH) de Elche. Thank you for the insightful comments, the fruitful collaboration, and your hospitality. I also want to acknowledge the interesting discussions and the joyful memories that I shared with my lab colleagues over the last four years.

A special thank you goes to my group of close friends "OFFESE AGGRATIS", to Andrea, and to Luigi. Your company, smiles, and inspiring comments have supported and motivated me on a daily basis. Last, I want to express my deepest gratitude to my family. Thank you Claudia, Stefano, and Alice for your unconditional love.

Contents

Introduction **1**

1 MoReV2X: Simulation of C-V2X Sidelink Communications **7**

1.1 Research Motivation and Main Contribution 8
1.2 System-Level and Link-Level Simulation Tools 9
1.3 Related Work 10
1.4 Overview of C-V2X SL Communications 11
1.4.1 LTE-V2X SL 11
1.4.1.1 PHY Layer 12
1.4.1.2 MAC Sublayer 13
1.4.2 NR-V2X SL 14
1.4.2.1 PHY Layer 15
1.4.2.2 MAC Sublayer 17
1.5 Evaluation Guidelines 19
1.5.1 Traffic Models 19
1.5.2 Channel Models 21
1.5.3 Performance Metrics 23
1.6 Ns-3 Implementation 24
1.6.1 Application Layer 25
1.6.2 MAC Sublayer 27
1.6.3 PHY Layer 27
1.6.4 Exemplary Results 29
1.7 Link-Level Evaluation of NR-V2X SL Communications 31
1.7.1 NR-V2X SL Link-Level Simulator 32
1.7.1.1 Transmitter Chain 33
1.7.1.2 Channel Model 36
1.7.1.3 Receiver Chain 37
1.7.2 NR-V2X SL Profile 38
1.7.3 Relevance of the NR-V2X SL LL Results 40
1.7.3.1 Decoding of 1st-stage SCI, 2nd-stage SCI and TB 40
1.7.3.2 Tx-Rx Relative Speed 41
1.7.3.3 Channel Models 44

1.7.3.4 MCS Selection . . . 45
1.7.3.5 Number of Subchannels . . . 46
1.8 Conclusions . . . 47

2 Assessing the Achievable Performance of C-V2X Sidelink Communications 49
2.1 Research Motivation and Main Contribution . . . 49
2.2 Related Work . . . 51
2.3 Coexistence of Periodic and Aperiodic Traffic in LTE-V2X Mode 4 . . . 52
2.3.1 Serving Aperiodic Traffic . . . 53
2.3.1.1 The Proposed Solution . . . 53
2.3.1.2 Limiting Analysis . . . 53
2.3.2 Simulation Environment . . . 55
2.3.3 Numerical Results . . . 56
2.4 NR-V2X Mode 2: MAC Analysis . . . 63
2.4.1 Impact of Different Traffic Types . . . 64
2.4.2 Impact of the PDB . . . 66
2.4.3 Analysis of the Re-Evaluation Mechanism . . . 67
2.4.3.1 Re-Evaluation Check . . . 67
2.4.3.2 Re-Evaluation Detection . . . 68
2.4.3.3 Resource Replacement . . . 74
2.5 NR-V2X Mode 2: Numerical Results . . . 74
2.5.1 Simulation Environment . . . 74
2.5.1.1 Traffic Models . . . 75
2.5.1.2 Performance Metrics . . . 76
2.5.2 Impact of Re-Evaluations . . . 77
2.5.2.1 SPS Strategy Without Retransmissions ($N = 1$) . . . 77
2.5.2.2 SPS Strategy With Retransmissions ($N = 2$) . . . 81
2.5.2.3 DS Strategy . . . 83
2.5.3 Comparative Analysis of the SPS and DS Strategies . . . 85
2.5.3.1 Impact of the PDB on the SPS and DS Strategies . . . 86
2.5.3.2 SPS and DS Performance Comparison: Single Model . . . 88
2.5.3.3 SPS and DS Performance Comparison: Mixed Models and Adaptive Scheduling . . . 91
2.6 Conclusions . . . 93

3 Machine Learning for Disseminating Awareness Messages 95
3.1 Research Motivation and Main Contribution . . . 96
3.2 Related Work . . . 97
3.3 ETSI-Generated CAM Sequences . . . 98
3.4 The Predictive Reservation Framework . . . 99

3.4.1 Modified SPS Implementation . . . 99
3.4.2 Machine Learning to Predict CAM Sequences . . . 101
3.5 Numerical Results . . . 103
3.5.1 PHY and MAC Layer Configuration . . . 103
3.5.2 Outcomes . . . 103
3.5.2.1 Suburban setting . . . 103
3.5.2.2 Highway setting . . . 110
3.6 Conclusions and Future Work . . . 113

4 Awareness Messages by VRUs and Vehicles: Field Tests via LTE-V2X 115
4.1 Research Motivation and Main Contribution . . . 115
4.2 Related Work . . . 117
4.3 Vulnerable Road Users . . . 118
4.3.1 VRU definition, features, and requirements . . . 118
4.3.2 VRU Awareness Messages . . . 119
4.3.2.1 VAM Generation . . . 119
4.3.2.2 VAM Format . . . 120
4.4 Field Tests . . . 122
4.4.1 LTE-V2X Boards . . . 122
4.4.2 Bicycle and E-scooter VAMs . . . 123
4.4.3 Vehicle and Motorcycle CAMs . . . 128
4.4.4 LTE-V2X SL: PRR Analysis . . . 131
4.4.5 Discussion . . . 133
4.4.5.1 VAM and CAM generation times . . . 133
4.4.5.2 PRR of LTE-V2X SL Communications . . . 134
4.5 Conclusions and Future Work . . . 134

5 Fundamental Limits on the Performance of PD-NOMA 137
5.1 Research Motivation and Main Contribution . . . 137
5.2 Related Work . . . 139
5.3 Performance Analysis . . . 140
5.3.1 System Model and Performance Metric Evaluation . . . 140
5.3.2 $P_{out}^{(j)}$ Approximation for $j \geq 2$ . . . 144
5.4 Fading Models . . . 146
5.4.1 Rayleigh Fading . . . 146
5.4.2 Rayleigh-lognormal Shadowed Fading . . . 147
5.5 Numerical Results . . . 148
5.6 Conclusions and Future Work . . . 153

6 Final Remarks 159

List of Figures

1.1 Resource reselection mechanism. 12
1.2 Subframe structure in LTE-V2X SL. 13
1.3 Slot structure in NR-V2X SL when $N_{\mathrm{DMRS}} = 2$. 16
1.4 Re-evaluation mechanism in NR-V2X Mode 2. 19
1.5 Average pathloss plus shadowing contribution ($PL + SH$) in the Highway and Urban propagation environments for different channel states. 22
1.6 Simplified Unified Modeling Language (UML) diagram of the MoReV2X implementation. 25
1.7 PRR as a function of the Tx-Rx distance D. 30
1.8 PIR CDF, $D_{max} = 520$ m. 30
1.9 PLR and CLR curves reported as a function of the Tx-Rx distance D. . . . 30
1.10 NR-V2X SL LL simulator workflow. 33
1.11 Transmitter chain. 36
1.12 Receiver chain. 38
1.13 1^{st}-stage SCI, 2^{nd}-stage SCI and TB: BLER vs SNR curves for different MCS choices. 42
1.14 Impact of the Tx-Rx relative speed on the TB BLER performance. 43
1.15 Impact of the Tx-Rx relative speed on the 2^{nd}-stage and 1^{st}-stage SCI BLER performance. 44
1.16 Comparison of channel models: BLER vs SNR curves. 45
1.17 Impact of the MCS on the TB BLER performance. 46
1.18 Impact of N_{sub} on the TB BLER performance. 47

2.1 PRR as a function of the Tx-Rx distance D, $X = 1000$ bytes. 56
2.2 PRR as a function of D, $X = 190$ bytes. 58
2.3 Aperiodic and periodic PRR as a function of D. 59
2.4 PRR for different aperiodic traffic models. 59
2.5 Comparison among alternative schemes, $\Theta = 50\%$. 60
2.6 Comparison among alternative schemes, $\Theta = 100\%$ and $X = 1000$ bytes. . . 61
2.7 PRR for different values of the PDB. 61
2.8 PRR when $\Theta = 100\%$. 62
2.9 Average PIR of aperiodic and periodic traffic as a function of D. 62

2.10 Aperiodic and periodic PIR CDF, $X = 1000$ bytes. 63
2.11 Collision probability associated with resource reselections. 64
2.12 Impact of different traffic types on NR-V2X Mode 2. 65
2.13 Persistent collisions not detected by re-evaluation when using SPS. 69
2.14 Re-evaluation detection when vehicles use the same RRI. 70
2.15 Re-evaluation detection when vehicles use different RRIs. 71
2.16 Impact of retransmissions on the re-evaluation detection. 72
2.17 AV traffic, $N = 1$: SPS performance in the mixed traffic scenario. 79
2.18 SPS strategy, single traffic scenario, $N = 1$: PRR curves for PF and AV traffic (minimum RRI strategy). Similar trends are observed in the mixed traffic scenario and with the average RRI strategy. 80
2.19 SPS strategy, PF traffic, mixed traffic scenario, $N = 1$: CLR and PRR-Re-evaluation curves obtained in the 100 veh/km setting. Similar trends are observed in the single traffic scenario. 81
2.20 SPS strategy (average RRI strategy), AV traffic, $N = 2$: PRR-Re-evaluation as a function of D in the 50 veh/km setting. Similar trends have been obtained for other densities. 82
2.21 AV traffic, $N = 2$, 50 veh/km: PRR performance of the SPS (average RRI strategy) and DS schemes in the single traffic scenario. Similar trends have been obtained for other densities. 83
2.22 DS strategy, AV traffic, single traffic scenario, $N = 2$: CLR as a function of D for two different vehicular densities. 85
2.23 SPS strategy: CLR as a function of D. 86
2.24 DS strategy: CLR as a function of D. 86
2.25 Single model scenario, PF traffic, Default PDB: PRR as a function of D. . . 88
2.26 Single model scenario, AF traffic: PRR as a function of D. 89
2.27 Single model scenario, AV traffic: PRR as a function of D. 90
2.28 Mixed models, 120 veh/km, $\Delta = 10, 25, 50$, and 75%: PRR as a function of D. 92

3.1 Correlation between vehicular speed and T_{CAM}. 98
3.2 SPS detection of potential collisions for different RRI_{TX} and RRI_{RX} settings. 100
3.3 Flowchart showing the proposed KNN-look ahead solution. 102
3.4 The examined suburban road topology. 103
3.5 Confusion matrix for the prediction of the T_{CAM_i} values. 104
3.6 Accuracy and macro-F1 score as a function of the T_{CAM} index. 105
3.7 PLR as a function of the Tx-Rx distance D. 106
3.8 PRR as a function of the Tx-Rx distance D, suburban scenario. 106
3.9 PRR comparison for different L_1 lists, suburban scenario. 107
3.10 CLR as a function of D, suburban scenario. 108
3.11 T_{CAM} PMF, suburban scenario. 109

3.12 PRR comparison of the proposed KNN-look ahead approach against SPS with three different RRI settings, suburban scenario. 109
3.13 PIR CDF, suburban scenario. 110
3.14 PRR as a function of D, highway scenario, CAM size $X = 470$ bytes. 111
3.15 PRR and CLR as a function of D, highway scenario, CAM trace Markov model. 111
3.16 T_{CAM} PMF, highway scenario. 112
3.17 PIR CDF, CAM size $X = 470$ bytes, highway scenario. 113

4.1 VAM structure. 121
4.2 Quectel evaluation board used for the field measurements. 122
4.3 Bicycle: VAM trace. 124
4.4 Bicycle: VAM triggers and T_{VAM} PMF for $\Delta_h = 4°$. 125
4.5 E-scooter: VAM triggers and T_{VAM} PMF for $\Delta_h = 4°$. 125
4.6 Bicycle: GNSS coordinates when heading-triggered VAMs are generated. . . 126
4.7 Bicycle: VAM triggers and T_{VAM} PMF for $\Delta_h = 10°$. 127
4.8 E-scooter: VAM triggers and T_{VAM} PMF for $\Delta_h = 10°$. 127
4.9 The examined scenarios. 128
4.10 Car: CAM triggers and T_{CAM} PMF, *urban* scenario. 129
4.11 Motorcycle: CAM triggers and T_{CAM} PMF, *urban* scenario. 129
4.12 Car: CAM triggers and T_{CAM} PMF, *suburban* scenario. 130
4.13 Car: CAM triggers and T_{CAM} PMF, *highway* scenario. 130
4.14 LTE-V2X SL PRR: employed equipment. 131
4.15 Bicycle-to-Vehicle communications: PRR. 132
4.16 V2V communications: PRR. 133

5.1 Users location along the cell radius when $n = 2, 3, 5$. 149
5.2 Rayleigh fading: $P_{out}^{(1)}$ as a function of the SNR, $n = 2, 3, 5$. 149
5.3 Rayleigh fading and log-normal shadowing: $P_{out}^{(1)}$ as a function of the SNR, $n = 2, 3, 5$. 150
5.4 $P_{out}^{(1)}$, Rayleigh fading and Rayleigh plus lognormal, $n = 3$. 150
5.5 Outage probability vs SNR, $n = 3$. 151
5.6 Outage probability as a function of D_2/R, $n = 3$, SNR = 30 dB. 152
5.7 Sum data rate of OMA and NOMA as a function of SNR, Rayleigh and lognormal shadowing. 152

List of Abbreviations

3GPP	3rd Generation Partnership Project
AI	Artificial Intelligence
BLER	BLock Error Rate
C-V2X	Cellular - V2X
CAM	Cooperative Awareness Message
CBR	Channel Busy Ratio
CDL	Clustered Delay Line
CLR	Collision Losses Ratio
CRR	Counter Reselections Ratio
DL	Downlink
DMRS	DeModulation Reference Signal
DS	Dynamic Scheduling
ETSI	European Telecommunications Standardization Institute
eV2X	enhanced V2X
GNSS	Global Navigation Satellite System
ITS	Intelligent Transportation Systems
LL	Link - Level
LOS	Line - Of - Sight
LUT	Look - Up Table
LRR	Latency Reselections Ratio
MAC	Medium Access Control
MCS	Modulation and Coding Scheme
ML	Machine Learning
NLOS	Non - Line - Of - Sight
NOMA	Non - Orthogonal - Multiple - Access
NR-V2X	New Radio - V2X
OFDM	Orthogonal Frequency Division Multiplexing
PD	Power - Domain
PDB	Packet Delay Budget
PHY	Physical
PIR	Packet Inter-Reception
PLR	Propagation Losses Ratio

PMF	Probability Mass Function
PRR	Packet Reception Ratio
PSCCH	Physical Sidelink Control CHannel
PSSCH	Physical Sidelink Shared CHannel
RB	Resource Block
RE	Resource Element
RRI	Resource Reservation Interval
RSRP	Reference Signal Received Power
SCI	Sidelink Control Information
SCS	SubCarrier Spacing
SIC	Successive Interference Cancellation
SINR	Signal-to-Interference-plus-Noise Ratio
SIR	Signal-to-Interference Ratio
SL	Sidelink
SNR	Signal-to-Noise Ratio
SPS	Semi-Persistent Scheduling
SRR	Size Reselections Ratio
TB	Transport Block
UE	User Equipment
UL	Uplink
URR	Unutilized Reservations Ratio
V2V	Vehicle-to-Vehicle
V2X	Vehicle-to-Everything
VAM	VRU Awareness Message
VRU	Vulnerable Road User

Abstract

According to the World Health Organization (WHO), 1.3 million people die every year as a result of road traffic accidents caused by human errors. While more severe traffic regulations and a safer road infrastructure design would substantially reduce the number of casualties, Vehicle-to-Everything (V2X) communications will play a crucial role in improving road safety and transportation efficiency. To this end, the Third Generation Partnership Project (3GPP) has introduced the first V2X cellular technology, known as LTE-V2X, in Release 14 specifications. LTE-V2X has been designed to support basic safety-related applications, such as the dissemination of awareness messages in out-of-coverage scenarios, where vehicles directly communicate without the support of the cellular infrastructure. In recent years, the advent of more sophisticated V2X use cases that rely on cooperative perception and maneuvering has prompted the development of the 5G New Radio (NR)-V2X technology. Standardized by 3GPP in Release 16, NR-V2X is characterized by a new physical (PHY) layer design and new Medium Access Control (MAC) features expected to guarantee improved robustness and flexibility.

System-level simulations are instrumental in assessing the large scale performance of LTE-V2X and NR-V2X networks, as real-world tests and measurements are often impracticable due to the limited availability and the high cost of hardware prototypes. This thesis introduces MoReV2X, an open-source ns-3 module for the simulation of LTE-V2X and NR-V2X communications. MoReV2X adheres to 3GPP specifications and features an accurate PHY layer abstraction model based on BLock Error Rate (BLER) curves obtained through a detailed link-level analysis. The MoReV2X simulator has been employed to thoroughly investigate the behavior of LTE-V2X and NR-V2X from different perspectives. To begin with, the coexistence of periodic and aperiodic traffic in LTE-V2X has been studied, focusing on the limitations which characterize its distributed resource allocation strategy. Then, the impact of aperiodic traffic dissemination on system performance has been further analyzed in the NR-V2X domain, concentrating on the effectiveness of the new MAC features introduced in Release 16.

As simulation results revealed that LTE-V2X and NR-V2X are not able to effectively disseminate aperiodic traffic, a novel Artificial Intelligence (AI)-based strategy to broadcast awareness messages has been put forth. According to it, each vehicle forecasts its message generation times and optimizes the MAC layer configuration. Exploiting AI predictions, the proposed approach outperforms legacy solutions in all respects.

Besides simulation studies, this thesis has analyzed the results of an extensive measurement campaign which investigated the generation of awareness messages in a real-world context. Field tests have been performed with connected Vulnerable Road Users (VRUs) and cars in the urban, suburban, and highway scenarios. As the findings indicate that many messages from VRUs were generated under non-relevant circumstances, this study has proposed an adjustment to reduce their dissemination frequency without missing relevant information about the VRU movements.

Last, this thesis presents a novel analytical approach to evaluate the outage probability that Power-Domain (PD) Non-Orthogonal Multiple Access (NOMA) achieves on the uplink. The fundamental limits of system performance are analytically assessed for an arbitrary number of simultaneously transmitting vehicles, and both the case of Rayleigh and lognormal-shadowed Rayleigh fading are examined. The obtained closed-form expressions disclose the potential of PD-NOMA in beyond 5G V2X communications.

Introduction

Present days witness an increased and widespread sensitivity to road safety and sustainable transports. To this end, great promise is shown by the advent of vehicular connectivity, which allows vehicles to exchange real-time information with their peers, with Vulnerable Road Users (VRUs), and with the roadside infrastructure. Collectively referred to as Vehicle-to-Everything (V2X), vehicular communications extend the vehicles' information horizon beyond the Line-Of-Sight (LOS) range of local sensors, e.g., cameras, LIDARs, and RADARs, granting the car and its driver an omniscient view of the surrounding environment. The onset of vehicular networking represents a major turning point, as it paves the way for the evolution of Intelligent Transportation Systems (ITS) towards connected and automated driving, where the vision of more secure, efficient, and environment-friendly transports will be accomplished.

The idea of leveraging vehicular connectivity in the ITS domain is not new, but dates back to the early 1980s [1], when a handful of projects across the US and Japan focused on the preliminary investigation of Inter-Vehicle Communications (IVC). Thanks to the development of affordable Global Navigation Satellite System (GNSS) receivers and radio transceivers, vehicular communications gained considerable momentum during the 1990s. Between 1986 and 1997, the US-based PATH project [2] begun to consider vehicular communications as a tool for increasing highway capacity and road safety while reducing traffic congestion and energy consumption. In Europe, several V2X-oriented projects were funded within the PROMETHEUS framework [3] between 1986 and 1995. PROMETHEUS was the largest project ever in the field of connected and autonomous driving. Then, the mobile ad-hoc networking paradigm was extended to the vehicular domain in the early 2000s, introducing the concept of Vehicular Ad-hoc NETwork (VANET). Under the VANET umbrella, the research community devoted significant efforts to the development of Vehicle-to-Vehicle (V2V) communications. V2V communications allow the direct data exchange among vehicles, without the cellular infrastructure orchestration, representing the baseline solution for the support of safety-critical applications. However, no dedicated technology had been developed for V2V communications until 2004, when the IEEE 802.11p task group was formed.

IEEE 802.11p is an amendment to the well-known IEEE 802.11 suite of standards, designed to allow wireless access in vehicular environments. IEEE 802.11p was the first V2X technology and it lied the foundations for the definition of the Dedicated Short Range

Communications (DSRC) and ITS-G5 standards in the US and in Europe, respectively. More recently, IEEE 802.11p has been paired by the Cellular V2X (C-V2X) Sidelink (SL) technology. Introduced by the Third Generation Partnership Project (3GPP) in Release 14 specifications (2017) [4], C-V2X SL is an adaption of traditional cellular communications to the vehicular domain. In C-V2X SL, vehicles employ a dedicated air interface to directly communicate, leveraging a custom set of Medium Access Control (MAC) sublayer and physical (PHY) layer features.

Both the IEEE 802.11p and C-V2X SL technologies target the ITS band centered at 5.9 GHz, but they are not interoperable. The lack of interoperability raised an intense debate on the technical, regulatory, and economical aspects related to the adoption of IEEE 802.11p and C-V2X SL. As of today, the debate has not settled. On one hand, IEEE 802.11p (and its evolution, IEEE 802.11bd) is considered a mature technology after several years of large-scale evaluation and testing. On the other hand, C-V2X SL is a more recent technology that has not been properly investigated yet.

In this regard, the goal of this thesis is to provide an exhaustive and accurate analysis of the C-V2X SL technology, investigating it from three different perspectives: simulative, experimental, and analytical.

Simulations are key in assessing the large-scale performance of C-V2X SL networks, as real-world tests and measurements are often impracticable due to the limited availability and the high cost of hardware prototypes. Moreover, simulations are instrumental to the development and evaluation of new solutions before their standardization and deployment. In this thesis, simulations are leveraged to thoroughly analyze the MAC sublayer scheduling schemes and the PHY layer features of the C-V2X SL technology. At MAC sublayer, simulations are also employed to assess the impact of a new Artificial Intelligence (AI)-based resource allocation strategy on the performance of C-V2X SL.

When a confined number of connected vehicles is considered, real-world tests and experiments can be profitably employed to validate simulation results and overcome the intrinsic limitations of a purely simulative approach. In this thesis, numerical simulations are complemented with the results of an experimental campaign that involved bicycles, e-scooters, motorbikes, and cars, conducted to analyze the dissemination of awareness messages via C-V2X SL boards. Awareness messages include basic information such as the position, speed, heading, and type of the transmitting road user, representing the fundamental elements to build safety-related V2X applications.

In addition to the simulative and the experimental approach, this thesis also investigates the C-V2X SL behavior from an analytical perspective. First, the impact of more stringent latency requirements on the performance of C-V2X SL is determined. The analytical results are successfully validated through numerical simulations. Then, this thesis puts forth a novel analytical approach to evaluate the performance of Power-Domain (PD) Non-Orthogonal Multiple Access (NOMA). PD-NOMA is a novel "beyond 5G" access technique able to accommodate multiple users on the same frequencies by carefully assigning different power levels. Although the adoption of PD-NOMA is usually confined to cellular

systems, this thesis discusses its potential in V2V applications.

Further details about the logical organization and the contributions of this thesis are reported below.

Implementation of a standard-compliant C-V2X SL simulator

The use of simulations plays a crucial role in the large-scale assessment of the C-V2X SL technology, providing valuable insights on its strengths and limitations before the deployment phase. This thesis introduces MoReV2X, an open-source ns-3 module for the simulation of C-V2X SL communications. Following the 3GPP evaluation guidelines summarized in [5], MoReV2X includes a specific set of V2V SL channel models, realistic traffic models, and evaluates a complete set of performance metrics. The implementation of MoReV2X adheres to 3GPP specifications from Release 14 up to Release 16 [6], concentrating on the LTE-V2X SL and New Radio (NR)-V2X SL technologies, which represent the 4G and 5G versions of the C-V2X SL standard, respectively. In this thesis, the main features of MoReV2X are reported.

This thesis also concentrates on the link-level analysis of C-V2X SL communications, assessing the impact of typical PHY layer aspects, e.g., variable Modulation and Coding Scheme (MCS) and transmitter-receiver relative speed, on system performance. The relevance of the obtained link-level results to the implementation of MoReV2X is also discussed.

The MoReV2X simulator and the C-V2X SL link-level analysis have been presented in the following works:

- [7] L. Lusvarghi and M. L. Merani, "MoReV2X - A New Radio Vehicular Communication Module for ns-3," *2021 IEEE 94th Veh. Technol. Conf. (VTC2021-Fall)*, 2021, pp. 1-7.

- [8] L. Lusvarghi, B. Coll-Perales, J. Gozalvez and M. L. Merani, "Link Level Evaluation of 5G NR V2X Sidelink Communications," submitted to *IEEE Transactions on Vehicular Technology*.

Performance Assessment of C-V2X SL

Leveraging the MoReV2X simulator, this thesis provides an accurate and exhaustive analysis of LTE-V2X Mode 4 and NR-V2X Mode 2, i.e., the distributed resource allocation modes employed by vehicles to autonomously select radio resources in C-V2X SL communications, without coordination from the infrastructure.

First, the key question of how to accommodate aperiodic traffic in LTE-V2X Mode 4 is addressed, and a novel strategy to jointly serve periodic and aperiodic data flows is put forth. The obtained results provide numerical evidence of the superiority of the proposed strategy, revealing that LTE-V2X Mode 4 is not able to effectively serve aperiodic traffic. As the design of LTE-V2X Mode 4 is tailored to the dissemination of periodic traffic,

the generation of aperiodic packets hampers the proper functioning of the MAC sublayer, deteriorating system performance.

Then, the MoReV2X simulator is employed to analyze the performance of NR-V2X Mode 2. NR-V2X Mode 2 has been designed to complement LTE-V2X Mode 4 and support more demanding V2X use cases, characterized by aperiodic traffic patterns and more stringent latency, reliability, and communication range requirements. Despite the introduction of new MAC sublayer features and a more flexible PHY layer design, numerical simulations reveal that aperiodic traffic can not be effectively served in NR-V2X Mode 2 either. In particular, simulation results show that the MAC sublayer of NR-V2X SL faces the same challenges which characterize LTE-V2X SL communications.

The C-V2X SL analysis presented in this thesis is based on the results reported in following contributions:

- [9] L. Lusvarghi and M. L. Merani, "On the Coexistence of Aperiodic and Periodic Traffic in Cellular Vehicle-to-Everything," *IEEE Access*, vol. 8, pp. 207076-207088, 2020.

- [10] A. Molina-Galan, B. Coll-Perales, L. Lusvarghi, J. Gozalvez and M. L. Merani, "How does 5G NR V2X Mode 2 Handle Aperiodic Packets and Variable Packet Sizes?," *2022 IEEE 23rd International Conference on High Performance Switching and Routing (HPSR)*, 2022, pp. 183-188.

- [11] A. Molina-Galan, L. Lusvarghi, B. Coll-Perales, J. Gozalvez and M. L. Merani, "On the Impact of Re-evaluation in 5G NR V2X Mode 2," submitted to *IEEE Transactions on Vehicular Technology*.

- [12] L. Lusvarghi, A. Molina-Galan, B. Coll-Perales, J. Gozalvez and M. L. Merani, "A Comparative Analysis of the Semi-Persistent and Dynamic Scheduling Schemes in NR-V2X Mode 2," submitted to *Elsevier Vehicular Communications*.

Design of an AI-based resource allocation strategy

In order to address the challenges related to the dissemination of aperiodic traffic in LTE-V2X Mode 4 and NR-V2X Mode 2, this thesis puts forth an AI-based resource allocation strategy to broadcast awareness messages. As demonstrated by the experimental results presented in [13, 14], awareness messages are characterized by an aperiodic profile, as the inter-arrival time between consecutive messages depends on the transmitting vehicle's dynamics. The proposed approach relies on Machine Learning (ML) to forecast the temporal pattern of awareness messages and accordingly tailor the MAC sublayer configuration. Numerical simulations indicate that ML achieves an excellent accuracy in predicting the inter-arrival time between awareness messages. Moreover, an exhaustive set of simulations performed using the MoReV2X simulator shows that the proposed ML-enhanced resource allocation strategy outperforms LTE-V2X Mode 4 under all points of view, guaranteeing a

reliable and effective dissemination of awareness messages. Although the simulation phase concentrated on LTE-V2X Mode 4, this thesis also discusses the impact of ML on NR-V2X Mode 2 performance.

The application of ML to the distribution of awareness messages is explored in the following work:

- [15] L. Lusvarghi and M. L. Merani, "Machine Learning for Disseminating Cooperative Awareness Messages in Cellular V2V Communications," *IEEE Transactions on Vehicular Technology*, vol. 71, no. 7, pp. 7890-7903, July 2022.

Dissemination of awareness messages via LTE-V2X SL

In this thesis, the dissemination of awareness messages via LTE-V2X SL has also been analyzed from an experimental perspective. Real-world field tests have been performed employing LTE-V2X SL prototype boards to address two different tasks: (i) accurately characterize the traffic generated by bicycles, e-scooters, motorbikes, and vehicles; (ii) assess the communication range of direct vehicle-to-vehicle and bicycle-to-vehicle communications.

With reference to the former goal, the inter-arrival time between consecutive awareness messages was experimentally determined, leveraging the results of a measurement campaign conducted in three different scenarios, namely: urban, suburban, and highway. As the obtained results revealed that many awareness messages originate under non-relevant conditions, especially when bicycles and e-scooters are considered, this thesis puts forth an adjustment to the current generation rules standardized by ETSI in [16, 17]. Experimental results show that the proposed amendment is able to avoid generating too many awareness messages without missing relevant information about the bicycle and e-scooter movements.

With reference to the latter goal, the experimental assessment of the vehicle-to-vehicle and the bicycle-to-vehicle communication range provided valuable insights on the coverage attained by LTE-V2X SL safety applications, complementing the conclusions drawn from numerical simulations. In particular, the performance of vehicle-to-vehicle communications revealed more pessimistic results with respect to simulation studies and other experimental campaigns performed in controlled environments.

Further details about the experimental results presented in this thesis can be found in the following contribution:

- [13] L. Lusvarghi, C. A. Grazia, M. Klapez, M. Casoni and M. L. Merani, "Awareness Messages by Vulnerable Road Users and Vehicles: Field Tests via LTE-V2X," submitted to *IEEE Transactions on Intelligent Vehicles*.

Analytical models of PD-NOMA

With the advent of V2X, the number of connected vehicles is expected to exponentially increase over the next years. Connected vehicles will generate an increasingly larger amount

of data to support a wide range of connected and automated driving applications. To sustain the ever growing demand for connectivity and to guarantee the coexistence between conventional radio User Equipments (UEs) and connected vehicles, novel Next Generation Multiple Access (NGMA) solutions have been proposed within the 6G framework. The goal of NGMA is to efficiently connect a significantly larger number of devices over the same portion of spectrum. For example, PD-NOMA proposes to accommodate multiple users (UEs and V2X-enabled vehicles) on the same frequencies, by carefully assigning them different power levels and employing Successive Interference Cancellation (SIC) receivers.

This thesis puts forth a novel analytical approach to evaluate the performance that PD-NOMA achieves on the uplink of a single cell when a dynamic-ordered SIC receiver is considered. The fundamental limits on the system performance are assessed analytically for an arbitrary number of simultaneously transmitting users, examining both the case of Rayleigh and lognormal-shadowed Rayleigh fading. The correctness and the excellent accuracy of the obtained closed-form expressions are validated through Monte Carlo simulations.

The analytical approach presented in this thesis is based on the results reported in the following contribution:

- [18] L. Lusvarghi and M. L. Merani, "Fundamental Limits on the Uplink Performance of the Dynamic-Ordered SIC Receiver," *IEEE Access*, vol. 10, pp. 73178-73189, 2022.

The outline of this thesis is reported next. Chapter 1 introduces the MoReV2X simulator, providing an overview of its features and reporting a preliminary set of numerical results. Furthermore, this Chapter includes the link-level analysis of C-V2X SL and discusses the relevance of the corresponding outcomes in the MoReV2X implementation. Chapter 2 leverages a comprehensive set of results obtained using the MoReV2X simulator to assess the achievable performance of C-V2X SL communications, from LTE-V2X SL to NR-V2X SL. Then, Chapter 3 presents the AI-enhanced resource allocation strategy to broadcast awareness messages, demonstrating its superiority with respect to the legacy solution through simulation. The dissemination of awareness messages via LTE-V2X SL is experimentally investigated in Chapter 4, where the results of the field tests campaign are discussed. Chapter 5 introduces the analytical approach proposed to evaluate the performance of PD-NOMA and, lastly, Chapter 6 draws the conclusions.

Chapter 1

MoReV2X: Simulation of C-V2X Sidelink Communications

As the evolution of ITSs has progressively shifted towards connected and automated driving, V2X communications have attracted a great deal of attention from both the academic and the industrial world, emerging as a key technology to improve road safety and transportation efficiency. To this end, the Third Generation Partnership Project (3GPP) devoted significant efforts to develop C-V2X SL communications. C-V2X SL communications rely on the PC5 air interface defined in Release 12 specifications and allow the direct data exchange between vehicles. The definition of the PC5 interface paved the way for the standardization of the first C-V2X technology in Release 14 [4], known as LTE-V2X SL. The design of LTE-V2X SL has been mainly oriented towards the support of basic safety and traffic management applications, such as the dissemination of Cooperative Awareness Messages (CAMs), in both in-coverage and out-of-coverage scenarios. To do so, LTE-V2X SL introduced two different modes for the assignment of radio resources to the vehicles, namely Mode 3 and Mode 4. Mode 3 is a centralized approach in which vehicles delegate the selection of collision-free radio resources to the evolved Node B (eNB). However, the eNB central orchestration requires all vehicles to be under cellular coverage, a condition clashing with the requirements of safety applications, that cannot depend on the availability of such coverage. On the other hand, in Mode 4 vehicles autonomously select radio resources for their V2V communications, without network assistance, thus making Mode 4 the baseline solution for safety applications.

Recently, V2X services evolved from basic safety-related functionalities to enhanced V2X (eV2X) applications [19] that rely on cooperative perception and maneuvering. The wide range of demanding latency, reliability, and data rate requirements that eV2X use cases will exhibit prompted the definition of 5G NR-V2X SL communications. Standardized by 3GPP within Release 16 [6], NR-V2X SL is characterized by a new PHY layer design and new MAC sublayer strategies expected to guarantee an increased flexibility in the management of resources. At MAC sublayer, 3GPP has introduced a new mandatory re-evaluation mechanism and two new distributed resource allocation modes, Mode 1 and

Mode 2, in order to complement, and not to replace, legacy LTE-V2X Mode 3 and Mode 4, respectively. However, the definition of a more flexible PHY and MAC layer design in Release 16 does not allow the compatibility between NR-V2X SL and LTE-V2X SL. For this reason, 3GPP started to concentrate on the development of new solutions, based on time or frequency division multiplexing, able to guarantee the coexistence and the cooperation between the two technologies [32].

1.1 Research Motivation and Main Contribution

In the V2X domain, real-world tests and measurements are often impracticable due to the limited availability and the high cost of hardware prototypes. In this regard, network simulators are instrumental to accurately analyze C-V2X SL communications and provide valuable insights on the expected system performance before their actual deployment.

This Chapter presents MoReV2X[1], an open-source ns-3 module for the simulation of LTE-V2X and NR-V2X Sidelink communications [7]. The development of MoReV2X concentrates on the LTE-V2X Mode 4 and NR-V2X Mode 2 distributed resource allocation strategies, and includes the most relevant PHY layer and MAC sublayer features standardized from Release 14 to Release 16. Following the evaluation guidelines summarized in [5], MoReV2X includes an accurate characterization of the V2V sidelink channels that captures large-scale and small-scale fading effects. Furthermore, MoReV2X implements a new set of realistic traffic models and evaluates a complete set of performance metrics. MoReV2X can also be easily interfaced with SUMO [20] for the realistic simulation of vehicular mobility. In this Chapter, the main features of MoReV2X are presented and a preliminary set of exemplary results obtained using the MoReV2X module is reported. To the author's knowledge, MoReV2X is the only open-source simulator based on ns-3 that allows the simulation of both LTE-V2X and NR-V2X Sidelink communications.

As the PHY layer of MoReV2X does not explicitly consider the transmission of the individual bits, but resorts to abstraction techniques to confine the simulation complexity, this Chapter also concentrates on the Link-Level (LL) analysis of NR-V2X SL communications [8]. Link-level simulations are instrumental to obtain BLock Error Rate (BLER) vs Signal-to-Noise Ratio (SNR) curves that capture the impact of fast-fading effects on the PHY layer performance. In MoReV2X, BLER curves are leveraged to implement simplified PHY layer abstraction models and to assess the impact of typical PHY layer aspects, e.g., variable MCS and transmitter-receiver relative speed on system performance. To the author's knowledge, the only existing NR-V2X SL BLER curves have been reported by a handful of 3GPP companies [21–23] during the standardization phase, and no LL simulators have been made publicly available. The lack of dedicated BLER curves limits the number of system configurations that can be analyzed, prohibiting an exhaustive and accurate evaluation of NR-V2X SL communications.

This Chapter presents the first complete set of standard-compliant NR-V2X SL BLER

[1]The code is available at https://github.com/LLusvarghi/MoReV2X.

vs SNR curves and provides a detailed LL analysis of NR-V2X SL communications. The BLER curves reported in this thesis have been obtained using a custom MATLAB-based LL simulator and cover the parameter settings included in the NR-V2X SL profile defined by 3GPP [5] and ETSI [24] evaluation guidelines.

The remainder of this Chapter is organized as follows: Section 1.2 outlines the main differences between link-level and system-level simulation tools. An overview of the available C-V2X SL simulators is provided in Section 1.3 and Section 1.4 presents the most relevant features of the C-V2X SL technology, from LTE-V2X to NR-V2X specifications. Next, Section 1.5 lists the most relevant points of the evaluation methodology suggested by 3GPP, i.e., traffic models, channel models, and performance metrics. Section 1.6 illustrates the MoReV2X implementation and the link-level analysis of C-V2X SL is reported in Section 1.7. Lastly, Section 1.8 draws the conclusions.

1.2 System-Level and Link-Level Simulation Tools

In principle, C-V2X SL simulations could be performed by a single simulation tool. However, constraints on the hardware resources and on the computational power of modern computers limit the complexity of simulations, which must rely on distinct simulation tools to analyze the PHY layer (link-level) and network (system-level) performance.

Link-level simulations concentrate on PHY layer aspects and are used to simulate a point-to-point communication link, featuring an extremely detailed implementation of the transmitter-receiver chain. The transmitter and the receiver chains are implemented up to the individual signal samples, taking into account every detail of the encoding, modulation, channel estimation and synchronization techniques. Link-level simulations employ frequency and time selective channel models to capture the impact of small-scale fading effects on the correct recovery of the transmitted packet. For a given average SNR value, the transmission-reception operations are performed several times in order to experience several channel and noise realizations, i.e., employing the Monte Carlo approach. Link-level simulations are used to evaluate typical PHY layer performance metrics such as the bit error rate and the BLER as a function of the SNR.

It is worth pointing out that link-level simulators have no notion of transmitter-receiver distance, interference, pathloss or shadowing. All the mentioned contributions, which determine the average SNR at the receiver, are modeled in system-level simulations. With respect to their link-level counterpart, system-level simulations are leveraged to examine the upper layers of the protocol stack, concentrating on network-related aspects such as scheduling algorithms and resource management mechanisms. This type of simulation involves a large number of connected vehicles and is used to analyze system performance when different scheduling strategies, channel bandwidths, data traffic sources, and propagation environments are examined. Typical performance metrics obtained during system-level simulations are the end-to-end latency, the throughput, and the fraction of correctly received packets.

In order to confine their complexity, system-level simulations do not explicitly consider the actual transmission of bits, as link-level simulations do, but resort to Link-to-System Mapping (L2SM) techniques to implement simplified PHY layer abstraction models. L2SM is used by system-level simulators to decide whether a received packet is correctly decoded or not based on the SNR level. The SNR is mapped into a BLER value on the basis of Look-Up-Tables (LUTs) that represent the BLER vs SNR curves obtained during link-level simulations.

1.3 Related Work

As of today, system-level simulations have been leveraged by several works for the investigation of LTE-based vehicular communications. Open-source simulators were employed by the authors of [25] and [26]. In [25], Eckermann *et al.* investigated the behavior of LTE-V2X Mode 4 in the urban setting using a system-level simulator based on ns-3. In [26], the OMNeT++-based OpenCV2X simulator was employed to analyze congestion control mechanisms in LTE-V2X Mode 4. Both simulators have not been extended according to Release 16 specifications and, unlike MoReV2X, cannot be employed to analyzed NR-V2X SL communications. To the author's knowledge, the only system-level simulator able to analyze both LTE-V2X and NR-V2X SL communications has been developed by Bazzi *et al.* [27, 28]. However, their simulator leverages a simplified PHY layer error model based on SNR thresholds and does not rely on conventional L2SM techniques like MoReV2X does.

Among 5G studies, the 5G-LENA simulator [29] has been recently extended with a NR-V2X SL module [30]. 5G-LENA is an open-source system-level simulator based on ns-3. Yet, the PHY abstraction model employed by the NR-V2X module leverages BLER curves obtained with a proprietary LL simulator that follows 5G NR Uplink (UL) and Downlink (DL) specifications [31]. The BLER curves obtained using UL or DL simulators cannot be used in NR-V2X SL system-level simulations, because NR-V2X SL is characterized by different PHY layer design choices, as highlighted in [32]. With respect to its UL and DL counterparts, NR-V2X SL features a peculiar PHY layer structure that multiplexes the control and data transmission on the same time-frequency resources. According to it, NR-V2X SL employs specific encoding mechanisms and DMRS patterns. Another NR-V2X SL network simulator based on ns-3 was presented by the authors of [10]. In this case, the BLER vs SNR curves used to represent the control channel and shared channel PHY layer performance are extracted from [33]. Although the employed curves adhere to NR-V2X SL specifications, they cover a small set of system settings and restrict the number of possible configurations that can be analyzed.

When we concentrate on NR-V2X SL communications, scanning the available literature sheds light on the need for standard-compliant LL simulators and BLER vs SNR curves. To the author's knowledge, the only open-source 5G NR LL simulation tools are those presented in [34] and [35]. In [34], the authors introduced the MATLAB-based GTEC

5G link-level simulator. This LL tool allows the simulation of two different PHY layer waveforms and can be easily integrated with a hardware testbed developed by the same authors. The Vienna 5G LL simulator was presented in [35]. The Vienna 5G LL simulator allows an exhaustive evaluation of the numerous PHY layer proposals that characterized the 5G standardization phase. With this tool, link-level simulations can compare four OFDM-based waveforms, different OFDM numerologies and alternative MIMO processing algorithms. Although based on MATLAB, its implementation relies on custom functions developed by the authors. Both the GTEC 5G and the Vienna 5G LL simulators focus on the 5G NR UL/DL analysis, and do not include any additional module based on NR-V2X SL specifications.

In recent years, the BLER vs SNR performance of NR-V2X communications has been analyzed by a handful of papers using proprietary link-level simulators. In [36] the authors compared the PHY layer performance of NR-V2X SL and IEEE 802.11bd to assess which technology is most suited for the support of V2V communications. However, the NR-V2X SL PHY layer is modeled following NR UL specifications in this work. Similarly, the implementation of the LL simulator presented in [37] is based on LTE-V2X SL specifications and lacks some relevant features introduced within the PHY layer of NR-V2X SL. As highlighted in [32], the standardization of NR-V2X included some important novelties with respect to 5G NR UL and LTE-V2X SL. Therefore, the BLER curves presented in [36] and [37] cannot be used to accurately represent the PHY layer performance in NR-V2X SL system-level simulations. A standard-compliant LL simulator is employed by the authors of [38] to evaluate the new NR-V2X SL control channel design. Yet, the BLER vs SNR curves reported in this work have been obtained using a proprietary LL simulator and the analysis cannot be extended to the shared channel, where the transmission's payload is accommodated. With respect to the control channel, the shared channel has a specific PHY layer design and can employ a wider range of MCS values; thus, it requires dedicated BLER vs SNR curves during NR-V2X SL system-level simulations.

To the authors' knowledge, MoReV2X is the only open-source system-level simulator that allows the simulation of LTE-V2X and NR-V2X Sidelink communications leveraging an accurate PHY layer error model based on BLER vs SNR curves.

1.4 Overview of C-V2X SL Communications

1.4.1 LTE-V2X SL

The 3GPP introduced the first C-V2X SL technology, known as LTE-V2X SL, in Release 14. LTE-V2X SL communications occur over the PC5 interface standardized in LTE Device-to-Device (D2D) specifications (Release 12) and reuse the upper V2X layers and protocols specified by the European Telecommunications Standardization Institute (ETSI) and the Society of Automotive Engineers (SAE). At PHY layer, LTE-V2X SL inherits several aspects and functionalities from UL/DL specifications. On the other hand, it

introduces a new distributed resource allocation mode at MAC sublayer, Mode 4, that allows vehicles to autonomously select their transmission resources.

1.4.1.1 PHY Layer

At PHY layer, LTE-V2X SL is allowed to operate in the 5.9 GHz ITS band over 10 MHz or 20 MHz wide channels employing half-duplex radios. It adopts Single-Carrier Frequency Division Multiple Access (SC-FDMA) with a fixed SubCarrier Spacing (SCS) of 15 kHz. According to it, transmission resources are organized over the time-frequency grid reported in Fig. 1.1, where the fundamental time unit is the subframe, whose duration is $t_s = 1$ ms, and the basic frequency unit is the Resource Block (RB), 180 kHz wide. A group of adjacent RBs within the same subframe defines a subchannel, which represents the smallest time-frequency unit for data transmission and reception.

Every time a vehicle has a packet to transmit, it encapsulates its content in a Transport Block (TB). The TB transmission is accommodated over the Physical Sidelink Shared Channel (PSSCH). Depending on the TB size and the employed MCS, the PSSCH transmission requires a variable number of subchannels within the same subframe. TBs are encoded using turbo codes and transmitted using QPSK, 16QAM, or 64QAM modulations. The transmission of each TB is associated with the transmission of the corresponding Sidelink Control Information (SCI), which includes relevant elements for the correct decoding of the TB. For example, the SCI includes the adopted MCS, indicates the RBs occupied by the TB, and specifies whether it is a first transmission or a blind retransmission of the TB. A TB cannot be properly decoded if the associated SCI is not correctly received. The SCI is accommodated over the Physical Sidelink Control Channel (PSCCH), which is transmitted in the same subframe as its associated PSSCH and it occupies 2 consecutive RBs. The PSCCH and the PSSCH are transmitted on frequency-adjacent resources, as shown in Fig. 1.1.

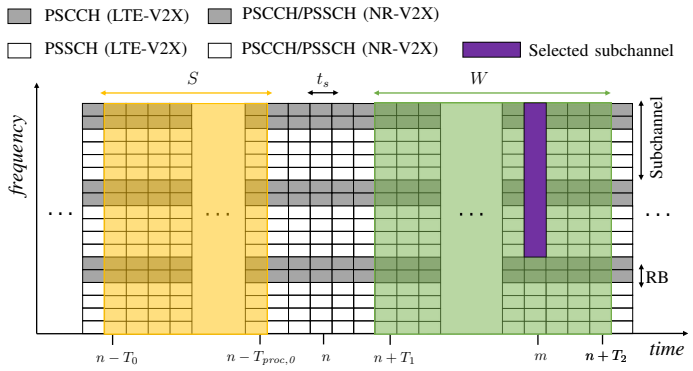

Figure 1.1: Resource reselection mechanism.

In LTE-V2X SL, each subframe consists of 14 OFDM symbols preceded by a normal Cyclic Prefix (CP). The first symbol is used for Automatic Gain Control (AGC) purposes and the last symbol is left idle, allowing vehicles to switch between transmission and

reception mode across consecutive subframes. The number of OFDM symbols used to transmit demodulation reference signals (DMRSs) is equal to 4. DMRS symbols are used to guarantee a robust channel estimation at the receiver and to combat the Doppler effect. With respect to LTE-D2D specifications, where only 2 symbols are reserved for the DMRS transmission, LTE-V2X SL is able to cope with larger transmitter-receiver relative speeds. The remaining OFDM symbols are used to accommodate the PSSCH and PSCCH transmission. The subframe structure is illustrated in Fig. 1.2.

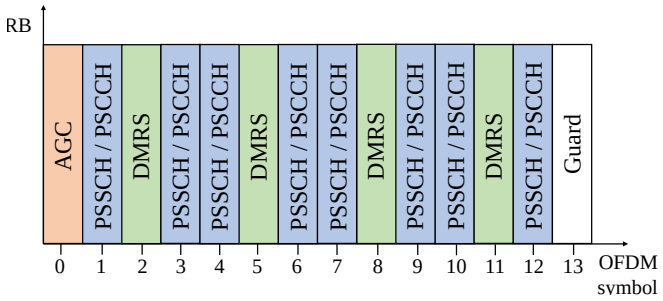

Figure 1.2: Subframe structure in LTE-V2X SL.

1.4.1.2 MAC Sublayer

To support vehicular communications also in out-of-coverage scenarios, LTE-V2X introduced a distributed resource allocation scheme known as Mode 4. In LTE-V2X Mode 4, vehicles autonomously select their radio resources employing the fairly sophisticated Semi-Persistent Scheduling (SPS) strategy, without the support of the cellular infrastructure. Vehicles using the SPS scheme periodically reserve the selected subchannel(s) for a number of reselection counter, C_{resel}, consecutive transmissions. The time period between consecutive reservations is called Resource Reservation Interval (RRI), and LTE-V2X SL supports any RRI value in the $\{20, 50, [100{:}100{:}1000]\}$ ms range, where in the $[x{:}y{:}z]$ notation, x denotes the minimum allowed value, z the maximum and y the incremental step between consecutive values. Vehicles select the RRI from a list of allowed RRI values that can include up to 16 entries, and employ the SCI to broadcast the adopted RRI configuration and inform neighboring vehicles about their next reservation. A vehicle sets the RRI equal to 0 to announce that it will not reserve the currently utilized subchannels for its future TB transmission.

Depending on the selected RRI, the value of the reselection counter is set as follows: if RRI $\geq$ 100 ms, C_{resel} is randomly set in the $[5, 15]$ interval. Otherwise, the value of the reselection counter is randomly set in the $[10, 30]$ and $[25, 75]$ interval when RRI = 50 ms and RRI = 20 ms, respectively. After every transmission, the value of the reselection counter is decremented by one; when $C_{resel} = 0$, the ego-vehicle selects new subchannels with probability $1 - P$, where P is the keep probability, $P \in [0, 0.8]$.

To select new subchannels, the SPS strategy employs a two-step mechanism called resource reselection. Let us call n the subframe at which the ego-vehicle generates a new

message and triggers the resource reselection process: in step 1, the ego-vehicle initially builds a list of the candidate resources that lie within the so-called selection window, W. The selection window extends from subframe $n + T_1$ to subframe $n + T_2$, where T_1 is not larger than 4 subframes and $T_{2min} \leq T_2 \leq$ PDB. T_{2min} depends on the priority of the transmission and its value is in the $[10, 20]$ ms range, whereas the Packet Delay Budget (PDB) is set to satisfy the latency requirements of the generated packet.

During step 1, the ego-vehicle creates a list, L_1, that contains all the available resources included within W. To do so, it removes from L_1 the subchannels that have already been reserved by the SCI of a neighboring vehicle during the ego-vehicle's sensing window, S. The sensing window is the time interval identified by the $[n - T_0, n - T_{proc,0}]$ range of subframes. According to the standard, $T_{proc,0}$ is equal to 1 subframe and T_0 corresponds to 1000 subframes. Subchannels are excluded from L_1 only if the associated SCI is received with a Reference Signal Received Power (RSRP) larger than a pre-configured threshold. The ego-vehicle also excludes from list L_1 all the subchannel(s) that lie on a subframe in which it was previously transmitting. If the ego-vehicle was transmitting during the sensing window, e.g., at subframe s, it could not sense the reservations announced by its neighbors during the same subframe due to its half-duplex limitations. Therefore, the ego-vehicle removes all the subchannels located at subframe $s + \text{RRI}_i$ within W, where RRI_i corresponds to every RRI included in the list of allowed values. The RSRP threshold is increased by 3 dB, and the resource reselection process is re-executed, until L_1 contains at least the β% of the resources initially included in W. Depending on the priority of the TB, β can be set to $20, 35$, or 50.

Next, the ego-vehicle sorts in ascending order the candidate resources of L_1 on the basis of the average Received Signal Strength Indicator (RSSI), and includes the top 20% of the sorted resources in a second list, L_2. If a candidate resource lies at subframe s, its average RSSI is computed in a periodic fashion over the previous $s - T_{L_2} \cdot j$ subframes within S, where j is a positive integer, $T_{L_2} = 100$ ms if RRI ≥ 100 ms, and $T_{L_2} =$ RRI otherwise.

In step 2, the ego-vehicle randomly selects a number of adjacent subchannels from L_2, e.g., at subframe m in Fig. 1.1, able to accommodate the transmission of the TB and the associated SCI. The most relevant elements of the resource reselection mechanism are illustrated in Fig. 1.1.

1.4.2 NR-V2X SL

Under the 5G umbrella, the advent of more sophisticated eV2X services, such as cooperative perception and maneuver coordination [19], has prompted the development of the NR-V2X SL technology. At PHY layer, NR-V2X SL supports a flexible subcarrier spacing configuration, employs more sophisticated DMRS patterns, and multiplexes the transmission of control and data information on time and frequency adjacent resources. At MAC sublayer, 3GPP also devoted significant efforts to develop a new re-evaluation mechanism

and a new resource allocation mode, known as NR-V2X Mode 2, more flexible and effective than Mode 4, its LTE-V2X predecessor.

1.4.2.1 PHY Layer

In NR-V2X SL, the information is transmitted using a CP - Orthogonal Frequency Division Multiplexing (OFDM) waveform that supports different subcarrier spacing settings. The supported SCS values are $2^{\mu} \times 15$ kHz, where μ is the OFDM numerology index, $\mu = 0, 1, 2$. These values correspond to a subcarrier spacing of 15, 30 and 60 kHz, respectively. Transmission resources are organized on a time-frequency grid like the one reported in Fig. 1.1, in which the smallest time and frequency units are the time slot and the RB. The duration of the time slot is defined as $t_s = 2^{-\mu}$ ms. Depending on the adopted OFDM numerology, t_s is equal to 1, 0.5, or 0.25 ms. A time slot includes up to 14 OFDM symbols when normal CP is utilized. In the frequency domain, an RB consists of 12 adjacent OFDM subcarriers and it is 180, 360, or 720 kHz wide depending on the employed SCS. On the time-frequency grid of NR-V2X SL radio resources, the Resource Element (RE) represents the fundamental unit. A RE consists of a single subcarrier and a single OFDM symbol. A group of M_{sub} consecutive RBs within the same time slot defines a subchannel, which represents the smallest time-frequency unit for data transmission and reception. According to the standard [39], the subchannel size M_{sub} can be equal to 10, 12, 15, 20, 25, 50, 75, or 100 RBs. With respect to LTE-V2X SL, where the smallest time unit is a 1 ms long subframe and the RB width is fixed, the use of a flexible SCS at PHY layer guarantees an enhanced flexibility in the arrangement of transmission resources on the OFDM time-frequency grid. For instance, vehicles might use a large SCS for low latency applications (due to the shorter slot duration) or to improve the robustness against Doppler effects. It is worth pointing out that all the vehicles operating in a given NR-V2X system employ the same OFDM numerology.

When a new message is generated, its content is encoded using Low Density Parity Check (LDPC) codes and encapsulated within a TB. The TB transmission is accommodated over the PSSCH and occupies an integer number n of subchannels, depending on the TB size and on the employed MCS. In NR-V2X SL, the supported modulation schemes range from QPSK to 256QAM and the MCS can be selected from either Table 5.1.3.1-1, Table 5.1.3.1-2 or Table 5.1.3.1-3 [40]. Each TB transmission is associated with dedicated SCI. The SCI is transmitted in two stages in NR-V2X SL: the 1$^{\text{st}}$-stage SCI includes relevant information for the correct decoding of the TB and it is transmitted on the PSCCH. The 2$^{\text{nd}}$-stage SCI is transmitted together with the TB on the PSSCH, and it carries information for supporting retransmissions and channel state reports.

In NR-V2X SL, the number of OFDM symbols within a slot, l_d, is flexible and can range from 7 to 14, occupying the full slot. The first OFDM symbol is used for AGC purposes and the last is used as guard symbol, like in LTE-V2X SL. Specifically, the AGC symbol is filled with a copy of the second OFDM symbol, whereas the guard symbol is left idle

to allow the switching between transmission and reception mode. The remaining symbols are used to accommodate the PSSCH and the PSCCH together with the corresponding DMRSs, as shown in Fig. 1.3 for $l_d = 14$. The multiplexing of PSSCH, PSCCH, and

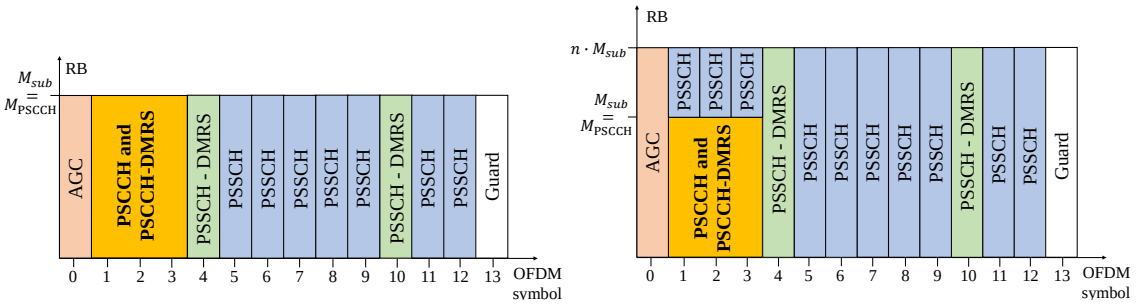

(a) 1 subchannel. (b) n subchannels.

Figure 1.3: Slot structure in NR-V2X SL when $N_{\text{DMRS}} = 2$.

DMRSs is performed on adjacent REs in NR-V2X SL, as illustrated in Fig. 1.3, and it is a specific feature of NR-V2X SL which guarantees an enhanced flexibility at the expense of a more sophisticated PHY layer design:

- **PSCCH:** The PSCCH is mapped to the lowest RBs of the first subchannel utilized for the NR-V2X SL transmission, as shown in Fig. 1.3(b), and it accommodates the 1$^{\text{st}}$-stage SCI. The number M_{PSCCH} of RBs utilized by the PSCCH can be 10, 12, 15, 20 or 25, as long as $M_{\text{PSCCH}} \leq M_{sub}$. In the time domain, the PSCCH can utilize 2 or 3 OFDM symbols starting from the 2$^{\text{nd}}$ symbol of the slot. In Fig. 1.3, the PSCCH occupies 3 OFDM symbols. The content of the 1$^{\text{st}}$-stage SCI follows the SCI format 1-A defined in [41], Section 8.3.1. According to it, the 1$^{\text{st}}$-stage SCI carries information about the TB priority, the RRI used at MAC sublayer, the MCS employed by the TB, the DMRS pattern and the 2$^{\text{nd}}$-stage SCI format among others. The size of the 1$^{\text{st}}$-stage SCI ranges from 20 to 43 bits depending on the configuration of higher layer parameters.

- **PSSCH:** The PSSCH accommodates the transmission of both the 2$^{\text{nd}}$-stage SCI and the TB. This is a distinctive feature of NR-V2X SL with respect to LTE-V2X SL and 5G NR UL/DL specifications. The 2$^{\text{nd}}$-stage SCI can be formatted according to the SCI format 2-A (35 bits long) or 2-B (48 bits long) defined in [41], Section 8.4.1. Both formats include information related to Hybrid Automatic Repeat ReQuest (H-ARQ) retransmissions and a cast type indicator that allows to discriminate between unicast, groupcast and broadcast NR-V2X SL communications. The 2$^{\text{nd}}$-stage SCI is also used to transmit channel state information reports. The TB and the 2$^{\text{nd}}$-stage SCI are separately encoded and multiplexed into a single codeword prior to OFDM modulation. The multiplexing operation depends on the number of Multiple Input Multiple Output (MIMO) transmission layers [41], which is limited to 2 in NR-V2X SL.

- **DMRS Pattern:** According to its peculiar PHY layer structure, NR-V2X SL features dedicated DMRS patterns for the PSSCH and the PSCCH. This is a specific feature of NR-V2X SL and it is a relevant difference with respect to LTE-V2X SL, where a fixed 4 DMRS pattern is employed for both the PSSCH and the PSCCH, as shown in Fig. 1.2. The number of OFDM symbols employed to accommodate DMRSs associated with the PSSCH, N_{DMRS}, can be equal to 2, 3 or 4. Vehicles can employ a larger number of PSSCH-DMRS symbols to enhance the robustness of the V2V link, sacrificing the amount of information transmitted on the PSSCH. The time-domain location of the PSSCH-DMRS symbols within a slot depends on the duration of the PSCCH (2 or 3 OFDM symbols) and on the value of l_d, as reported in Table 1.1. A complete list of the allowed PSSCH-DMRS configurations is available in Table 8.4.1.1.2-1 of [42]. In Fig. 1.3, $l_d = 14$, $N_{\text{DMRS}} = 2$, and the PSCCH duration is equal to 3 OFDM symbols.

  Note that PSSCH-DMRS configurations with $N_{\text{DMRS}} > 2$ multiplex the first PSSCH-DMRS symbol (always located at symbol 1) with the PSCCH on frequency-adjacent RBs. When the NR-V2X SL transmission employs a single sub-channel, configurations with $N_{\text{DMRS}} > 2$ are allowed only if $M_{sub} \geq 20$ RBs, as explained in Section 8.2.2 of [40]. If $M_{sub} < 20$ RBs, the first subchannel only supports the configuration with 2 PSSCH-DMRS symbols, i.e., $N_{\text{DMRS}} = 2$. It is also important to note that PSSCH-DMRSs are transmitted on 1 every 2 REs on the allocated OFDM symbols, as explained in Section 8.4.1.2.2 of [42]. The REs not used to transmit PSSCH-DMRSs are employed to accommodate the PSSCH.

  In the PSCCH case, the DMRS signals (PSCCH-DMRS) are transmitted on the same OFDM symbols allocated for the PSCCH and occupy 1 every 4 REs, see Section 8.4.1.3.2 of [42]. Like in the PSSCH case, the PSCCH is transmitted in the REs that are not utilized to accommodate the PSCCH-DMRS transmission.

Table 1.1: PSSCH-DMRS time domain location in NR-V2X SL.

l_d	PSCCH duration: 2 OFDM symbols			PSCCH duration: 3 OFDM symbols		
	$N_{\text{DMRS}} = 2$	$N_{\text{DMRS}} = 3$	$N_{\text{DMRS}} = 4$	$N_{\text{DMRS}} = 2$	$N_{\text{DMRS}} = 3$	$N_{\text{DMRS}} = 4$
12	3,10	1,5,9	1,4,7,10	4,10	1,5,9	1,4,7,10
13	3,10	1,5,9	1,4,7,10	4,10	1,5,9	1,4,7,10
14	3,10	1,6,11	1,4,7,10	4,10	1,6,11	1,4,7,10

1.4.2.2 MAC Sublayer

At MAC sublayer, NR-V2X SL relies on the Mode 2 distributed resource allocation technique to allow vehicles to autonomously select their transmission resources. In NR-V2X Mode 2, vehicles employ either an SPS or a Dynamic Scheduling (DS) strategy for selecting the subchannel(s) that will accommodate the TBs transmission. The SPS strategy

has been inherited from LTE-V2X specifications with minor modifications, whereas the DS strategy is a new reservation-less solution that forces the selection of new transmission resources for every generated message. Vehicles using the SPS scheme periodically reserve the selected subchannel(s) for a number of reselection counter, C_{resel}, consecutive transmissions. NR-V2X Mode 2 supports any RRI value in the $\{0, [1{:}1{:}99], [100{:}100{:}1000]\}$ ms range. With respect to its LTE-V2X Mode 4 predecessor, NR-V2X Mode 2 provides higher flexibility to fit the requirements and characteristics of eV2X services by allowing any integer RRI between 1 and 99 ms. Also in NR-V2X Mode 2, vehicles select the RRI from a list of allowed values that can include up to 16 entries. Depending on the selected RRI, the value of the reselection counter is set as follows: if RRI $\geq$ 100 ms, C_{resel} is randomly set between 5 and 15. Otherwise, C_{resel} is randomly selected in the $[5 \cdot C, 15 \cdot C]$ interval, $C = 100/\max(20, \text{RRI})$.

On the other hand, vehicles using the DS scheme select new subchannels every time a new message is generated, and are not allowed to place any reservation. In other words, the DS strategy is the reservation-less variant of the SPS scheme which sets $C_{resel} = 1$ and $P = 0$.

Despite being characterized by a totally different reservation policy, the SPS and the DS scheme employ the same two-step resource reselection process for selecting new subchannels. The resource reselection process is inherited from LTE-V2X Mode 4 with a major difference: list L_2 is not considered and, during step 2, vehicles randomly select their resources directly from L_1. In NR-V2X Mode 2, the value of the T_1 parameter is not fixed, but depends on the adopted SCS. T_1 is not larger than $T_{proc,1} = \{3, 5, 9\}$ slots for a SCS of $15, 30, 60$ kHz. Similarly, the value of $T_{proc,0}$ is equal to 1 slot when the SCS is 15 or 30 kHz, and equal to 2 slots when the SCS is 60 kHz. According to the standard [39], also the width of the sensing window is not fixed in NR-V2X Mode 2. The value of T_0 can be set to a number of slots equivalent to either 1100 ms or 100 ms. The most relevant elements of the resource reselection mechanism are reported in Fig. 1.1.

In addition to the resource reselection process, 3GPP has introduced a new mandatory feature in NR-V2X Mode 2 to guarantee a more reliable message delivery, i.e., the re-evaluation mechanism. The re-evaluation mechanism forces the ego-vehicle to keep monitoring the status of its selected resources before transmitting the TB. In principle, this allows the detection of collisions that had not been identified during the resource reselection phase. Note that, in this thesis, we distinguish between reserved and selected resources: we term reserved resources the subchannel(s) whose reservation have been announced by the vehicle through the SCI before being employed for the TB transmissions. On the other hand, we call selected resources the subchannel(s) that have not been reserved by vehicles through SCI broadcasting. Selected resources cannot be announced before being employed and are characterized by higher collision probability. For example, the subchannel(s) chosen at the end of a resource reselection are selected resources.

Following the notation of Fig. 1.1, the functioning of the re-evaluation mechanism is illustrated in Fig. 1.4: let us assume that the ego-vehicle has generated a message

at slot n and that, after performing a resource reselection, has selected a subchannel at slot m. As the resources selected during a reselection cannot be announced beforehand, the re-evaluation mechanism forces the ego-vehicle to re-execute step 1 of the resource reselection process at slot $n_{re} = m - T_3$ to check whether the selected resources are still available. According to the standard, T_3 is set equal to $3, 5, 9$ slots for an SCS of $15, 30, 60$ kHz. The re-execution of step 1 at slot n_{re} is referred to as a *re-evaluation check* in 3GPP specifications. During the re-evaluation check, the ego-vehicle defines a new selection window, W_{re}, that extends from slot $n_{re} + T_1$ to slot $n_{re} + T_2' = n + T_2$ and partially overlaps with the selection window W initially defined at slot n. After removing the unavailable subchannels from W_{re}, the ego-vehicle builds a new list of candidate resources L_1. If the re-execution of step 1 reveals that the originally selected resources, at slot m, have been excluded from list L_1 due to, e.g., a previously undetected reservation from a neighboring vehicle, then the re-evaluation check leads to the *re-evaluation detection* phase according to 3GPP terminology [40]. The re-evaluation detection phase triggers the re-execution of step 2 of the resource reselection mechanism to select new resources. As a result, the initially selected resources are replaced by new resources, e.g., the subchannel located at slot m_{re} in Fig. 1.4. The re-execution of step 2 as part of the re-evaluation mechanism is referred to as *resource replacement*.

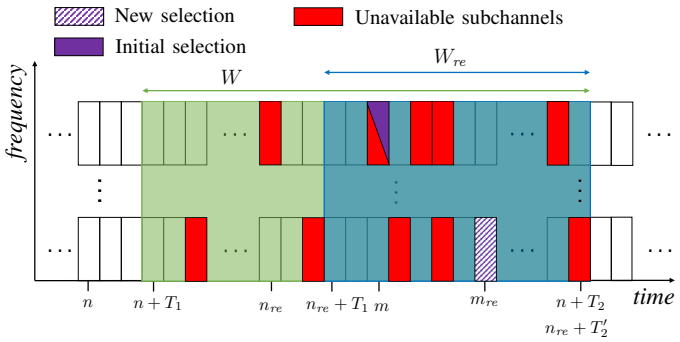

Figure 1.4: Re-evaluation mechanism in NR-V2X Mode 2.

1.5 Evaluation Guidelines

1.5.1 Traffic Models

Among the evaluation guidelines reported in [5], 3GPP included a set of periodic and aperiodic traffic models for the performance assessment of C-V2X SL communications:

- Periodic traffic: in this case, packets are periodically generated every T ms and three different traffic models can be leveraged to reflect varying traffic intensity conditions. In the low intensity *Model 1*, $T = 100$ ms and the packet size follows a $[300, 190, 190, 190, 190]$ bytes pattern. In the medium density *Model 2*, $T = 10$ ms and the packet size X is sampled from a discrete binary distribution in which

$\mathrm{P}\{X = 800 \text{ bytes}\} = 0.8$ and $\mathrm{P}\{X = 1200 \text{ bytes}\} = 0.2$. In the high density case, i.e., *Model 3*, $T = 30$ ms and the packet size is uniformly distributed in the $[30 \text{ kB}, 60 \text{ kB}]$ range with a 10 kilobytes step. Note that the simple assumption of periodic traffic perfectly suits the SPS strategy of LTE-V2X Mode 4 and NR-V2X Mode 2, especially when fixed size packets are considered, but it hinders a thorough evaluation of its scheduling capability. As a matter of fact, the periodic generation of messages does not reflect the nature of the traffic originating from the majority of driving applications. More realistically, the performance of LTE-V2X and NR-V2X SL communications has to be evaluated considering different traffic profiles, in which neither the packet inter-arrival time nor the packet size is constant.

- Aperiodic traffic: in order to mimic the traffic profiles that will likely characterize eV2X applications, 3GPP has introduced a set of aperiodic traffic models which exhibit variable packet size and a non-deterministic packet inter-arrival time, τ. In detail, τ is a random variable modeled as follows:

$$\tau = c + r \tag{1.1}$$

  where c is a constant and r is an exponentially distributed random variable with mean $\bar{r} = c$. Also in this case, 3GPP has introduced two distinct models to reflect different traffic intensities. In the medium intensity *Model 1*, $c = 50$ ms and the packet size is uniformly distributed in the $[200, 1200]$ bytes range with a 200 bytes step. In the high intensity *Model 2*, $c = 10$ ms and the packet size is uniformly distributed in the $[10 \text{ kB}, 30 \text{ kB}]$ with a 4 kilobytes step.

Although more realistic than their periodic counterparts, the aperiodic traffic models generate random traffic profiles that do not correspond to any V2X service. For this reason, MoReV2X features an additional traffic model that replicates the generation of CAMs. Standardized by ETSI in [16], CAMs are facility-layer packets devised to regularly broadcast information about the transmitting vehicle status, e.g., position, speed and type of vehicle. When CAM traffic has to be examined, the packet inter-arrival time should be modeled using the ETSI algorithm, avoiding the use of simplified periodic models. According to it, the minimum time elapsed between consecutive CAM generation events has to be equal or larger than T_{GenCam}, where T_{GenCam} falls in the interval $[T_{GenCamMin}, T_{GenCamMax}] = [100, 1000]$ ms. The default value for T_{GenCam} is $T_{GenCamMax}$. Within such limits, CAMs are triggered depending on the transmitting vehicle dynamics, which have to be sampled every $T_{CheckGenCam}$ milliseconds, $T_{CheckGenCam} \leq T_{GenCamMin}$. The typical setting is $T_{CheckGenCam} \leq T_{GenCamMin} = 100$ ms. Specifically, a new CAM shall be immediately generated every time one of the following conditions is met:

(i) the absolute difference between the current heading and the heading included in the previous CAM is greater than $\Delta H = 4°$;

(ii) the absolute distance between the current position and the position included in the previous CAM is greater than $\Delta D = 4$ m;

(iii) the absolute difference between the current speed and the speed included in the previous CAM is greater than $\Delta S = 0.5$ m/s;

(iv) the time elapsed since last CAM generation is equal to or greater than $T_{GenCamMax}$.

As reported by the experimental results in [14] and [13], the CAM inter-arrival time, T_{CAM}, can assume any value in the $[100, 200, \ldots, 900, 1000]$ ms range when $T_{CheckGenCam} = 100$ ms, generating aperiodic CAM patterns. The ETSI specifications in [16] also provide indications on the mandatory and optional information which shall be conveyed within a CAM. The presence of non-mandatory fields leads to a variable CAM size, depending on the specific implementation. In [14], the smallest and largest observed CAM sizes are 182 and 807 bytes.

1.5.2 Channel Models

A key aspect of the LTE-V2X SL and NR-V2X SL evaluation guidelines reported in [5] concerns the definition of new Sidelink channel models. In this regard, 3GPP devoted significant efforts to model the large and small scale fading effects that characterize vehicular communications in *Highway* and *Urban* propagation environments. For both the Highway and Urban scenarios, 3GPP has introduced two possible channel states, namely:

- **LOS:** the link between the transmitting and receiving vehicle is characterized by a direct LOS path;
- **NLOSv:** the LOS path is blocked by the presence of other vehicles.

For every received packet, the channel state is determined according to the LOS and NLOSv state probabilities, $P(\text{LOS})$ and $P(\text{NLOSv}) = 1 - P(\text{LOS})$. According to [5], $P(\text{LOS})$ is defined as

$$P(\text{LOS}) = \begin{cases} \min\{1, a \cdot b^2 + b \cdot D + c\} & D \leq 475 \text{ m} \\ \max\{0, 0.54 - 0.001 \cdot (D - 475)\} & D > 475 \text{ m} \end{cases} \tag{1.2}$$

with $a = 2.1013 \cdot 10^{-6}$, $b = -0.002$, $c = 1.0193$ and D being the distance between the transmitting and the receiving vehicle.

In the Urban scenario, the LOS and NLOSv states are employed when the transmitting and receiving vehicles are in the same street. When the transmitting and receiving vehicles are in different streets, the V2V Sidelink channel is modeled using an additional Non-Line-Of-Sight (NLOS) state:

- **NLOS:** the direct LOS path is blocked by the presence of buildings,

and $P(\text{NLOS}) = 1$.

As a result, 3GPP guidelines have introduced a total of five different channel models for the evaluation of LTE-V2X and NR-V2X SL communications: Highway LOS and NLOSv, Urban LOS, NLOSv, and NLOS. As illustrated next, specific pathloss, shadowing, and fast fading parameters have been defined for each channel model.

In the Highway scenario, the pathloss (PL) which characterizes both the LOS and the NLOSv states is:

$$PL = 32.4 + 20 \cdot \log_{10}(D) + 20 \cdot \log_{10}(f_c) \tag{1.3}$$

where f_c is the center frequency. In the Urban setting, the pathloss is defined as

$$PL = 38.77 + 16.7 \cdot \log_{10}(D) + 18.2 \cdot \log_{10}(f_c) \tag{1.4}$$

for the LOS and NLOSv states. An additional blockage loss term, B, is added to PL when the transmitting and receiving vehicles are in the NLOSv state [5]. When all vehicles have the same height, the value of B is modeled as a lognormally distributed random variable with standard deviation $\sigma_B = 4$ dB and mean μ_B. The value of μ_B is defined as follows:

$$\mu_B = 5 \text{ dB} + \max\{0, 15 \cdot \log_{10}(D) - 41\}\,. \tag{1.5}$$

When the Urban NLOS channel model is considered, the pathloss is computed as:

$$PL = 36.85 + 30 \cdot \log_{10}(D) + 18.9 \cdot \log_{10}(f_c)\,. \tag{1.6}$$

In (1.3), (1.4), and (1.6) the value of PL is expressed in dB.

In both the Highway and Urban scenarios, shadowing (SH) is modeled as a zero-mean lognormally distributed random variable with standard deviation σ_{SH}. In the LOS and NLOSv states, $\sigma_{SH} = 3$ dB. In the NLOS state, $\sigma_{SH} = 4$ dB. The average pathloss plus shadowing contribution ($PL + SH$) that characterizes the five different channel models is reported in Fig. 1.5.

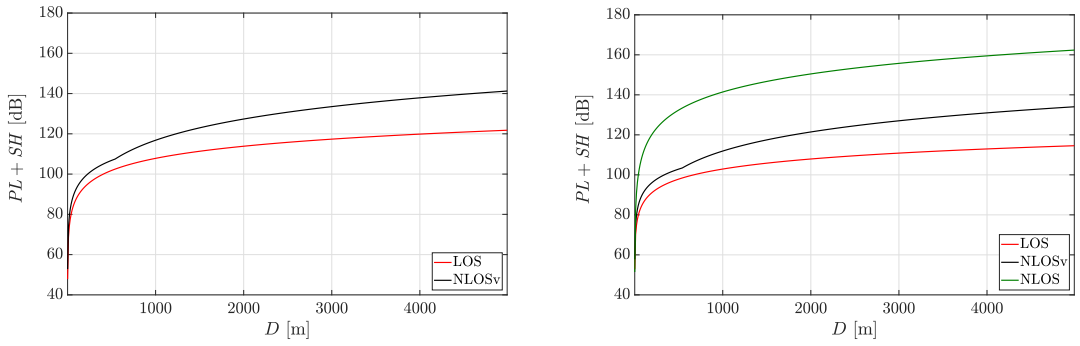

(a) Highway. (b) Urban.

Figure 1.5: Average pathloss plus shadowing contribution ($PL + SH$) in the Highway and Urban propagation environments for different channel states.

As indicated in [5], the state transition between LOS, NLOSv, and NLOS is checked every 100 ms for each Tx-Rx pair during system-level simulations. Accordingly, also the pathloss and the shadowing experienced by each transmission link are updated every 100 ms, based on the current Tx-Rx distance. The shadowing contribution update follows the spatial correlation model defined by 3GPP in [43]. For each Tx-Rx pair, the new shadowing value $S(n)$ is determined as follows:

$$S(n) = e^{-D_{update}/D_{corr}} \cdot S(n-1) + \sqrt{1 - e^{-2D_{update}/D_{corr}}} \cdot SH_n \,, \tag{1.7}$$

where SH_n is a sample drawn from the shadowing lognormal distribution, $S(n-1)$ indicates the previous shadowing sample, and D_{update} is the update distance. D_{update} measures the transmitter-receiver distance variation over the last 100 ms. In (1.7), D_{corr} represents the decorrelation distance and it is equal to 10 m and 25 m in the Urban and Highway scenarios, respectively

Besides the characterization of pathloss and shadowing, 3GPP evaluation guidelines include five different Clustered Delay Line (CDL) channel models in order to capture the small-scale fading effects which impair V2X Sidelink communications in the Highway LOS, NLOSv and Urban LOS, NLOSv, NLOS propagation environments. A CDL channel models the received signal as the superposition of different clusters, or echoes. Each cluster is defined with a specific delay, attenuation, arrival and departure angle. The full set of parameters used to define the CDL channels is reported in [5], Tables 6.2.3.1-1 through 6.2.3.1-5.

1.5.3 Performance Metrics

Along with the definition of traffic and channel models, 3GPP introduced the performance metrics that shall be used for the analysis of LTE-V2X SL and NR-V2X SL communications [5].

For broadcast communications, the first metric is the Packet Reception Ratio (PRR). The PRR is a widely adopted figure of merit that measures the average fraction of correctly received TBs with respect to the total number of transmitted TBs. Its definition relies on the notion of distance slice: the i-th distance slice is defined as the set of transmitter-receiver distances that fall within the $[a_i, b_i]$ range, $a_i = i \cdot 20$ m and $b_i = (i+1) \cdot 20$ m. For the i-th slice, the PRR is computed as:

$$\text{PRR} = \frac{\sum_{j=1}^{N} X_i^j}{\sum_{j=1}^{N} Y_i^j} \tag{1.8}$$

where X_i^j indicates the number of vehicles within the i-th slice that correctly decoded the j-th TB, Y_i^j represents the number of vehicles within the i-th slice when the j-th TB was transmitted, and N denotes the total number of transmitted TBs. In this thesis, the PRR is reported as a function of D, which represents the average transmitter-receiver distance

within the i-th slice.

The second metric recommended by 3GPP is the Packet Inter-Reception (PIR). The PIR is an additional metric designed to highlight the impact of persistent collisions on those applications that require higher reliability with respect to basic safety-related LTE-V2X services. For a given transmitter-receiver pair, the PIR is defined as the time between two consecutive successful receptions of TBs belonging to the same application flow, when the transmitter-receiver distance falls within the $(0, D_{max}]$ range at reception time. Both the PIR Cumulative Distribution Function (CDF) and the average PIR are recommended by 3GPP in [5]. The average PIR, within a given distance D_{max}, is computed as:

$$\overline{\text{PIR}} = \frac{1}{M} \cdot \sum_{j=1}^{M} T_j \tag{1.9}$$

where M denotes the number of collected PIR values during the simulation and T_j indicates the generic PIR value.

MoReV2X computes two additional metrics for the assessment of the LTE-V2X SL and NR-V2X SL system-level performance, namely: the Propagation Losses Ratio (PLR) and the Collision Losses Ratio (CLR). Also the PLR and the CLR definition relies on the notion of distance slice. For the i-th slice, the PLR is computed as:

$$\text{PLR} = \frac{N_{PL}}{N_{PL} + N_{CL} + N_{SR}} \tag{1.10}$$

and the CLR is defined as

$$\text{CLR} = \frac{N_{CL}}{N_{PL} + N_{CL} + N_{SR}} \tag{1.11}$$

where, within the i-th slice of distances: (i) N_{PL} is the number of uncorrectly received TBs because of poor propagation conditions, i.e., N_{PL} is the number of packets that did not collide, but experienced an insufficient SNR at the receiver; (ii) N_{CL} is the number of TBs that were lost in a collision and could not be successfully decoded because of an insufficient Signal-to-Interference-plus-Noise Ratio (SINR); (iii) N_{SR} is the number of correctly received TBs.

Another metric MoReV2X determines is the Channel Busy Ratio (CBR). Defined by 3GPP in [44], the CBR is a physical layer indicator that estimates the channel load. At slot n, the CBR is measured as the fraction of subchannels whose RSSI is larger than a threshold over the $[n - 100 \cdot 2^{\mu}, n - 1]$ range of slots.

1.6 Ns-3 Implementation

In this Section, the main features of the MoReV2X module are presented. The focus is on the implementation of NR-V2X SL and its Mode 2 distributed resource allocation scheme, as it required the development of several new functionalities with respect to its LTE predecessor. LTE-V2X SL simulations rely on simplified versions of the NR-V2X SL

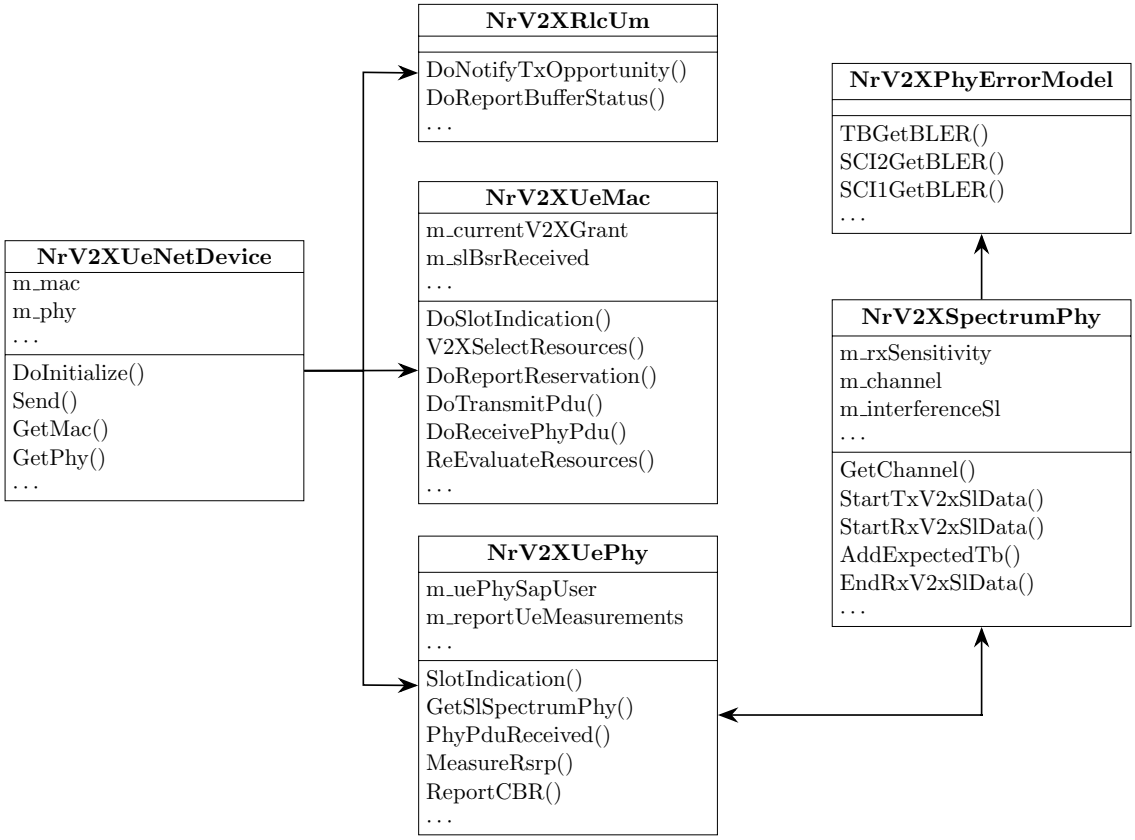

Figure 1.6: Simplified Unified Modeling Language (UML) diagram of the MoReV2X implementation.

classes, methods and functions illustrated in this Section and summarized in the UML diagram reported in Fig. 1.6.

1.6.1 Application Layer

The application layer of the MoReV2X module has been significantly upgraded with respect to the LTE D2D simulator [45, 46] that represented the starting point for the implementation of MoReV2X. The two ns-3 classes `UdpClient` (transmitter side) and `PacketSink` have been modified to support the traffic models mentioned in Subsection 1.5.1. Accordingly, every simulation can be configured using four different command line arguments, namely:

- `Periodic`: when this option is chosen, all vehicles generate periodic traffic following the *3GPP Periodic traffic Model 1* defined in Subsection 1.5.1. According to it, the packet size is modeled using a pattern of five values that endlessly repeats. For each vehicle, the pattern's starting point is randomly set, as indicated in [5]. The PDB is set equal to the generation period, i.e., PDB $= T = 100$ ms. The remaining periodic traffic models, i.e., *Model 2* and *Model 3*, have been implemented but are not available by default.

- `Aperiodic`: when this option is selected, all vehicles generate aperiodic traffic following the *3GPP Aperiodic traffic Model 1*. Based on the system bandwidth and on the employed MCS, the largest packet size value might be reduced to fit in the allocated resources. By default, the PDB is set equal to the value of the constant term c, i.e., PDB $= c = 50$ ms.

- `Mixed`: this configuration allows for the coexistence of `Periodic` and `Aperiodic` traffic sources within the same simulation. The percentage of vehicles generating periodic traffic with respect to the total number of vehicles, termed Δ, is set using the `Percentage` command line argument. The default value is $\Delta = 50\%$.

- `CAM`: this setting allows for the generation of CAM traffic as specified in Subsection 1.5.1. When vehicular mobility is simulated through SUMO, the ETSI-compliant CAMs are produced using a suite of custom `Python` scripts developed by the author. At the beginning of each `CAM` simulation, the following two-step process takes place:

  1. the SUMO simulation is started and its output mobility traces are exploited to determine the sequence of T_{CAM} values for each vehicle, on the basis of the vehicle speed, position and heading variables. The CAM size is fixed to either 190 or 470 bytes, which are the smallest and the largest statistically relevant sizes reported in [14]. The obtained CAM traces are stored in dedicated array variables;

  2. the `LoadCAMTraces` function pairs every ns-3 node to a CAM trace, setting the `m_interval` and `m_size` attributes of each `UdpClient` instance.

  For CAM traffic, a further alternative is offered: in step (1), the CAM generation can exploit the empirical models presented in [47]; these models consists of m-th order Markov sources that reproduce CAM temporal and size patterns in accordance with real-world CAM traces [14], experimentally collected in different scenarios. The empirical models are particularly useful when SUMO-generated mobility traces are not available, or when the impact of more realistic CAMs needs to be evaluated. With respect to the simulative approach based on SUMO, the empirical models allow the generation of variable size CAMs. In `CAM` simulations, the SUMO-based approach is the default configuration, but the empirical models can be enabled using the `EmpModels` command line argument.

A dedicated `Tag` object is attached to every transmitted packet, employing the new `NrV2XTag` ns-3 class, that has been developed from scratch. The packet tag contains high level information inspected by the receiving application. The `NrV2XTag` content includes the packet timestamp, the ID of the transmitting vehicle and its position at the time the packet was generated. Additional fields within the `NrV2XTag` allow for cross-layer communication within the same ns-3 node. For example, the PDB and RRI fields can be

dynamically set on a per-packet basis, and employed by lower layers during the resource selection process in LTE-V2X Mode 4 and NR-V2X Mode 2 simulations.

In addition, the `UdpClient` and `PacketSink` classes have been modified to save the corresponding transmission and reception events in a dedicated format, in order to evaluate the performance metrics mentioned in Subsection 1.5.3.

1.6.2 MAC Sublayer

In MoReV2X, the MAC sublayer features and mechanisms are implemented within the `NrV2XUeMac` class. Specifically, the resource reselection process is implemented by a new method of the `NrV2XUeMac` class, called `V2XSelectResources`. Before the simulation starts, the sensing window width, the RSRP threshold and the minimum amount of candidate resources β are set employing dedicated variables. The width of the selection window W is dynamically adjusted, depending on the latency deadline which characterizes the packet that triggered the resource reselection. The semi-persistent strategy is the default configuration of the `V2XSelectResources` method. However, the dynamic strategy can be enabled via the ns-3 attribute system, using the `Dynamic` command line argument.

Every time new radio resources have to be selected for the transmission of a packet, the `V2XSelectResources` method is invoked for the generation of a valid scheduling grant. Scheduling grants are implemented by means of the `V2XSidelinkGrant` data structure, and contain relevant information for the scheduling of each vehicle transmission. The content of a grant includes the position of the selected (amd reserved) resources on the time-frequency grid, the employed MCS, the RRI configuration and the reselection counter value. A scheduling grant is considered valid if $C_{resel} > 0$ and the associated reservation satisfies the size and latency requirements of the generated packet.

At the beginning of each slot, the `NrV2XUeMac` class employs the `DoSlotIndication` method to check if there is a new packet ready to be transmitted. If so, the transmission opportunity is notified to the PHY layer using the `DoTransmitPdu` method. Otherwise, if there is no packet waiting for transmission, `NrV2XUeMac` switches to reception mode and listens for incoming received packets which are forwarded from lower layers. Inside the `NrV2XUeMac` class, the re-evaluation mechanism is implemented by the `ReEvaluateResources` method.

1.6.3 PHY Layer

The PHY layer of the MoReV2X simulator has been significantly modified with respect to its original implementation to support a flexible subcarrier spacing, different DMRS patterns, and to retrieve the physical layer indicators exploited by NR-V2X Mode 2. For example, the `NrV2XUePhy` and `NrV2XUePhySapUser` classes have been enhanced to evaluate the RSRP of each received TB and to forward such measurement to the upper MAC sublayer, where the resource reselection mechanism runs. At PHY layer, the `NrV2XUePhy` class has been further modified to continuously monitor the RSSI on all the time-frequency

grid resources. This information is periodically exploited by the `ReportCBR` method, which evaluates the CBR level as indicated in Subsection 1.5.3.

In addition, the physical layer of MoReV2X includes the definition of the Highway LOS, NLOSv and Urban LOS, NLOSv, NLOS channel models presented in Subsection 1.5.2. The physical layer impairments associated with each channel model are implemented within the `NrV2XSpectrumPhy` class, which is employed to interface the `NrV2XUePhy` objects associated with each vehicle over a common wireless channel object. To evaluate the impact of large and small scale fading effects on the correct decoding of a TB, `NrV2XSpectrumPhy` runs the following steps:

1) the average SINR of the PSSCH carrying the TB and the 2$^{\text{nd}}$-stage SCI and of the PSCCH used to transmit the 1$^{\text{st}}$-stage SCI are evaluated taking into account the pathloss plus shadowing contribution ($PL + SH$) defined in Subsection 1.5.2. When there is no interference from neighboring vehicles, the SINR corresponds to the SNR and its average value is computed as follows:

$$\text{SINR} = \text{SNR} = P_t - (PL + SH) - N_0 \tag{1.12}$$

where P_t is the transmission power and N_0 is the noise level measured across a single RB. The value of PL depends on the transmitter-receiver distance D (see Subsection 1.5.2), whereas N_0 depends on the noise power spectral density and on the transmission bandwidth. Note that, in (1.12), MoReV2X assumes that every RB is characterized by the same noise and shadowing contributions.

In case of overlapping transmissions from neighboring vehicles, the interference on each RB is interpreted as additional noise and the SINR is averaged over the RBs employed to accommodate the PSSCH and PSCCH transmission:

$$\text{SINR} = \frac{1}{n_{RB}} \sum_{i=1}^{n_{RB}} \text{SINR}_i \tag{1.13}$$

where SINR_i represents the SINR experienced by each RB and n_{RB} is the number of employed RBs.

2) if the average SINR is below the receiver sensitivity, the TB cannot be recovered. Otherwise, the SINR levels of the PSSCH and of the PSCCH are separately mapped into BLER values on the basis of LUTs modeling small-scale fading effects. As mentioned in Section 1.2, the use of simplified PHY layer abstraction models allows to significantly reduce the complexity of the MoReV2X simulator. To do so, `NrV2XSpectrumPhy` invokes the `TBGetBLER`, `SCI2GetBLER`, and `SCI1GetBLER` methods of the `NrV2XPhyErrorModel` class to retrieve the BLER associated with the TB (BLER_{TB}), the 2$^{\text{nd}}$-stage SCI ($\text{BLER}_{\text{SCI2}}$), and the 1$^{\text{st}}$-stage SCI ($\text{BLER}_{\text{SCI1}}$).

3) Next, the 1$^{\text{st}}$-stage SCI is declared successfully decoded with probability $1-\text{BLER}_{\text{SCI1}}$.

As the 1st-stage SCI contains critical information for the correct decoding of the PSSCH, the TB and the 2nd-stage SCI can be decoded only if the 1st-stage SCI has been correctly retrieved. If so, the TB and the 2nd-stage SCI are declared successfully decoded with probability $1 - \text{BLER}_{\text{TB}}$ and $1 - \text{BLER}_{\text{SCI2}}$, respectively.

1.6.4 Exemplary Results

In this Subsection, a set of exemplary results obtained through the MoReV2X module is reported. The simulation scenario consists of a 5 km long highway segment with 3 lanes per driving direction. Vehicular mobility has been recreated using SUMO. In all simulations, the vehicles' speed is 140 km/h and the average vehicular density is 60 vehicles/km.

The system bandwidth is 10 MHz wide, centered at 5.9 GHz in the ITS band; the SCS value is 15 or 30 kHz, thus resulting in a total number of RBs equal to 52 and 24, respectively [48]. In the frequency domain, RBs are organized in subchannels made of 12 RBs each. The TB and the SCI are multiplexed in non-overlapping resources and transmitted during the same time slot, employing QPSK and a code rate 0.5 and 0.23, respectively. The number of symbols occupied by the SCI in each slot is equal to 2, and 4 DMRS symbols are employed for the decoding of the TB. The transmission power is set to 23 dBm for all vehicles.

The results refer to broadcast communications and have been obtained considering periodic and aperiodic traffic profiles. The `Periodic` profile utilizes a packet periodicity equal to $T = 100$ ms and the packet size pattern is $[300, 190, 190, 190, 190]$ bytes. In the `Aperiodic` profile, the constant term c in (1.1) is set equal to 50 ms, and the packet size X is sampled from a discrete binary distribution in which $\text{P}\{X = 190 \text{ bytes}\} = 0.8$ and $\text{P}\{X = 300 \text{ bytes}\} = 0.2$; 190 and 300 bytes long messages are accommodated over 1 and 2 subchannels, respectively. The average traffic generated by the two profiles is the same.

The PDB is equal to the minimum packet inter-arrival time, i.e., 100 ms for the `Periodic` profile, 50 ms for the `Aperiodic` alternative. All simulations adopt the NR-V2X Mode 2 SPS strategy: the number of reserved subchannels is set so as to accommodate the largest expected packet size, i.e., 300 bytes, and subchannels are reserved with a periodicity $RRI = PDB$. New resources are selected every time the reselection counter expires, i.e., $P = 0$. The Mode 2 initial RSRP threshold is configured to -128 dBm and the β threshold is set to 20%.

Fig. 1.7 reports the average PRR as a function of the distance D between the transmitting and the receiving vehicle. The SPS strategy achieves higher PRR levels when periodic traffic is considered, as constant packet inter-arrival times perfectly fit the periodic, semi-persistent resource reservation approach. This holds true irrespective of the SCS value. The figure also reveals that for periodic traffic, the PRR performance improves when increasing the SCS from 15 to 30 kHz. This effect has to be mainly ascribed to a different number of slots in the selection window, which is twice as high in the 30 kHz case; in turn, this means a lower probability that two or more vehicles select a subchannel

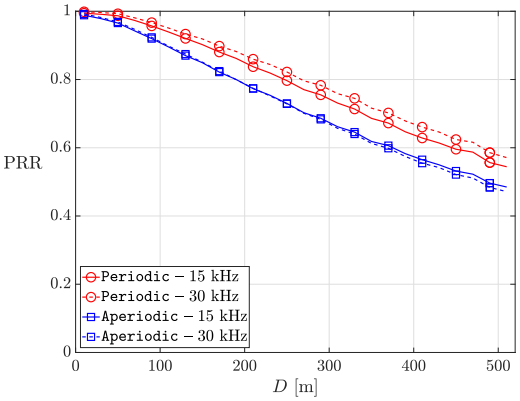

Figure 1.7: PRR as a function of the Tx-Rx distance D.

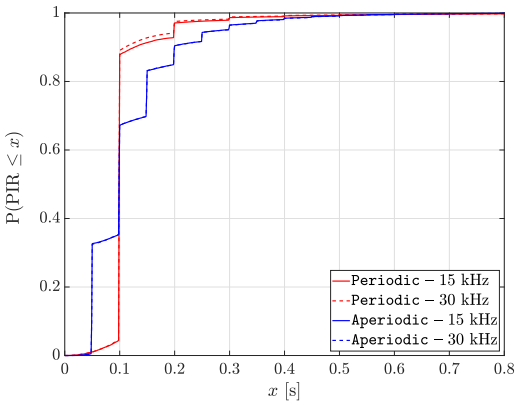

Figure 1.8: PIR CDF, $D_{max} = 520$ m.

in the same slot. When this event happens, vehicles transmitting within the same slot are "blind" to each other and persistently collide if they select the same subchannels. On the other hand, the PRR performance for aperiodic traffic does not improve when increasing the subcarrier spacing. Here, the main cause of PRR degradation is the mismatch between the aperiodic packet generation and the periodic allocation of resources.

Next, Fig. 1.8 reports the PIR CDF for $D_{max} = 520$ m. For all simulations, the CDF step behavior reflects the periodic nature of the reservations placed using the SPS strategy, and indirectly validates MoReV2X accurate implementation. For periodic traffic, the probability of observing PIR values lower than 100 ms is close to 0.9, further demonstrating the NR-V2X Mode 2 ability to cope with this traffic type.

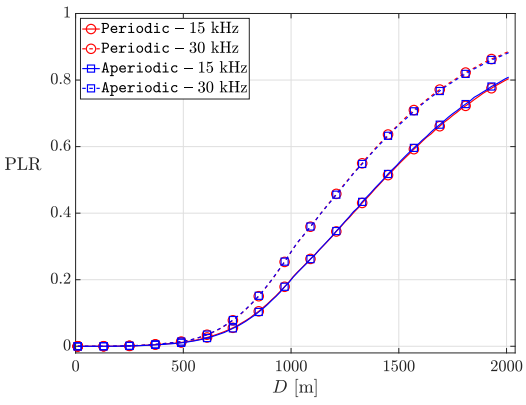

(a) PLR.

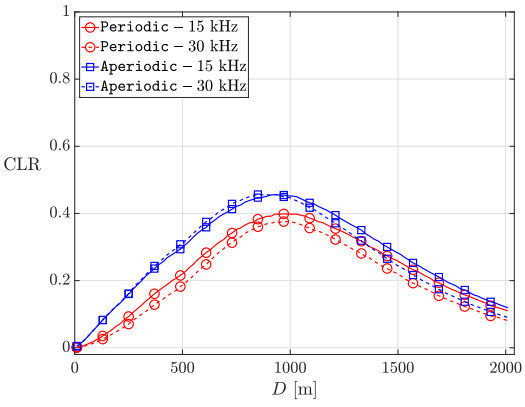

(b) CLR.

Figure 1.9: PLR and CLR curves reported as a function of the Tx-Rx distance D.

In Fig. 1.9(a), the PLR as a function of the distance D is reported, where it is recalled that the PLR measures the fraction of packets exclusively lost because of poor propagation conditions. As expected, this figure shows that the PLR does not depend on the traffic type, but it exclusively depends on the SCS. In detail, it reveals that propagation losses

are negligible for D values lower than 500 m and that above this value the SCS = 15 kHz choice guarantees lower PLR values than SCS = 30 kHz. As a matter of fact, the noise power of a subchannel is twice as high when SCS = 30 kHz, for a given noise power spectral density. Therefore, packets transmitted in the system adopting SCS = 30 kHz exhibit an SNR value which is 3 dB lower than in the system where SCS = 15 kHz, on average. Moreover, given the vehicles relative speed is not excessively high (its maximum value is 280 km/h), the 30 kHz choice does not improve system robustness against Doppler effects.

Then, Fig. 1.9(b) reports the CLR as a function of D. In the presence of periodic traffic, collisions are constantly lower than for the aperiodic traffic case. This figure corroborates the intuition that a smaller time slot duration, due to the adoption of a larger subcarrier spacing, reduces the number of collisions only in the case of periodic traffic.

At PHY layer, the measured CBR levels are equal to 0.38 for periodic traffic and to 0.36 in the case of aperiodic traffic. The RSSI threshold employed for the CBR evaluation is 0.5 dB greater than the receiver sensitivity, i.e., -89.9 dBm. Note that the CBR for is larger when periodic traffic is considered. As a matter of fact, the periodic and the aperiodic profiles pour the same amount of traffic on the channel, but the SPS strategy is able to allocate resources more effectively in the former circumstance. For the same channel load, a smaller number of collisions results in a more efficient occupation of the channel, hence leading to a larger CBR.

1.7 Link-Level Evaluation of NR-V2X SL Communications

System-level simulators like the MoReV2X module do not explicitly consider PHY layer aspects. Instead, they resort to abstraction techniques to confine the simulation complexity. The adopted PHY layer abstraction technique can significantly affect the results of system-level simulations, as shown in [49]. The most common solution relies on the use of BLER curves obtained using dedicated link-level simulators. As evidenced in Section 1.3, the only 5G NR-based LL simulators that have been made publicly available consider either UL or DL communications. However, the BLER curves obtained using these tools cannot be used in NR-V2X SL system-level simulations because the SL technology is characterized by different PHY layer design choices, as highlighted in [32]. With respect to its UL and DL counterparts, NR-V2X SL features a peculiar PHY layer structure and dedicated functionalities:

- The PSCCH and the PSSCH are multiplexed on time-and-frequency adjacent resources.
- The SCI is transmitted in two distinct stages and each stage employs specific encoding and decoding procedures.
- Similarly, the TB undergoes specific encoding and decoding mechanisms. The maximum TB size is computed using a new and dedicated procedure.

- A new set of DMRS patterns is separately defined to protect the PSCCH and the PSSCH transmission.

During the standardization phase, NR-V2X SL BLER curves have been reported by several companies within 3GPP working groups [21–23,33]. Yet, these curves have been generated using proprietary LL simulators and cover only a limited set of parameter settings, thus restricting the scope of the associated NR-V2X SL analysis. The lack of dedicated BLER vs SNR curves has a two-fold effect. First, it does not allow an exhaustive and accurate system-level evaluation of NR-V2X SL communications. Second, it does not guarantee consistent and comparable results between system-level studies that employ curves from different sources. In this regard, the contribution of this Section is two-fold:

- Present the first complete set of standard-compliant NR-V2X SL BLER vs SNR curves.
- Provide a detailed LL analysis of NR-V2X SL to assess the impact of typical PHY layer aspects, e.g., variable MCS and transmitter-receiver relative speed, on the BLER performance.

The BLER curves presented in this Section have been obtained using a custom MATLAB-based LL simulator and cover the parameter settings included in the NR-V2X SL profile. Defined by the 3GPP and ETSI evaluation guidelines reported in [5] and [24], the NR-V2X SL profile details a recommended configuration of the most important system settings (e.g., subchannel size, channel bandwidth, modulation scheme, code rate) in order to establish a common evaluation framework. Following the NR-V2X SL profile indications, we have obtained more than 1900 BLER vs SNR curves. The curves are openly accessible[2] in the form of LUTs and can be easily integrated into any system-level simulator using L2SM techniques [8]. The PHY layer of the MoReV2X module relies on the BLER vs SNR curves presented in this Section.

The remainder of this Section is organized as follows: Subsection 1.7.1 provides a functional description of the NR-V2X SL link-level simulator employed in this thesis, highlighting the main differences with respect to 5G NR UL and DL simulation tools. Subsection 1.7.2 illustrates the most relevant settings included within the NR-V2X SL profile and Subsection 1.7.3 underlines the relevance of the obtained results by reporting an exemplary set of BLER vs SNR curves.

1.7.1 NR-V2X SL Link-Level Simulator

In this Subsection, we provide a description of the NR-V2X SL LL simulator employed to assess the link-level performance of NR-V2X SL communications and to obtain the BLER curves presented in this study. The simulator is implemented in MATLAB and it is

[2]Unfortunately, we are not yet able to share the link of the open-source repository since the corresponding paper is under review. The repository link is included in the paper and will be available once the review process is completed.

based on the 5G Toolbox functionalities. As the 5G toolbox provides standard-compliant functions for 5G NR UL and DL, significant efforts were devoted to the implementation of the new features which characterize NR-V2X SL communications at PHY layer, namely:

- We have developed a new set of functions to determine the TB size and to properly format the 1st-stage and 2nd-stage SCI.
- We have customized the encoding, decoding, modulation and demodulation functionalities of MATLAB's 5G toolbox according to 3GPP specifications.
- We have implemented from scratch the PSCCH and PSSCH multiplexing mechanism.
- We have implemented from scratch the PSCCH-DMRS and PSSCH-DMRS patterns that have been specifically designed for NR-V2X SL.
- We have included the NR-V2X SL channel models provided by 3GPP evaluation guidelines.

The general structure of the simulator and its workflow are summarized in Fig. 1.10. The three main building blocks of the simulator are the transmitter chain, the channel model, and the receiver chain. TBs and their associated SCIs are generated and processed by the transmitter chain, passed through the channel model, and decoded by the receiver chain. For each SNR value, the TB plus SCI transmission-reception operations are performed N_{slots} consecutive times to experience several channel and noise realizations, i.e., employing the Monte Carlo approach. As a result, the BLER curves capture the average PHY layer performance of NR-V2X Sidelink communications over the specified range of SNR values. The SNR range is a configurable parameter of our simulator. For each SNR point, the BLER is computed as the fraction of incorrectly received transmissions with respect to the total. The NR-V2X SL LL simulator separately evaluates the 1st-stage SCI, 2nd-stage SCI and TB BLER performance.

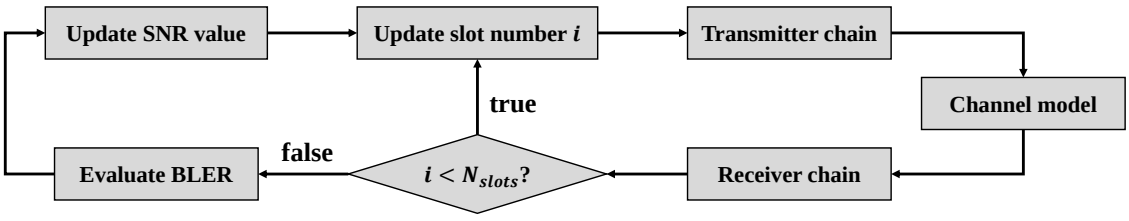

Figure 1.10: NR-V2X SL LL simulator workflow.

1.7.1.1 Transmitter Chain

The transmitter chain is responsible for the encoding, modulation, and transmission of the generated data. Differently with respect to 5G NR UL and DL, NR-V2X SL multiplexes the transmission of PSCCH and PSSCH on time and frequency adjacent REs, as explained

in Subsection 1.4.2.1. Therefore, the transmitter chain of our NR-V2X SL LL simulator features a more complex structure with respect to existing NR UL/DL link-level simulators [31, 50]. Fig. 1.11 outlines the different modules included within the transmitter chain of the NR-V2X SL LL simulator.

1) *1st-stage SCI Encoding*: The operations at the transmitter chain start with the encoding of the 1$^{\text{st}}$-stage SCI. The "1$^{\text{st}}$-stage SCI Encoding" module implements a custom function that formats the content of the 1$^{\text{st}}$-stage SCI following the SCI format 1-A definition reported in Subsection 1.4.2.1. Following NR-V2X SL specifications (Section 8.3.2-8.3.4 in [41]), the "1$^{\text{st}}$-stage SCI Encoding" block performs the following sequence of operations on the content of the 1$^{\text{st}}$-stage SCI: 24-bit Cyclic Redundancy Check (CRC) attachment, Polar encoding, and Polar code rate matching. These operations are inherited from NR DL specifications, although there are some differences. Scrambling is not performed during the CRC attachment step, and the target code rate evaluation is based on the number of REs available within the PSCCH (excluding the PSCCH-DMRS REs).

2) *2nd-stage SCI Encoding*: The "2$^{\text{nd}}$-stage SCI Encoding" module includes custom functions that allow the formatting of the 2$^{\text{nd}}$-stage SCI according to either the SCI format 2-A or the SCI format 2-B reported in Subsection 1.4.2.1. Regardless of the format, the 2$^{\text{nd}}$-stage SCI bits undergo the same sequence of CRC attachment, Polar coding and rate matching operations employed for the 1$^{\text{st}}$-stage SCI, although with some minor modifications. According to 3GPP specifications (Section 8.4.2-8.4.4 in [41]), 2$^{\text{nd}}$-stage SCI bits do not undergo interleaving during the rate matching phase. Moreover, the target code rate of the 2$^{\text{nd}}$-stage SCI is determined as a function of several parameters such as the associated TB code rate, the size of the selected 2$^{\text{nd}}$-stage SCI format, and the number of subcarriers occupied by the PSSCH. The formatting of the 2$^{\text{nd}}$-stage SCI and the calculation of its target code rate have been implemented from scratch.

3) *TB Encoding*: Also in this case, the implementation of the "TB encoding" module has required custom functions that are not available within MATLAB's 5G Toolbox. These functions are responsible for generating and encoding the TB information bits. First, the TB size is determined following the two-step process defined in Section 8.1.3.2 of [40]. The TB size depends on the length of the rate matched 2$^{\text{nd}}$-stage SCI, the PSSCH-DMRS pattern, the number of PSCCH REs, the selected MCS index, and the number of available REs in the PSSCH. Next, TB encoding is performed following a sequence of operations inherited from NR UL specifications. In our NR-V2X SL LL simulator, TB encoding consists of 24-bit CRC attachment, LDPC encoding and rate matching (see Section 8.2 of [41]). With respect to NR UL, the NR-V2X SL implementation disables the limited buffer rate matching option during the rate matching step.

4) *Data and Control Multiplexing*: The "Data and control multiplexing" module is responsible for multiplexing the output of the "2nd stage SCI Encoding" and "TB encoding" blocks into a single codeword. This is a specific feature of NR-V2X SL which has been implemented from scratch in our simulator. The data and control multiplexing operations depend on the number of transmission layers (Section 8.2.1 in [41]). When only 1 transmission layer is considered, the 2nd-stage SCI and TB bits are concatenated. In the second layer, the 2nd-stage SCI is replaced by placeholder bits.

5) *PSCCH-PSSCH Multiplexing*: Inside the "PSCCH-PSSCH Multiplexing" module, the control channel and the shared channel codewords are modulated and allocated on the NR time-frequency grid together with their associated DMRS symbols. The control channel codeword is provided by the "1st-stage SCI Encoding" block and accommodated over the PSCCH. The shared channel codeword is generated by the "Data and control multiplexing" module and allocated over the PSSCH. As shown in Fig. 1.11, the PSCCH and the PSSCH are processed independently until they are mapped on physical resources. In this module, the implementation of each function follows the NR-V2X SL specifications reported in [42]. The operation of each block is summarized here below:

   - *PSCCH Modulation.* Prior to modulation, the bits of the 1st-stage SCI codeword undergo a scrambling operation. In this function, the scrambling is implemented by means of an XOR operation between the codeword and a pseudo-random scrambling sequence, as in NR UL/DL. In our NR-V2X SL implementation, the pseudo-random sequence of bits is initialized with a fixed value following the specifications in Section 8.3.2.1 of [42]. Then, the block of scrambled bits is QPSK-modulated to generate the PSCCH OFDM symbols.
   - *PSCCH-DMRS.* This module exploits a new function to generate the DMRS symbols associated with the PSCCH. The implementation of the PSCCH-DMRS symbols follows the SL specifications in [42]. Accordingly, the PSCCH-DMRS symbols are generated using a pseudo-random sequence whose initialization value is set by the scramble ID provided by higher layers. In our implementation, we set the scramble ID to 0.
   - *PSSCH Modulation.* In the "PSCCH modulation" module, the 2nd-stage SCI and TB bits are separately scrambled employing two different pseudo-random scrambling sequences. Each scrambling sequence is initialized using the decimal representation of the CRC associated with the 1st-stage SCI. Then, the scrambled bits of the 2nd-stage SCI are modulated using QPSK modulation, whereas the modulation order of the TB bits depends on the selected MCS. After modulation, the PSSCH OFDM symbols undergo a layer mapping and a non-codebook-based precoding procedure before being allocated on the corresponding resources.

- *PSSCH-DMRS*. The "PSSCH-DMRS" module has been implemented to allow all the DMRS time pattern configurations reported in SL specifications (see Table 8.4.1.1.2-1 in [42]). A new function has been developed to generate the DMRS symbols associated with the PSSCH. In the NR-V2X SL case, the generation of the PSSCH-DMRS symbols is based on a pseudo-random sequence which is initialized using the decimal representation of the 1st-stage SCI CRC.
- *Mapping to physical resources*. This module represents the last operation of the "PSSCH-PSCCH Multiplexing" block, and it is responsible for the allocation of the PSSCH, PSCCH, PSSCH-DMRS and PSCCH-DMRS OFDM symbols on the corresponding REs within the NR-V2X SL time-frequency grid. To do so, we have developed a new function that implements the time-frequency organization of radio resources illustrated in Subsection 1.4.2.1. In our implementation, the NR-V2X SL time-frequency grid is represented as matrix where each element corresponds to a distinct RE. The "Mapping to physical resources" block assigns the RE indexes to the different components. AGC and Guard symbols are also included (see Fig. 1.3).

Once that the PSSCH and the PSCCH have been mapped to their respective REs along with the associated DMRS symbols, the time-frequency resource grid is converted into a time-domain CP - OFDM waveform using MATLAB's built-in UL/DL functions. The transmission of the CP-OFDM waveform is the last step of the transmitter chain.

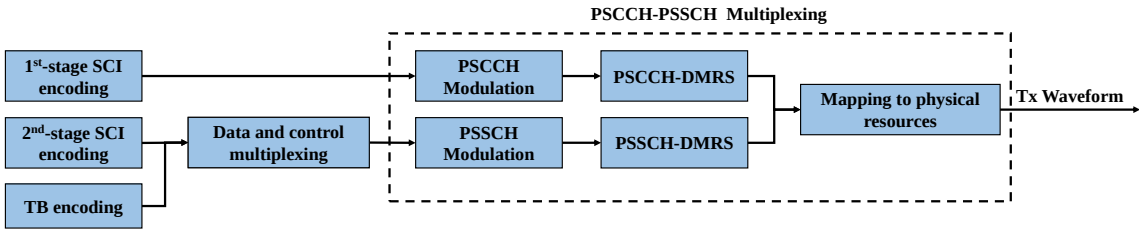

Figure 1.11: Transmitter chain.

1.7.1.2 Channel Model

In order to capture the small-scale fading effects which impair V2X communications, the NR-V2X SL LL simulator employs the CDL channel models defined by 3GPP and reported in Subsection 1.5.2.

In the NR-V2X SL LL simulator, each channel model is implemented with a custom configuration of MATLAB's built-in CDL system object, originally defined for 5G NR UL/DL simulations. In addition, each CDL channel model can be customized with specific maximum Doppler shift and Tx-Rx antenna array configurations. The setting of the maximum Doppler shift reflects different relative speeds between the transmitting and receiving vehicles. The Tx-Rx antenna array property allows the configuration of the

number of radiating elements and their radiation pattern. With respect to 5G NR UL/DL simulations, the CDL channel models employed in our simulator feature a horizontal, rather than vertical, MIMO arrangement of isotropic or V2X-specific antenna models [5]. In addition to the impairments introduced by the CDL models, the channel model employed in our NR-V2X SL LL simulator features an Additive White Gaussian Noise (AWGN) contribution. The noise variance depends on the considered SNR value.

1.7.1.3 Receiver Chain

The receiver chain processes the noisy and channel-impaired CP-OFDM waveform that is captured by each of the receiving antennas following the sequence of operations described in Fig. 1.12. Fig. 1.12 shows that the recovery of the PSSCH and PSCCH information is carried out on two different receiver chain branches. This is the case because the PSCCH and the PSSCH employ distinct DMRS patterns and undergo separate channel estimation and equalization operations. This is a relevant difference with respect to 5G NR UL and DL simulations, and it required significant implementation efforts. In both the PSSCH and PSCCH case, the DMRS symbols are used to perform timing and channel estimation. After having estimated the channel response, equalization is used to reconstruct the originally transmitted PSSCH and PSCCH symbols constellations. Specifically, the NR-V2X SL LL simulator employs Minimum Mean Square Error (MMSE) equalization.

After the equalization step, demodulation and 1$^{\text{st}}$-stage SCI decoding are the next operations within the receiver chain, if we concentrate on the PSCCH branch. The "PSCCH demodulation" and "1$^{\text{st}}$-stage SCI decoding" modules perform the inverse operations with respect to their transmitter chain counterparts, i.e., the "PSCCH modulation" and "1$^{\text{st}}$-stage SCI encoding" modules reported in Fig. 1.11. Specifically, the Polar decoding process uses a CRC-aided successive-cancellation list decoder of length L. In our implementation, we set L to its maximum allowed value, i.e., $L = 8$, as it guarantees the best error correction performance [51]. On the PSSCH branch, the "PSSCH demodulation", "2$^{\text{nd}}$-stage SCI decoding" and "TB decoding" modules follow the "PSSCH equalization" step. The "PSSCH demodulation" module includes a custom piece of code that separates the 2$^{\text{nd}}$-stage SCI and TB symbols. This is performed to guarantee a proper demodulation process, since the 2$^{\text{nd}}$-stage SCI and TB symbols might employ different modulation orders. Recall that the TB modulation depends on the selected MCS, whereas the 2$^{\text{nd}}$-stage SCI always utilizes QPSK modulation. After the demodulation step, the 2$^{\text{nd}}$-stage SCI and TB bits undergo their corresponding decoding operations. The 2$^{\text{nd}}$-stage SCI employs the same Polar decoder configuration used for the 1$^{\text{st}}$-stage SCI. On the other hand, LDPC decoding is carried out using the "Normalized min-sum" algorithm and by setting the maximum number of LDPC decoding iterations to 6. The configuration of the LDPC decoder follows the default settings employed by MATLAB's 5G toolbox in 5G NR UL simulations. Both the PSSCH and the PSCCH receiver branch are based on soft-decision decoding techniques. It is worth pointing out that the PSSCH decoding can

be performed only when the associated PSCCH information has been correctly decoded, according to the standard. This is the case since the 1st-stage SCI is transmitted on the PSCCH and it carries relevant information for the decoding of the PSSCH (see Subsection 1.4.2.1). Such decoding procedure is implemented in the NR-V2X SL LL simulator. In addition, our simulator allows to evaluate the PSSCH performance independently of the PSCCH. Accordingly, we can provide 2nd-stage SCI and TB BLER curves that are not conditioned on the correct decoding of the associated 1st-stage SCI information.

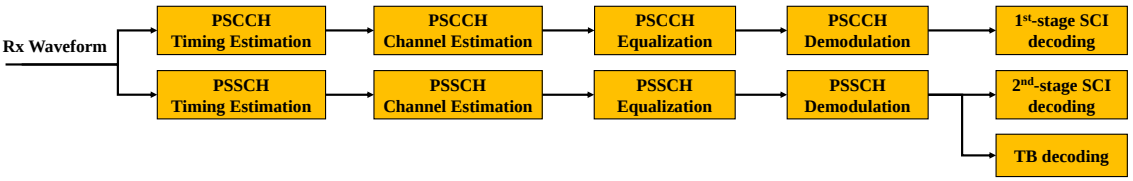

Figure 1.12: Receiver chain.

1.7.2 NR-V2X SL Profile

NR-V2X SL communications are characterized by a large number of configurable parameters. As a result, a vast number of different system choices has been proposed and analyzed over the last years, e.g., see [28,30,52], hindering the development of a common evaluation framework for the assessment of the NR-V2X SL achievable performance. In this direction, 3GPP and ETSI working groups devoted significant efforts towards the definition of a common NR-V2X SL profile, as witnessed by the documents in [5] and [24], respectively. The introduction of a common NR-V2X SL profile fosters the development of common evaluation guidelines, hence allowing a more fair and accurate analysis of the strengths and limitations which characterize NR-V2X SL communications. In addition, the availability of a common NR-V2X SL profile is particularly useful when out-of-coverage V2V communications are considered, as it provides a default system configuration that allows vehicles to exchange information also without the intervention of the cellular infrastructure.

Table 1.2 summarizes the common NR-V2X SL profile parameters that we employed in our LL simulations. Due to space constraints, we cannot report the full set of parameters, but the complete list is available in [5] and [24]. The NR-V2X SL profile considers a 20 MHz wide channel located at 5.9 GHz, a 30 kHz subcarrier spacing and 14 OFDM symbols per slot. According to [48], the total number of available RBs is equal to 51. Since the recommended subchannel size is 12 RBs [24], the total number of available subchannels in the NR-V2X SL profile is 4. In each LL simulation, the number of employed subchannels, N_{sub}, can vary from 1 to 4. The PSCCH is 12 RBs long and occupies 3 consecutive OFDM symbols.

Despite providing a default configuration for many parameters, it is worth pointing

Table 1.2: Default NR-V2X SL Profile.

Parameter	**Value**
Bandwidth	20 MHz
Operating Frequency	5.9 GHz
Subcarrier spacing	30 kHz
OFDM symbols per slot	14
Subchannel size	12 RBs
Available subchannels	4
PSCCH RBs	12 RBs
PSCCH OFDM symbols	3
MCS Table (c)	Table 5.1.3.1-2
PSSCH-DMRS Pattern (c)	{2}
PSSCH-DMRS scramble ID	0
Vehicle speed, Highway (c)	70, 140 km/h
Vehicle speed, Urban (c)	60 km/h
Number of Tx layers (c)	1
Number of Tx-Rx antennas (c)	2 Tx 4 Rx
Antenna type (c)	Isotropic

out that some settings are still left to each specific implementation. Such configurable parameters have been highlighted with the letter (c) in Table 1.2. For example, the MCS used for the TB transmission can be selected from either Table 5.1.3.1-1 or Table 5.1.3.1-2 according to the standard. Both tables are reported in [40]. In our LL simulations we decided to use Table 5.1.3.1-2 because it provides a wider, and more exhaustive, range of MCS choices. In detail, the selected table allows MCS values that range from QPSK modulation with 120/1024 target code rate (indicated as QPSK-120) to QPSK-602/1024 (more simply indicated as QPSK-602), from 16QAM-378 to 16QAM-658, from 64QAM-466 to 64QAM-873, and from 256QAM-682.5 to 256QAM-948. Similarly, the choice of the PSSCH-DMRS time pattern is up to each implementation. Vehicles can either use 2, 3, or 4 PSSCH-DMRS symbols in the same slot. However, when a single subchannel is considered, the PSSCH does not support the 3 and 4 DMRS options if the subchannel size is smaller than 20 RBs (see Subsection 1.4.2.1). The subchannel size is 12 RBs in the NR-V2X SL profile. For this reason, we decided to employ the 2 DMRS time pattern (indicated as {2} by the standard [40]) in all our LL simulations, regardless of the considered subchannels' number.

Another relevant simulation setting concerns the vehicles' speed. The relative speed between vehicles affects the maximum Doppler shift experienced by a V2V link, as anticipated in Subsection 1.7.1.2. According to [5], the vehicles' speed in the Highway environment can be either set to 70 or 140 km/h. As a result, the relative speed (v_{rel})

between vehicles can be 0, 70, 140, or 280 km/h. In the Urban environment, the vehicles' speed is fixed to 60 km/h and therefore limits the v_{rel} values to 0, 60, and 120 km/h. Last, our LL simulations consider one transmission layer and 2 Tx - 4 Rx antennas, as this is the recommended configuration in [5]. The radiation pattern of each antenna element is modeled as isotropic.

1.7.3 Relevance of the NR-V2X SL LL Results

In this Subsection, we highlight the relevance of the NR-V2X SL LL results presented in this Chapter. To do so, we analyze typical circumstances that require the use of dedicated and accurate BLER vs SNR curves during NR-V2X SL system-level simulations. Due to space constraints, our analysis is supported only by a selection of BLER vs SNR curves, but the full set of results is publicly available [8]. A subset of the presented results is validated against the few NR-V2X SL BLER curves available in the literature. Unfortunately, more sophisticated validation tools are not available. Unless otherwise stated, we consider the Highway-LOS channel model, $v_{rel} = 0$ km/h, $N_{sub} = 1$ and four different MCS selections, namely: QPSK-308, 16QAM-490, 64QAM-616, and 256QAM-797. We also assume that the decoding of the PSSCH does not depend on the correct decoding of the PSCCH. Therefore, the 2$^{\text{nd}}$-stage SCI and TB BLER curves presented in this study are not conditioned on the correct recovery of the associated 1$^{\text{st}}$-stage SCI information.

1.7.3.1 Decoding of 1$^{\text{st}}$-stage SCI, 2$^{\text{nd}}$-stage SCI and TB

NR-V2X SL transmissions include a TB and its associated 1$^{\text{st}}$-stage and 2$^{\text{nd}}$-stage SCI. Besides being allocated on different REs, we observed earlier that these three components undergo specific encoding and modulation mechanisms. Therefore, they are characterized by a different PHY layer performance and system-level simulations require dedicated BLER vs SNR curves to separately evaluate the 1$^{\text{st}}$-stage SCI, 2$^{\text{nd}}$-stage SCI, and TB decoding probability. The NR-V2X SL LL simulator employed in this study is able to separately evaluate the 1$^{\text{st}}$-stage SCI, 2$^{\text{nd}}$-stage SCI and TB BLER curves, as Fig. 1.13 demonstrates. The curves reported in Fig. 1.13 show that the 1$^{\text{st}}$-stage SCI transmission is the most robust, experiencing the best BLER performance. As a matter of fact, the 1$^{\text{st}}$-stage SCI employs the most robust MCS, i.e., QPSK-84. Moreover, since the 1$^{\text{st}}$-stage SCI length and the size of the PSCCH do not depend on the TB's MCS, the 1$^{\text{st}}$-stage SCI BLER curve does not modify when moving from Fig. 1.13(a) to (d). Conversely, the code rate of the 2$^{\text{nd}}$-stage SCI depends on the MCS of the associated TB. In detail, the code rate of the 2$^{\text{nd}}$-stage SCI is always larger than the one employed by the 1$^{\text{st}}$-stage SCI, thus resulting in a less robust transmission as highlighted by Fig. 1.13. On the other hand, these figures show that the 2$^{\text{nd}}$-stage SCI is always characterized by a better PHY layer performance than the associated TB. In Fig. 1.13(a), the 2$^{\text{nd}}$-stage SCI and the TB employ the same modulation scheme, but the 2$^{\text{nd}}$-stage SCI features a more robust, i.e.,

lower, code rate. In Fig. 1.13(b)-(d), the 2$^{\text{nd}}$-stage SCI is always transmitted using QPSK modulation, whereas the TB employs less robust modulation schemes, thus achieving a worse BLER performance. It is worth highlighting that the TB BLER curves reported in Fig. 1.13(a) and Fig. 1.13(c) are in line with the results presented in [33]. Note that the distance between the TB and the 1$^{\text{st}}$-stage SCI BLER curves increases as we move from Fig. 1.13(a) to Fig. 1.13(d). This is the case since the TB BLER performance significantly deteriorates as more spectrally efficient MCS choices are considered. It is important to note that the TB BLER departs from 1 only when the correct decoding of the 1$^{\text{st}}$-stage SCI is guaranteed, i.e., BLER = 0, in all figures. Such result reveals that the content of the 1$^{\text{st}}$-stage SCI can be correctly recovered also when the decoding of the associated TB fails. In NR-V2X SL communications, vehicles employ the 1$^{\text{st}}$-stage SCI to reserve their transmission resources and its correct decoding is crucial to guarantee the proper functioning of the resource allocation schemes employed in NR-V2X Mode 2 (see [32] for further details). It is also worth pointing out that, since the PSSCH can be correctly decoded only when the PSCCH BLER is equal to zero, the curves reported in Fig. 1.13 do not modify when the decoding of the PSSCH is conditioned on the correct decoding of the PSCCH, as the standard mandates.

1.7.3.2 Tx-Rx Relative Speed

When a realistic system-level simulation environment is considered, the vehicles' dynamics is characterized by different travelling velocities and directions. Depending on the speed and travelling direction of the transmitting and receiving vehicles, the V2V link can experience different relative speed values. At PHY layer, variations in the relative speed affect the maximum Doppler shift that impairs the transmitted waveform. Depending on the (PSSCH/PSCCH)-DMRS pattern and on the MCS employed by the TB, Doppler shifts can significantly deteriorate the BLER performance. In this regard, NR-V2X SL LL curves are instrumental to accurately capture the relative speed impact in system-level simulations. The BLER performance of the TB and of the 1$^{\text{st}}$ and 2$^{\text{nd}}$-stage SCIs considering different Tx-Rx relative speed values is examined next.

First, let us focus on Fig. 1.14 where the TB PHY layer performance is reported considering the four relative speed values examined in the Highway environment, i.e., $v_{rel} = \{0, 70, 140, 280\}$ km/h. Fig. 1.14(a) quantifies the impact of the relative speed on the TB BLER curves in the QPSK-308 case. Despite being a robust MCS choice, the BLER performance deteriorates as the relative speed increases. The employed PSSCH-DMRS pattern is not able to effectively combat the Doppler shift impairments which affect the V2V communication. For example, the SNR corresponding to BLER = 10^{-2} moves from -1.8 dB to 5.5 dB when the relative speed increases from $v_{rel} = 0$ km/h to $v_{rel} = 280$ km/h. Such 7.3 dB shift can have a significant impact on system-level simulations. When the MCS increases to 16QAM-490, the impact of the Doppler shift on the TB PHY layer performance is more evident. In Fig. 1.14(b), the BLER settles at $2 \cdot 10^{-3}$ when $v_{rel} = 280$

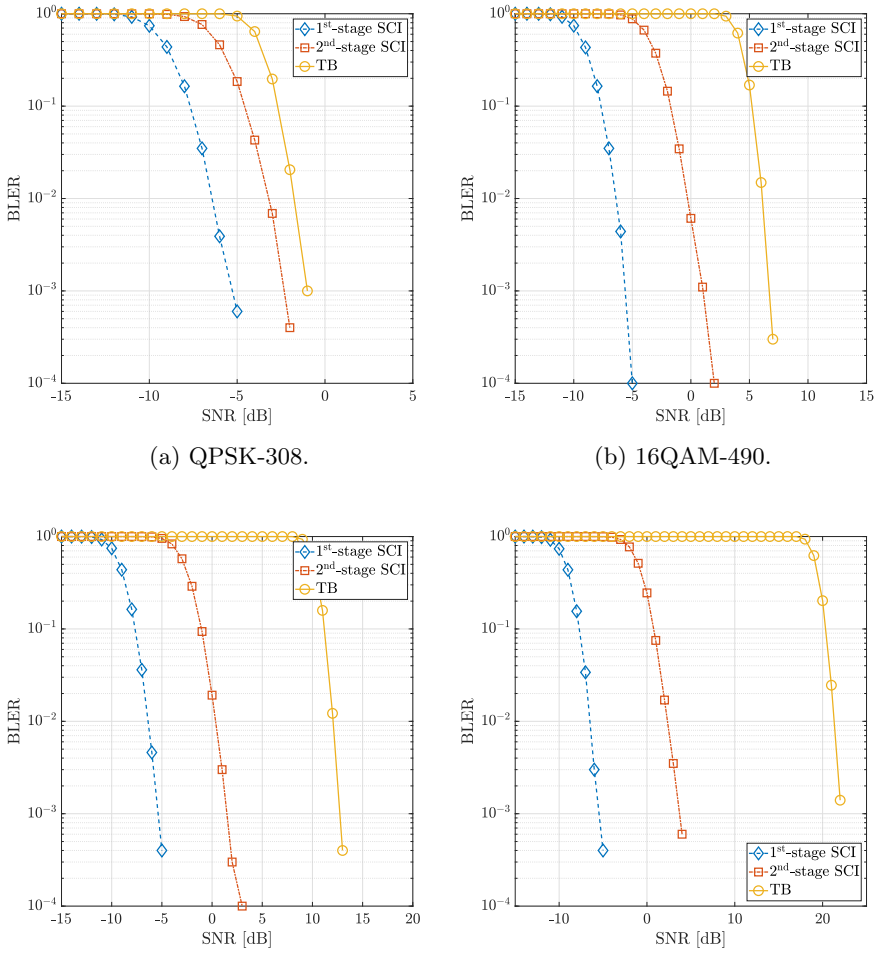

(a) QPSK-308.

(b) 16QAM-490.

(c) 64QAM-616.

(d) 256QAM-797.

Figure 1.13: 1st-stage SCI, 2nd-stage SCI and TB: BLER vs SNR curves for different MCS choices.

km/h and the SNR is higher than 25 dB. The presence of an error floor means that the correct decoding of the received packets cannot be always guaranteed. The negative impact of the Tx-Rx relative speed on the PHY layer performance is further highlighted in Fig. 1.14(c) and Fig. 1.14(d), where the 64QAM-616 and the 256QAM-797 MCS choices are analyzed, respectively. In Fig. 1.14(c), the correct decoding of the received packet can only be guaranteed if v_{rel} is not larger than 140 km/h. When $v_{rel} = 280$ km/h, the considered PSSCH-DMRS pattern is not sufficiently robust to compensate for the Doppler shift and the BLER is fixed at 1, thus not allowing the correct reception of any packet, regardless of the SNR value. When the least robust 256QAM-797 MCS is considered, Fig. 1.14(d) shows that the BLER is fixed at 1 starting from $v_{rel} = 140$ km/h. In this case, a V2V communication is feasible only if the Tx-Rx relative speed is smaller than 140 km/h. Although evaluated on a different channel model, the impact of the Tx-Rx relative speed on the TB PHY layer performance illustrated in Fig. 1.14 is characterized by the same trends reported in [33].

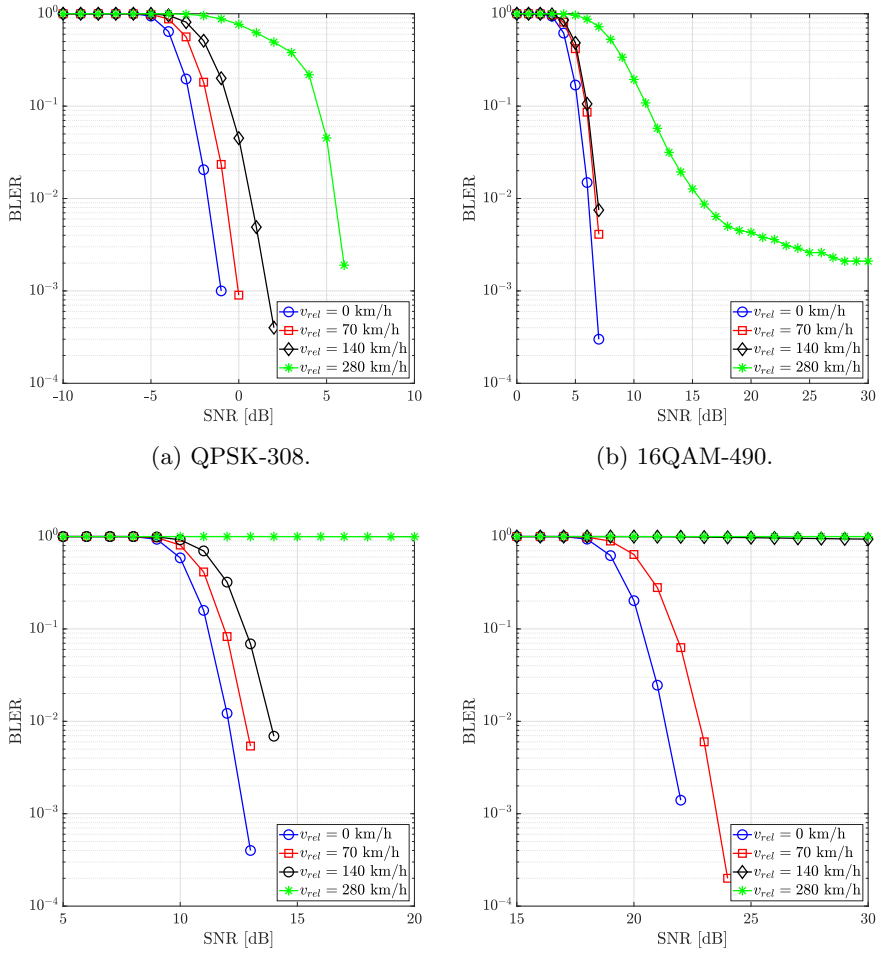

(a) QPSK-308.

(b) 16QAM-490.

(c) 64QAM-616.

(d) 256QAM-797.

Figure 1.14: Impact of the Tx-Rx relative speed on the TB BLER performance.

Next, Figs. 1.15(a) and 1.15(b) analyze the impact of the relative speed on the 2$^{\text{nd}}$-stage and 1$^{\text{st}}$-stage SCI. The 2$^{\text{nd}}$-stage SCI is multiplexed on the PSSCH together with the associated TB and its decoding process leverages the same PSSCH-DMRS pattern considered in the TB analysis. Since the 2$^{\text{nd}}$-stage SCI always employs QPSK modulation, the impact of the relative speed on its performance is similar to the one reported in Fig. 1.14(a), where the TB BLER was analyzed considering the same modulation order. On the other hand, the 1$^{\text{st}}$-stage SCI is transmitted on the PSCCH and employs a more robust PSCCH-DMRS pattern with respect to its PSSCH counterpart. In addition, the 1$^{\text{st}}$-stage SCI is encoded with a high-redundancy code rate and is transmitted using QPSK modulation. As a result, the transmission of the 1$^{\text{st}}$-stage SCI is extremely robust and can cope with high relative speeds without suffering any BLER degradation, as Fig. 1.15(b) shows. In this figure, the 1$^{\text{st}}$-stage SCI BLER vs SNR curve does not vary for increasingly large relative speeds. The robustness of the 1$^{\text{st}}$-stage SCI transmission against Doppler shifts has also been demonstrated in [53]. This Subsection has showed that, when large

relative speed values are considered, the 1st-stage SCI can be correctly decoded even when the associated TB and 2nd-stage SCI decoding inevitably fails due to error floors in the BLER performance. As already mentioned in Subsection 1.7.3.1, the correct decoding of the 1st-stage SCI is instrumental to maximize the effectiveness of the resource allocation schemes, as it minimizes the risk of collisions in NR-V2X SL communications.

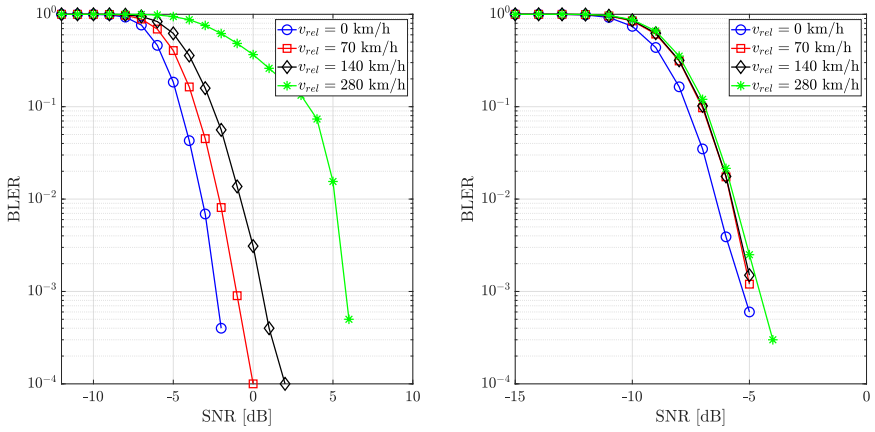

(a) 2nd-stage SCI. (b) 1st-stage SCI.

Figure 1.15: Impact of the Tx-Rx relative speed on the 2nd-stage and 1st-stage SCI BLER performance.

1.7.3.3 Channel Models

As highlighted in Subsection 1.5.2, 3GPP guidelines have introduced five different CDL channel models for the evaluation of NR-V2X SL communications. In system-level simulations, a V2V communication can occur in different propagation environments, i.e., Highway or Urban, and different channel states, i.e., LOS, NLOSv, or NLOS. It is important to highlight that, during system-level simulations, the SNR experienced by the 1st-stage SCI, 2nd-stage SCI, and TB transmission depends on the pathloss and shadowing ($PL + SH$) contributions reported in Subsection 1.5.2. Accordingly, the average SNR that characterizes each channel model can be significantly different. For example, when the Tx-Rx distance is 1000 m, the average SNR is 16.8 dB and 7.8 dB in the Highway LOS and NLOSv cases. At 1000 m, the SNR is 21.7 dB, 12.7 dB and -16.7 dB in the Urban LOS, NLOSv and NLOS cases.

In addition to different pathloss and shadowing contributions, each channel model is also defined by different fast-fading parameters, i.e., different CDL descriptions. In this regard, the results reported in this Subsection highlight the importance of having dedicated BLER curves able to capture the impact of fast-fading impairments on the PHY layer performance. First, Fig. 1.16(a) illustrates the impact of a different channel state on the BLER performance in the Highway setting. When the attention is confined to a single MCS choice, Fig. 1.16(a) shows that the difference between LOS and NLOSv

curves can be as large as 2.6 dB when the BLER is 10^{-2} and QPSK-308 is considered. As expected, NLOSv curves have a worse BLER performance with respect to their LOS counterparts. Note that all the MCS values reported in this figure are characterized by the same trend and that the gap between LOS and NLOSv curves reported in Fig. 1.16(a) is in line with the results presented in [23].

Next, Fig. 1.16(b) concentrates on the Urban setting. As in Fig. 1.16(a), the Urban-NLOSv channel has a worse BLER performance with respect to its Urban-LOS counterpart. In this case, the difference between LOS and NLOSv curves is 3 dB when the BLER is 10^{-2} and QPSK-308 is examined. With respect to the Highway case, the Urban setting is characterized by an additional NLOS channel state. In the NLOS state, the BLER performance is the worst, and the gap between LOS and NLOS curves can become quite relevant: for example, when SNR = 15 dB, the 64QAM-616 NLOS BLER is $4 \cdot 10^{-2}$ whereas the corresponding LOS BLER has already reached zero, thus guaranteeing the correct decoding of the received packet.

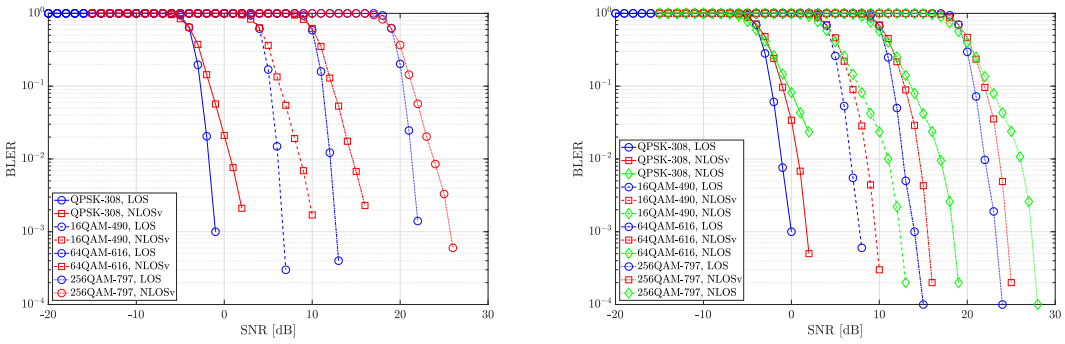

(a) Highway. (b) Urban.

Figure 1.16: Comparison of channel models: BLER vs SNR curves.

1.7.3.4 MCS Selection

In NR-V2X SL specifications, the standard does not provide any indication regarding the MCS selection, which is left to each specific implementation [54]. As a result, vehicles are allowed to change the MCS on a TB-basis during system-level simulations: for example, a vehicle might use a more spectrally efficient MCS to fit a larger amount of information in the same number of subchannels. However, the MCS cannot be naively changed without properly modeling its implications at PHY layer. In this regard, the impact that the MCS selection has on the TB BLER performance is quantified next. Recall from Subsection 1.4.2.1 that the 1$^{\text{st}}$-stage SCI employs a fixed MCS and the 2$^{\text{nd}}$-stage SCI is always transmitted using QPSK modulation. Therefore, the MCS employed by the TB does not affect the 1$^{\text{st}}$-stage SCI PHY layer performance and has little impact on the 2$^{\text{nd}}$-stage SCI BLER.

Fig. 1.17 reports the TB BLER curves obtained considering a representative set of

MCS choices selected from the NR-V2X SL profile MCS table. Due to space constraints, we could not report all the MCS values included in the table. For each modulation scheme, Fig. 1.17 reports the BLER performance that characterizes the smallest and the largest code rates. As a result, this figure allows to identify the interval of SNR values within which all the BLER curves corresponding to a given modulation scheme lie.

Fig. 1.17 sheds light on the BLER degradation associated with the selection of a more spectrally efficient MCS. Such effect is particularly relevant for system-level simulations. For example, a TB transmitted using 16QAM-658 requires an approximately 7 dB higher SNR to achieve the same BLER $= 10^{-1}$ performance than a TB employing QPSK-602 as MCS.

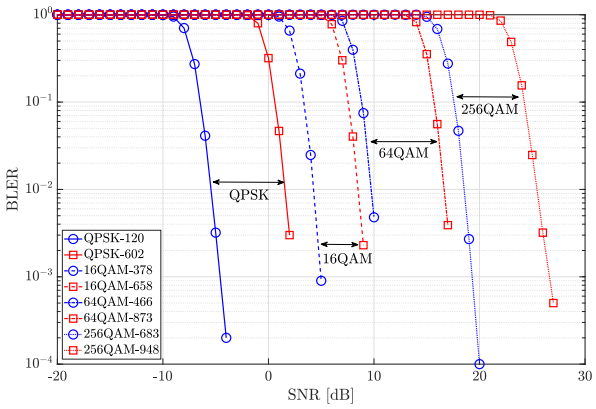

Figure 1.17: Impact of the MCS on the TB BLER performance.

1.7.3.5 Number of Subchannels

When the MCS is fixed, the number of subchannels required for the transmission of a TB in a system-level simulation changes depending on the size of the generated TBs. In this regard, the impact that a variable number of subchannels has on the LL performance of the TB is examined next. Fig. 1.18 reports the BLER curves obtained when the number of occupied subchannels ranges from 1 to 4, i.e., $N_{sub} = \{1, 2, 3, 4\}$.

When QPSK-308, 16QAM-490 and 64QAM-616 modulation and coding schemes are considered, Fig. 1.18 shows that the number of occupied subchannels does not affect the TB BLER performance, as all curves perfectly overlap. In this case, increasing the number of occupied subchannels does not have any impact on the correct decoding of the packet at PHY layer. The same trend also characterizes the BLER curves reported in [31] for NR-V2X UL/DL communications.

Conversely, the results obtained with the 256QAM-797 MCS show that the BLER performance deteriorates as N_{sub} increases from 1 to 4. In Fig. 1.18, the distance between the BLER curves can become quite relevant: for example, a TB accommodated over $N_{sub} = 4$ subchannels requires a 2 dB higher SNR to achieve a BLER of 10^{-2} with respect to the $N_{sub} = 1$ setting. Although a single code rate has been reported in this figure, we

verified that the impact of N_{sub} does not vary for different code rate values.

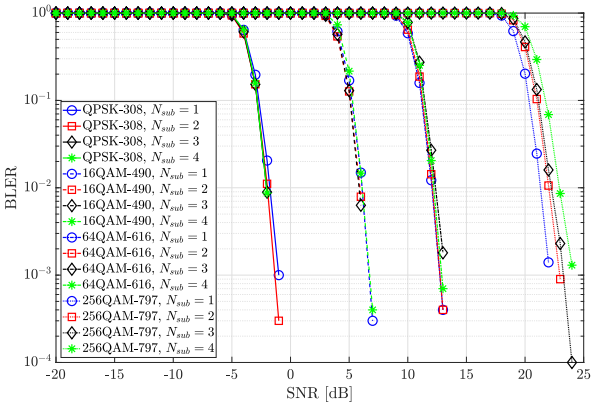

Figure 1.18: Impact of N_{sub} on the TB BLER performance.

1.8 Conclusions

This Chapter has first presented MoReV2X, an ns-3 module for the analysis of LTE-V2X and NR-V2X Sidelink communications. Its development adheres to the 3GPP specifications and performance assessment guidelines: as such, MoReV2X implements standard-compliant traffic and channel models, and evaluates a complete set of performance metrics. An exemplary set of results offered an interesting insight on the achievable performance of NR-V2X SL for two alternative SCS choices, when periodic and aperiodic traffic patterns are examined.

In addition, this Chapter has presented the first complete set of NR-V2X SL BLER vs SNR curves. The results are openly available and have been obtained with a custom LL simulator based on MATLAB. The obtained BLER curves have been employed to extend the PHY layer abstraction model of MoReV2X, and cover the parameter settings recommended by the NR-V2X SL profile. Grounded on a representative selection of simulation results, this Chapter also highlighted the impact of PHY layer aspects on the BLER performance. The LL analysis of NR-V2X SL quantified the impact of the transmitter-receiver relative speed, channel model, MCS selection, and subchannel number on the BLER performance. The LL analysis revealed that simplified and incomplete PHY abstraction models can lead to inaccurate results during system-level simulations.

Chapter 2

Assessing the Achievable Performance of C-V2X Sidelink Communications

The digital transformation of ITSs witnessed a fast-paced evolution over the last years. V2X services evolved from basic safety-related functionalities to eV2X applications [55] that rely on cooperative perception and maneuvering.

However, the research on the C-V2X SL technology has not proceeded with the same impetus. Many existing works have investigated the performance of C-V2X SL employing simplified simulation assumptions, or through real-world field tests performed in controlled environments. As of today, the achievable performance of the C-V2X SL technology has not been exhaustively assessed, and the impact of several MAC and PHY layer aspects on the system performance needs to be characterized. As a result, the superiority of the C-V2X SL technology over its IEEE 802.11p competitor has not been clearly demonstrated yet, and the C-V2X SL technology is being deployed only in some parts of the world. For example, the C-V2X SL technology has already been adopted in China and, although it was initially discarded by the US, the Federal Communications Commission (FCC) has recently allocated a portion of the ITS band to C-V2X SL.

2.1 Research Motivation and Main Contribution

This Chapter leverages the MoReV2X simulator to provide an exhaustive analysis of C-V2X SL communications. The goal is to assess the achievable performance of LTE-V2X SL and NR-V2X SL communications, identifying the strengths and the limits of the two different technologies. In particular, this Chapter concentrates on the MAC sublayer features which characterize the LTE-V2X Mode 4 and NR-V2X Mode 2 distributed resource allocation techniques.

First, this Chapter concentrates on LTE-V2X Mode 4 and its SPS strategy, and addresses the key question of how to accommodate aperiodic traffic in LTE-V2X Mode 4 [9].

To the authors' knowledge, this is an aspect that had not been investigated in the literature at the time of this study, and it was therefore worth being explored. The performance of LTE-V2X Mode 4 had been thoroughly analyzed by several works in the presence of periodic packet transmissions. Yet, the hypothesis of periodic traffic is quite unrealistic when V2X applications are considered. For example, the periodic generation of CAMs has been questioned through experimental data [14], and new models have been introduced to generate realistic, aperiodic CAM traces [47]. Moreover, the eV2X use cases introduced in Release 15 will foresee a mixture of periodic and aperiodic traffic. Indeed, information detected through on-board cameras, LIDARs, or RADARs, generates traffic flows that are stochastic in nature.

The obtained results reveal that LTE-V2X Mode 4 is not able to effectively serve aperiodic traffic. The periodic reservation of resources employed by the SPS strategy perfectly suits the dissemination of periodic traffic, but it fails when the generated traffic is aperiodic. For this reason, a novel strategy to jointly serve periodic and aperiodic traffic is put forth, under the design constraint of preserving the performance of periodic flows. The comparison between the proposed solution and two schemes that accommodate aperiodic traffic in accordance to the original SPS scheme is offered, providing numerical evidence of the superiority of the new strategy.

Then, this Chapter delves into the 5G domain and analyzes the latest MAC sublayer features introduced within the NR-V2X SL technology. These features are expected to mitigate the inefficiencies which characterized LTE-V2X Mode 4 [15, 56], and guarantee the effective dissemination of fixed and variable size aperiodic traffic. In this regard, this Chapter provides a qualitative analysis of the re-evaluation mechanism and sheds light on the impact that different traffic types and PDB requirements have on the operation of the SPS and DS strategies. The analytical insights are corroborated by simulation results:

1) This Chapter quantifies the impact of different traffic types on the NR-V2X Mode 2 SPS strategy. The obtained results reveal that the SPS strategy is not able to effectively serve aperiodic traffic also in NR-V2X Mode 2, despite the introduction of the re-evaluation mechanism.

2) The operation of the re-evaluation mechanism is also examined in this Chapter. The goal of this study is to advance the state of the art with an exhaustive analysis of the re-evaluation mechanism, assessing its impact on NR-V2X Mode 2 performance. To the authors' knowledge, this is the first study that analyses the effectiveness of the re-evaluation mechanism in detecting and avoiding packet collisions when both the SPS and DS strategies are considered.

   This study analyzes the re-evaluation mechanism considering periodic and (fixed or variable size) aperiodic packets, generated according to the 3GPP traffic models reported in Subsection 1.5.1. The study also evaluates the impact of retransmissions on the effectiveness of re-evaluation. Simulation results show that the re-evaluation mechanism is able to detect a negligible number of collisions in the periodic traffic

case. Conversely, re-evaluation can detect a larger number of collisions when vehicles generate aperiodic traffic, but it is not able to select new collision-free resources. In both cases, the re-evaluation mechanism does not improve the NR-V2X Mode 2 performance.

3) Last, this Chapter compares the performance of the SPS and DS strategies [12]. Differently from previous investigations, the distinctive intent of this study is to provide clear guidelines to understand for which traffic types the SPS approach can be profitably employed and when, on the contrary, the DS choice has to be preferred. To accomplish such a goal, the performance of the SPS and DS strategies is analyzed considering different traffic models and, for the first time, varying latency requirements. The SPS and DS schemes are also compared when fixed size periodic traffic and variable size aperiodic traffic sources coexist within the same NR-V2X system. In this case, a novel solution, termed Adaptive Scheduling (AS), is put forth: the proposed AS strategy allows vehicles to adapt their scheduling strategy on the basis of the traffic type they generate. The main findings that this study provides are the following: (i) differently from its SPS counterpart, the performance of the DS scheme is independent of the PDB requirements, no matter what traffic model is considered; (ii) the SPS approach represents the best solution for serving fixed size periodic traffic, whereas the DS strategy is the winning choice when fixed or variable size aperiodic traffic is considered; (iii) the newly devised AS approach consistently achieves better performance than the SPS and DS schemes.

The remainder of this Chapter is organized as follows: Section 2.2 reviews the C-V2X SL literature and Section 2.3 presents the study on the coexistence of periodic and aperiodic traffic in LTE-V2X Mode 4. Next, the qualitative analysis of NR-V2X Mode 2 and its MAC sublayer features is reported in Section 2.4. The performance of NR-V2X Mode 2 is exhaustively assessed in Section 2.5, where a vast set of simulation results quantifies the impact of different traffic types on the SPS strategy operation, provides a thorough analysis of the re-evaluation mechanism, and compares the performance of the SPS and DS strategies. Lastly, Section 2.6 draws the conclusions.

2.2 Related Work

Various works have investigated the performance of the LTE-V2X Mode 4 SPS strategy considering the transmission of fixed size periodic packets. In [57], Masegosa *et al.* assessed the system-level behavior of LTE-V2X Mode 4 in a Manhattan grid urban topology, comparing it against that of a random scheduling strategy. The authors of [58] provided an accurate analysis of the impact that different MAC and PHY layer parameters have on LTE-V2X Mode 4 when the periodic CAM dissemination is targeted. The authors of [59] proposed a novel set of analytical models to evaluate the SPS performance, whereas [60] highlighted the positive effects of full-duplex radios on LTE-V2X Mode 4. The perfor-

mance of the SPS strategy has been analytically characterized also by the authors of [61], concentrating on the hidden terminal problem. A novel geo-based scheduling scheme that allows vehicles to autonomously select their radio resources based on their location has been proposed in [62]; these studies too, considered periodic traffic only.

Over the last years, the impact of more realistic traffic patterns on the LTE-V2X Mode 4 performance has been analyzed by a handful of works. In [46], the authors analyzed the effectiveness of LTE-V2X Mode 4 for multi-hop delivery of asynchronous DENMs, in the presence of background, periodic CAM traffic. A similar issue was studied in [63], considering the benefit of short term sensing. However, DENMS are by no means comparable to stochastic packet transmissions. The impact of variable size aperiodic traffic on LTE-V2X Mode 4 has been analyzed in [56].

Under the 5G umbrella, few papers investigated the performance of the SPS strategy in NR-V2X Mode 2 [28, 30], showing that the effective dissemination of aperiodic traffic cannot be guaranteed, especially when the size of the generated packets is not fixed. In this regard, novel AI-based solutions [64] and alternative techniques [65,66] have been proposed to overcome the limitations that characterize the design of the SPS strategy. On the other hand, no work has analyzed the DS strategy so far. To the authors' knowledge, the only study that concentrated on the comparison between the SPS and DS strategies is [67], although it assumed fixed size packets without considering the 3GPP-defined traffic models [5]. Further and more importantly, all these works did not consider the re-evaluation mechanism, which is a mandatory MAC sublayer feature in NR-V2X Mode 2.

2.3 Coexistence of Periodic and Aperiodic Traffic in LTE-V2X Mode 4

In this Section, the coexistence of periodic and aperiodic traffic in LTE-V2X Mode 4 is analyzed [9]. The issue of coping with aperiodic flows in LTE-V2X Mode 4 is faced and a simple, standard-compliant approach is put forth. The throughput of the proposed resource allocation technique is analytically determined in the absence of transmission impairments, for the limiting case of aperiodic traffic only. The analysis outcomes are then used to forecast what to expect in a realistic vehicular scenario, when periodic flows are also present. In this Section, the comparison between the proposed solution and two schemes that accommodate aperiodic traffic in accordance to the SPS strategy is offered, providing numerical evidence of the superiority of the new strategy. It is analytically proved and verified by simulation that the performance of aperiodic traffic is insensitive to different latency requirements set for aperiodic packets. Lastly, this Section demonstrates that LTE-V2X Mode 4 is not able to effectively serve aperiodic traffic. Vehicles generating such traffic type experience low performance levels, even when moderate traffic densities are examined, unless very small aperiodic packets are considered.

2.3.1 Serving Aperiodic Traffic

2.3.1.1 The Proposed Solution

In this study, we put forth a reservation-less policy for the allocation of aperiodic traffic within the time-frequency grid described in Subsection 1.4. We assume that aperiodic and periodic traffic coexist and that vehicles transmitting periodic traffic implement the SPS strategy.

We further suppose that all aperiodic TBs are characterized by the same PDB, corresponding to the maximum latency their delivery can tolerate, given by N_t ms. As the resource reselection mechanism mandates, we require that every vehicle generating aperiodic traffic monitors the channel with no interruptions; for doing so, the vehicle continuously slides forward its sensing window S and updates the information collected through the received SCIs, the RSSI and the RSRP values. When the aperiodic TB is ready for transmission, the ego-vehicle relies on the most recent sensing window to build its own, current view of the available subchannels. The ego-vehicle runs the resource reselection process with a selection window whose duration is exactly $W = N_t$ ms, in order to respect the PDB; if more candidate subchannels are present, the resource selection is random, in line with the original algorithm. In our proposal however, the ego-vehicle sets the reselection counter C_{resel} to 1 and the keep probability P to 0. In other words, the ego-vehicle selects new resources for every generated TB and does not place any periodic reservations, as if it were to perform a one-shot communication. It follows that vehicles generating aperiodic traffic build their own map of available resources taking into account ongoing periodic traffic only, monitoring its reservations, and that their aperiodic transmissions do not conflict with them.

2.3.1.2 Limiting Analysis

It is now instructive to focus on the limiting condition when periodic traffic is absent and the previous strategy is employed to accommodate aperiodic traffic. In this scenario, we model the selection window as an $N_f \times N_t$ grid, where N_t represents the number of subframes in W, and N_f the number of RBs in the frequency channel. Next, we observe that: (i) the sensing window is useless here, owing to the lack of periodic reservations; (ii) the resource access policy is totally random.

Let us further assume that the aperiodic traffic is Poisson distributed, with an overall average arrival rate given by λ TB/s, and that every TB plus the associated SCI requires the assignment of R RBs within a subframe, where for the sake of simplicity N_f is a multiple integer of R. If we neglect transmission impairments, that is, consider ideal transmission conditions, the evaluation of the aggregate aperiodic throughput $\mathcal{S}_{ap}$ is brought back to multi-channel slotted Aloha, where the number of channels is $\frac{N_f}{R}$. As a matter of fact, define G as the average number of aperiodic TBs collectively generated

within a subframe of duration t_s, i.e.,

$$G = \lambda \cdot t_s \,. \tag{2.1}$$

In the simple case where $N_f > 1$ and $N_t = 1$, observe that the vulnerable period of the access strategy coincides with the subframe duration; indeed, all aperiodic TBs generated during a subframe will be transmitted within the next subframe. It readily follows that $\mathcal{S}_{ap}$ is given by

$$\mathcal{S}_{ap} = \left(\frac{N_f}{R}\right) \cdot \left(\frac{G}{\frac{N_f}{R}} \cdot e^{-\frac{G}{\frac{N_f}{R}}}\right) = G \cdot e^{-\frac{G}{\frac{N_f}{R}}} \,. \tag{2.2}$$

When $N_t > 1$, i.e., when the selection window is made of N_t consecutive subframes, the transmission attempts of aperiodic TBs generated during a given subframe are equally distributed over the next N_t subframes, each being subject to a Poisson traffic whose rate is $\lambda_i = \frac{\lambda}{N_t}$. On any subframe, this term adds to other $N_t - 1$ Poisson flows, originated within the previous N_t subframes, with every of them exhibiting a rate $\lambda_i = \frac{\lambda}{N_t}$. That is to say, the traffic poured on each of the N_t subframes is still Poisson, with rate

$$\sum_{i=1}^{N_t} \lambda_i = N_t \cdot (\frac{\lambda}{N_t}) = \lambda \,, \tag{2.3}$$

leading to the conclusion that, when $N_f > 1$ and $N_t > 1$, the aperiodic throughput $\mathcal{S}_{ap}$ is still given by (2.2).

This outcome is interesting, as (2.2) evidences that the aperiodic traffic throughput heavily depends on the size of aperiodic packets: for a given N_f value, the larger the size, the worse. Moreover, it is even more illuminating to observe that $\mathcal{S}_{ap}$ does not depend on N_t, i.e., on how stringent (or loose) the PDB of aperiodic packets is. Pairing the last result with the remark that aperiodic TBs will not collide with ongoing periodic traffic, as outlined at the end of previous subsection, allows to infer that the overall throughput of aperiodic and periodic traffic is insensitive to the delay requirement on the delivery of aperiodic packets. Ultimately, the PRR is expected to be independent of such PDB and indeed, the numerical results presented in this Section corroborate this insight.

If N_f is not a multiple integer of R, then the number of available channels is $\lfloor \frac{N_f}{R} \rfloor$, which forces some radio resources in a subframe to be unused. Their number amounts to $U = N_f - \lfloor \frac{N_f}{R} \rfloor \times R$. For a generic R, it follows that (2.2) modifies in

$$\mathcal{S}'_{ap} = \frac{N_f - U}{N_f} \cdot \left(G \cdot e^{-\frac{G}{\left\lfloor \frac{N_f}{R} \right\rfloor}}\right) < \mathcal{S}_{ap} \,. \tag{2.4}$$

The two previous expressions reveal that the choice of the R value plays a non-negligible role in the throughput evaluation. The emerging guideline is to adopt a proper combination of packet size, MCS, to minimize U and therefore confine unused radio resources.

2.3.2 Simulation Environment

The results presented hereafter have been obtained using the MoReV2X simulator presented in Chapter 1. In accordance to the highway scenario defined in [5], the examined setting consists of a 4 km long highway trunk, where six 4 meter wide lanes are considered, three per each traveling direction. Vehicular mobility is modeled using SUMO [20] and vehicles travel on the lanes following the modified Krauss model [20]; their speed is 140 km/h. The highway trunk has been populated considering three vehicular density values, 60, 120 and 260 vehicles/km, giving rise to three distinct scenarios that from now on will be termed low, medium and high density, respectively.

Vehicles generate application-layer packets in accordance to the periodic and aperiodic traffic pattern defined by 3GPP in [5] and reported in Subsection 1.5.1; each packet is encapsulated within a single TB. Vehicles acting as periodic traffic sources generate 190 byte-long TBs every $T = 100$ ms. The PDB of periodic traffic coincides with the generation period T. Unless otherwise stated, vehicles acting as aperiodic traffic sources set $c = \bar{r} = 50$ ms. The PDB of aperiodic packets coincides with c. This implies that every vehicle generates packets at an average rate of 10 packets/s. Two alternative choices for the aperiodic packet size are considered: large or small packets, 1000 or 190 byte long, in order to better highlight the influence of this parameter on system performance.

Vehicle radios operate at 5.9 GHz, on a 10 MHz channel that is partitioned in 4 subchannels of 12 RBs each. Both the 190 and 1000 bytes TBs are transmitted assuming QPSK modulation with a 0.7 code rate; as a result, they occupy 1 and 4 subchannels, respectively. Vehicles broadcast packets with a transmission power equal to 23 dBm and the receiver sensitivity is set to -90.4 dBm. As regards the algorithm ruling the radio resource allocation, it is SPS for periodic traffic, with the keep probability P set to 0. Aperiodic traffic is served as indicated in Subsection 2.3.1.1. Table 2.1 summarizes the most relevant numerical choices made in the current study.

Table 2.1: Simulation parameters.

Parameter	Values
Traffic density	60, 120, 260 vehicles/km
Highway length	4 km
Number of lanes	6 (3 per driving direction)
Vehicles' speed	140 km/h
Aperiodic traffic percentage (Θ)	0, 10, 50, 90, and 100 %
Packet size (X)	190 bytes, 1000 bytes
Occupied subchannels	1 (190 bytes), 4 (1000 bytes)
Subchannel size	12 RBs
Channel bandwidth	10 MHz
Transmission power	23 dBm
Receiver sensitivity	-90.4 dBm
Modulation and coding scheme	QPSK-0.7

2.3.3 Numerical Results

The evaluation of the performance metrics introduced in Subsection 1.5.3 is presented next; for all results, an adequate number of simulations has been executed, in order to determine sufficiently tight 95% confidence intervals.

Figs. 2.1(a)-(c) report the PRR as a function of the Tx-Rx distance D. The size of the aperiodic packets is $X = 1000$ bytes, different percentages of aperiodic traffic, namely, $\Theta = 10\%, 50\%$ and 90%, are considered, as well as the case when periodic traffic only is present ($\Theta = 0\%$). The worsening in performance for increasing Θ values is manifest. In the low density scenario, when half of the vehicles generate asynchronous traffic, the PRR drops below 0.9 for distances greater than 90 m. Moreover, Fig. 2.1(b) and Fig. 2.1(c) allow to appreciate how markedly the PRR worsens in increasingly crowded settings. To better explain why the aperiodic traffic is so detrimental to the performance of the SPS strategy, Table 2.2 reports the values of the CBR for the same choice of the parameters considered in Figs. 2.1(a)-(c). In this study, the RSSI threshold used to determine the CBR is 0.5 dB larger than the receiver sensitivity level; the CBR is averaged over the central portion of the simulation time and over all vehicles.

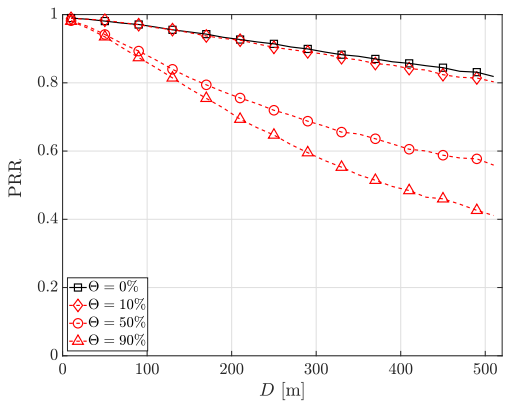

(a) Low density.

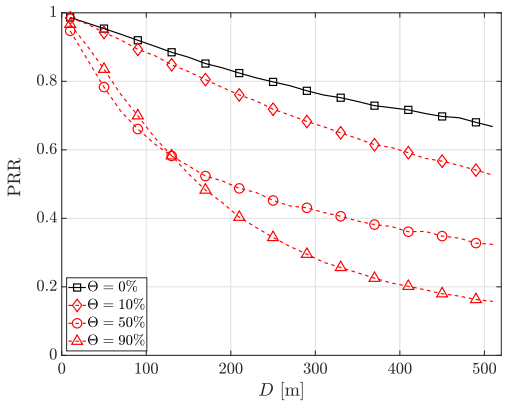

(b) Medium density.

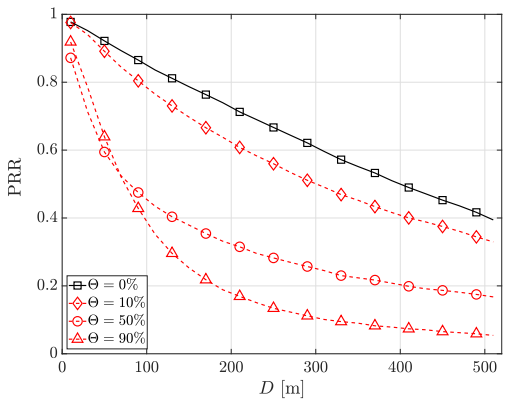

(c) High density.

Figure 2.1: PRR as a function of the Tx-Rx distance D, $X = 1000$ bytes.

The first row of Table 2.2 refers to the low density scenario and reveals that the CBR monotonically increases from 0.44 to 0.72, when the percentage of aperiodic traffic raises from $\Theta = 0\%$ to $\Theta = 90\%$. Moving to the second and third row of Table 2.2, i.e., considering the more crowded medium and high density settings, it is observed that the CBR takes on notably high values. Overall, the combined effect of increased Θ values and increased vehicular densities leads to remarkable loads on the radio channel, that inflate the CBR and are responsible for the PRR degradation evidenced in Figs. 2.1(a)-(c). Table 2.3 also helps in grasping what happens in the medium and high density settings portrayed in Figs. 2.1(b) and (c). These figures revealed that, for transmitter-receiver distances lower than 100 m, the PRR values for $\Theta = 50\%$ and $\Theta = 90\%$ are similar, with a slightly worse performance observed for $\Theta = 50\%$, in spite of a lighter load placed on the channel. Table 2.3 discloses that the average collision ratio exhibits its maximum exactly for $\Theta = 50\%$ and that this maximum is not so far from the value observed for $\Theta = 90\%$ (0.32 versus 0.28 for the medium density scenario, 0.46 versus 0.4 for the high density scenario). When comparing the two cases, it is observed that for $\Theta = 50\%$ the aperiodic traffic is lighter, but the large, aperiodic packets randomly compete to gain access over a smaller fraction of radio resources left unoccupied by periodic flows, whose transmissions are protected by the proposed strategy. Conversely, the overall traffic is heavier for $\Theta = 90\%$, yet more radio resources are available for aperiodic transmissions. Ultimately, this leads to a PRR deterioration that is comparable in the two cases.

Table 2.2: CBR values.

	$\Theta = 0\%$	$\Theta = 10\%$	$\Theta = 50\%$	$\Theta = 90\%$
Low density	0.44	0.49	0.55	0.72
Medium density	0.77	0.72	0.67	0.83
High density	0.92	0.93	0.75	0.88

Table 2.3: $\overline{C}_R$ values.

	$\Theta = 0\%$	$\Theta = 10\%$	$\Theta = 50\%$	$\Theta = 90\%$
Low density	0.11	0.14	0.19	0.22
Medium density	0.2	0.23	0.32	0.28
High density	0.32	0.35	0.46	0.4

Figs. 2.2(a)-(c) are the counterparts of Figs. 2.1(a)-(c), when the size of the periodic packets is reduced to $X = 190$ bytes, all other system choices being unmodified. As it had to be expected, the size plays a relevant role in the achievable performance. Small aperiodic packets occupy a limited amount of radio resources, making the coexistence with periodic transmissions almost unproblematic. In the low density scenario, the PRR curves obtained for a small aperiodic packet size and different percentages of aperiodic traffic reveal a modest dependence on Θ. Furthermore, the curves always lie in the region of the (D, PRR) plane where the PRR takes on values higher than 0.75. The curves obtained for

different Θ values begin to differentiate at medium traffic density. The difference in the relative position of the curves referring to $\Theta = 0\%$, $\Theta = 10\%$ and $\Theta = 50\%$ in Fig. 2.2(a) (low density) with respect to Fig. 2.2(b) (medium density) can be understood observing that: (i) higher Θ values, i.e., higher percentages of aperiodic traffic, more heavily penalize the PRR of aperiodic flows; (ii) in the medium density scenario, the number of aperiodic flows increases. Overall, this explains why, in Fig. 2.2(b), the overall PRR more markedly decreases moving from $\Theta = 0\%$ to $\Theta = 10\%$ and $\Theta = 50\%$ than it does in Fig. 2.2(a). When the extreme, high density landscape is examined, the curves are very far apart from one another and the PRR sharply degrades to unbearable values.

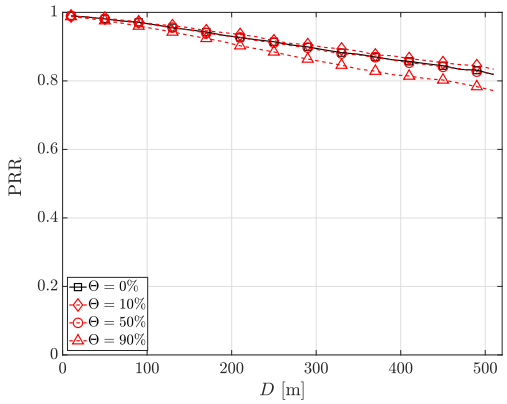

(a) Low density.

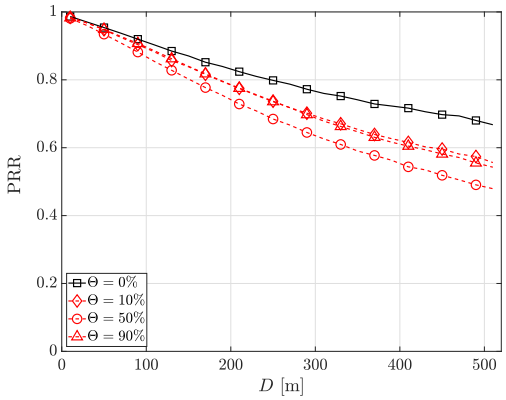

(b) Medium density.

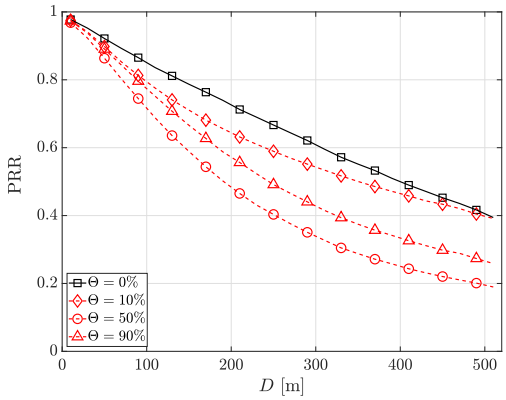

(c) High density.

Figure 2.2: PRR as a function of D, $X = 190$ bytes.

Unfortunately, aperiodic packets bred by general purpose eV2X applications are expected to exhibit a large size [5]. Unless otherwise stated, in what follows the main focus will therefore be on the attainable performance when relatively large-sized aperiodic packets ($X = 1000$ bytes) are considered. Moreover, the low density setting only will be examined.

For the same system parameters considered so far, Fig. 2.3 separately reports the PRR experienced by vehicles generating periodic and aperiodic traffic as a function of the D

when $\Theta = 50\%$ and 90%. It is observed that: (i) the PRR of periodic traffic is almost unaffected by Θ; (ii) the PRR of aperiodic traffic is significantly worse than the PRR of periodic traffic and it decreases for increasing Θ. The reason is that an aperiodic vehicle performs a random resource selection contending with periodic and aperiodic traffic, but it respects periodic reservations. When Θ increases from 50% to 90%, the effect is emphasized, as aperiodic traffic requires four times the resources that periodic traffic asks for. When $\Theta = 50\%$, the aperiodic PRR is down to 0.8 at 130 m from the transmitting vehicle, revealing the low level of reliability guaranteed to the delivery of aperiodic packets.

Fig. 2.4 delves into the impact that different arrival rates of aperiodic traffic have on the PRR of vehicles generating periodic and aperiodic packets. The results reported in this figure refer to $c = \bar{r} = 10$, 50 and 100 ms. The PDB of aperiodic packets is equal to 10, 50 and 100 ms, respectively. Note that the 10 and 50 ms choices adhere to the preliminary indications for eV2X applications detailed in [5] and [55]. The curves reveal that the PRR of both periodic and aperiodic traffic collapses when the arrival rate is significantly high.

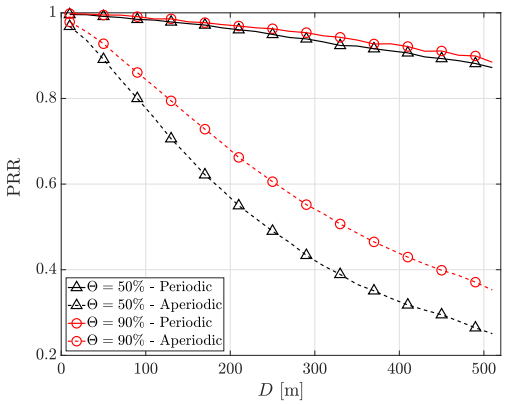

Figure 2.3: Aperiodic and periodic PRR as a function of D.

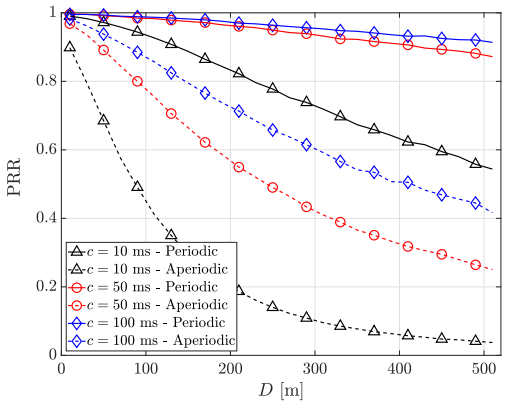

Figure 2.4: PRR for different aperiodic traffic models.

Next, in order to offer a further insight into the behavior of the proposed strategy, such solution is compared against two schemes that accommodate aperiodic traffic in accordance to SPS. We term such alternatives *submissive* and *aggressive* scheme and report their description below. For the submissive scheme, aperiodic traffic reserves radio resources strictly in accordance to SPS, hence periodically, employing a reservation interval RRI that matches the average inter-arrival time of aperiodic packets, i.e., RRI $= c + \bar{r}$. When an aperiodic packet is ready for transmission, the originating vehicle first checks if the next reserved radio resource becomes available within the PDB of aperiodic traffic. If it does, the TB of the aperiodic packet plus its SCI are allocated within such resource. Otherwise, the aperiodic packet is dropped, as its transmission within the next reserved resource would violate the delay requirement. The aggressive scheme differs from the submissive scheme only in the event that next reserved resource violates the PDB. In this circumstance, the aperiodic packet is not dropped; rather, the corresponding TB plus its SCI is transmitted over a free resource picked at random in a selection window W whose duration coincides

with the PDB. Both solutions represent interesting terms of comparison: the submissive scheme does not prevaricate over periodic flows, as aperiodic traffic adheres to periodic reservations whenever possible and is dropped otherwise; on the other hand, the aggressive strategy always serves aperiodic traffic, at the expense of periodic traffic.

Hence, Fig. 2.5(a) compares the PRR of periodic and aperiodic traffic obtained when aperiodic packets are served according to our one-shot proposal with the PRR performance achieved by the aggressive and the submissive strategy. Here, $\Theta = 50\%$, $c = \bar{r} = 50$ ms and $X = 1000$ bytes. The submissive scheme penalizes aperiodic traffic to a great extent, while the aggressive scheme heavily strikes on periodic traffic performance. On the other hand, our solution serves periodic packets in an excellent manner at the expense of aperiodic traffic, resulting in an intermediate approach. For the sake of completeness, Fig. 2.5(b) broadens the comparison, considering $X = 190$ bytes. For this size of aperiodic packets, the submissive strategy is inadequate too, whereas our solution overall performs better than the aggressive scheme. The PRR that our proposal guarantees periodic packets is as high as possible, while the PRR of aperiodic packets is only slightly lower than the PRR aperiodic traffic experiences if the aggressive strategy is adopted. To complete the comparison, Fig. 2.6 confronts our proposal, the submissive and aggressive strategies, when $\Theta = 100\%$ and $X = 1000$ bytes. As somewhat had to be expected, in this limiting case, the submissive strategy is largely unsuccessful, whereas the performance of the aggressive scheme gets close to our proposed solution.

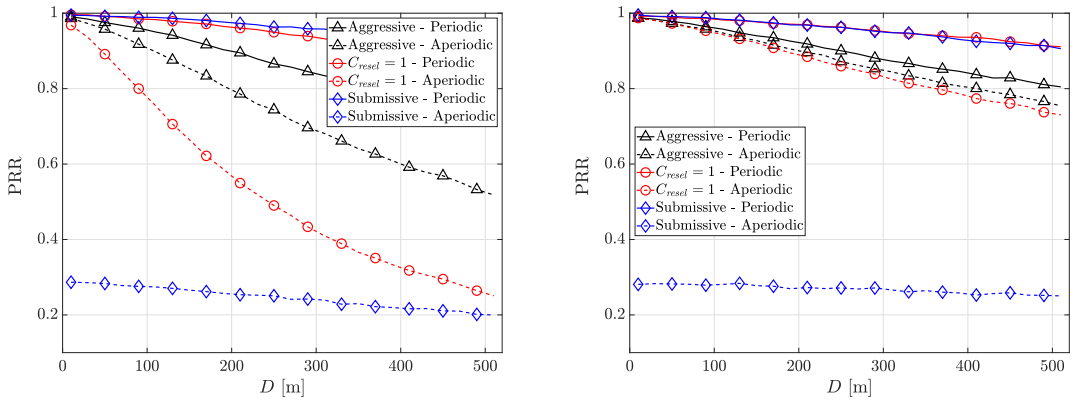

(a) $X = 1000$ bytes. (b) $X = 190$ bytes.

Figure 2.5: Comparison among alternative schemes, $\Theta = 50\%$.

Next, Fig. 2.7 reports the overall PRR as a function of the transmitter-receiver distance D for different values of the aperiodic traffic PDB, PDB $= 10, 20, 30$ and 50 ms. In this figure, we also assume $\Theta = 50\%$, $c = \bar{r} = 50$ ms and $X = 1000$ bytes. Fig. 2.7 shows that the PDB choice has no impact on the PRR curves, confirming the *a priori* indication provided by the throughput analysis in Subsection 2.3.1.2. The effect of more stringent delay requirements for aperiodic traffic is to shrink its selection window W; this might lead to the erroneous intuition that, for a given traffic density, a PDB decrease should

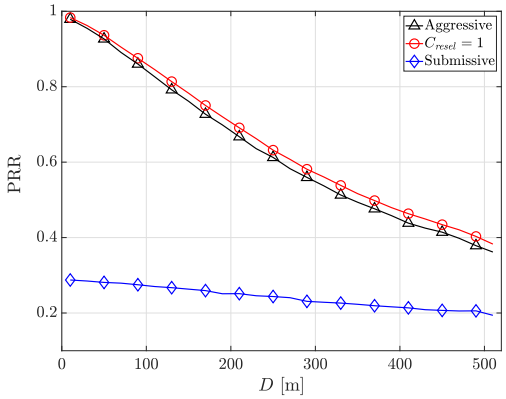

Figure 2.6: Comparison among alternative schemes, $\Theta = 100\%$ and $X = 1000$ bytes.

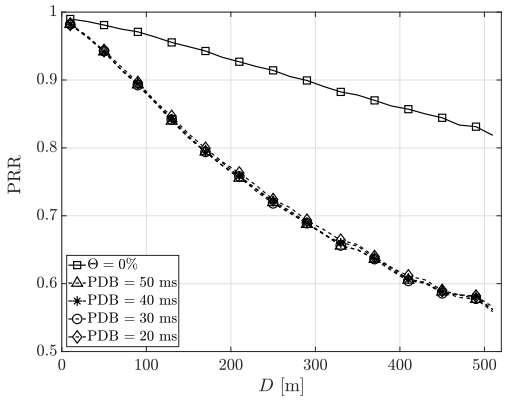

Figure 2.7: PRR for different values of the PDB.

reflect in a lower PRR. However, Fig. 2.7 demonstrates it is not true.

Fig. 2.8 considers the limiting case when only aperiodic traffic is present ($\Theta = 100\%$), and reports the PRR for three values of packet size, $X = 1000$, 720 and 190 bytes, when the low and medium density scenarios are examined. On one hand, increasing the packet size from $X = 190$ to 720 bytes has a remarkable negative effect on the PRR. The effect is amplified the more crowded the vehicular setting is. On the other hand, the PRR is not affected at all by the packet size increase from $X = 720$ to 1000 bytes. As a matter of fact, the PRR depends on the maximum number of packets that can be successfully allocated without collisions in every subframe, which is expressed by the ratio between the number of available subchannels and the number of required subchannels for packet. Such ratio specializes to $\lfloor \frac{4}{1} \rfloor$ when $X = 190$ bytes, and to the same unitary value for $X = 720$ bytes and $X = 1000$ bytes, $\lfloor \frac{4}{3} \rfloor$ and $\lfloor \frac{4}{4} \rfloor$, respectively. More generally, if larger packet sizes not fitting in one subchannel were to be considered, it is our belief that fragmentation should be avoided. Rather, a more spectrally efficient modulation scheme might be used; depending on the application type and on its requirements, a low code rate might also be employed.

In Fig. 2.9(a), the average PIR ($\overline{\text{PIR}}$) as a function of D is separately reported for

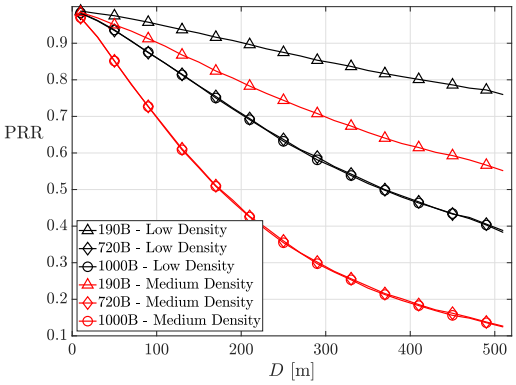

Figure 2.8: PRR when $\Theta = 100\%$.

aperiodic and periodic traffic in the low density scenario, for $\Theta = 10\%$, 50% and 90% and $X = 1000$ bytes. The same remarks applied to Fig. 2.3 hold, as the average PIR strongly depends on the PRR: when the former lowers, the latter inevitably climbs up. In Fig. 2.9(b), the same setting is considered, except for the packet size, which is $X = 190$ bytes; as already observed with reference to Fig. 2.2(a), the injection of aperiodic traffic is not problematic, as long as the size of its packets is small.

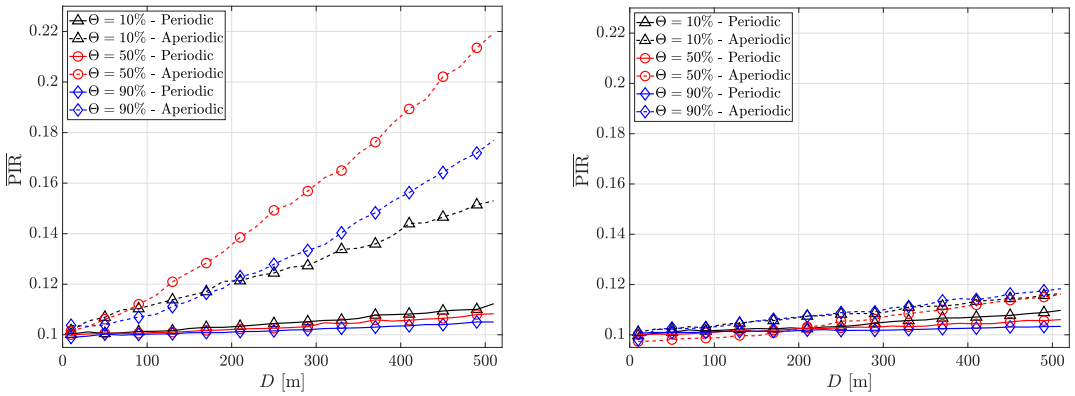

(a) $X = 1000$ bytes.

(b) $X = 190$ bytes.

Figure 2.9: Average PIR of aperiodic and periodic traffic as a function of D.

Fig. 2.10 concludes the investigation, reporting the PIR CDF for periodic and aperiodic traffic flows, for $\Theta = 10\%$, 50% and 90% in the low density scenario, when $X = 1000$ bytes. In this figure, $D_{max} = 520$ m; the CDF is therefore evaluated from the PIR values collected for all transmitter-receiver distances falling in the $(0, 520]$ range. The PIR CDF of periodic traffic exhibits a step behavior, that reflects the periodicity $T = 100$ ms of resource assignment; note that these curves are nearly independent on the aperiodic traffic percentage. Instead, the PIR CDF of aperiodic traffic smoothly varies, reflecting that aperiodic traffic is characterized by a random inter-packet arrival time and is randomly served. Unfortunately, this figure demonstrates that the PIR of aperiodic traffic cannot be

successfully confined: the probability of observing PIR values lower than 100 ms is only 0.4, unbearable for all future eV2X use cases.

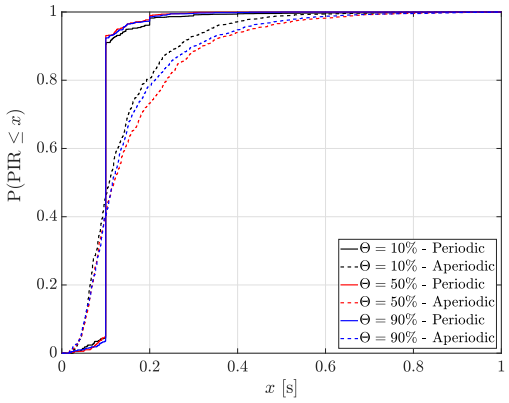

Figure 2.10: Aperiodic and periodic PIR CDF, $X = 1000$ bytes.

2.4 NR-V2X Mode 2: MAC Analysis

In NR-V2X Mode 2, vehicles employ a distributed access strategy based on the resource reselection and re-evaluation mechanisms to select, and in the SPS strategy also reserve, their transmission resources. This Section analyzes the functioning of the NR-V2X Mode 2 MAC sublayer. The goal of this Section is two-fold: (i) it sheds light on the impact that different traffic types and PDB requirements have on the operation of the SPS and DS strategies; (ii) it provides a qualitative analysis of the re-evaluation mechanism, discussing the impact of the most relevant NR-V2X Mode 2 parameters on the re-evaluation check, the re-evaluation detection, and the resource replacement phase introduced in Subsection 1.4.2.2.

We begin by observing that, during the resource reselection process, collisions may occur when the selection windows of two or more vehicles overlap and the vehicles end up selecting the same resources. As selected resources cannot be announced before being employed, this type of collision is unavoidable, regardless of the scheduling strategy adopted, SPS or DS. Such a situation is clarified in Fig. 2.11(a), where two vehicles, v_1 and v_2, perform a resource reselection and their selection windows, W_1 and W_2, partially overlap. If the two vehicles select the same resources in the overlapping region, a collision will occur. Note that, in the SPS case, the collision will persist for a number of consecutive transmissions, given the involved vehicles generate packets with the same periodicity and employ the same RRI. As the resource reselection process may generate collisions, the larger the number of resource reselections the higher the collision probability, as Fig. 2.11(b) suggests. In this figure, two additional vehicles, v_3 and v_4, perform a resource reselection and their selection windows, W_3 and W_4, overlap with W_1 and W_2. It follows that the number of potentially interfering vehicles grows from 2 to 4, therefore increasing the collision probability.

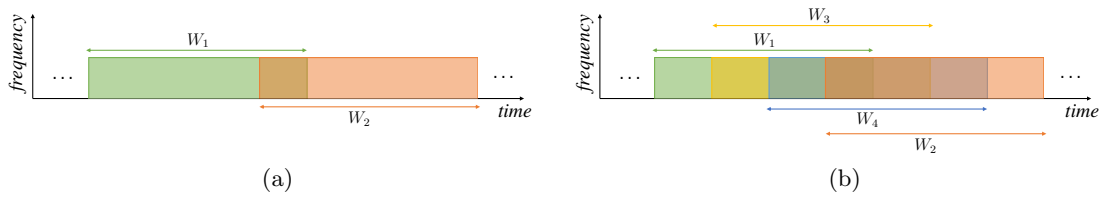

Figure 2.11: Collision probability associated with resource reselections.

2.4.1 Impact of Different Traffic Types

When fixed size periodic packets are considered, vehicles perform a resource reselection only when the reselection counter expires, i.e., $C_{resel} = 0$. We term such an event *counter reselection*. When the SPS scheme is adopted, the number of counter reselections depends on the value of the keep probability P and on the average value of the C_{resel} counter. On the other hand, in the DS scheme, a counter reselection is triggered every time a new packet is generated, regardless of the traffic type.

When the packet size or the inter-arrival time between packets are not constant, the scheduling capability that characterizes the SPS scheme is deteriorated by the occurrence of additional resource reselections, called size and latency reselections, that are illustrated next:

- A *size reselection* occurs when the generated packet does not fit the current reservation size and forces the reselection of a larger number of subchannels. A size reselection is illustrated in Fig. 2.12, where the ego-vehicle generates a 200 bytes long packet at slot s_{g_1}, transmits it in the selected resources at slot s_{r_1} and broadcasts the associated SCI to reserve the same subchannel after RRI slots, at slot s_{r_2}. However, the next packet generated by the ego-vehicle, at slot s_{g_2}, is 400 bytes long and it does not fit the reservation at slot s_{r_2}. Therefore, the ego-vehicle is forced to perform a resource reselection, select a larger number of subchannels at slot s_{r_3}, and leave the reservation at slot s_{r_2} idle.

- A *latency reselection* is triggered when the reserved resources do not satisfy the latency requirements of a TB, i.e., its PDB. Latency reselections only occur if the adopted RRI value is larger than the PDB of the generated packets, RRI > PDB. In Fig. 2.12, after transmitting the TB at slot s_{r_3}, the ego-vehicle reserves the same resources after RRI slots, at slot s_{r_4}. Next, the ego-vehicle generates a 200 bytes long packets at slot s_{g_3} whose latency limit is identified by the s_{lim} slot. Since $s_{lim} < s_{r_4}$, the reservation at slot s_{r_4} cannot satisfy the PDB of the generated packet, thus forcing the reselection of new resources able to accommodate the packet transmission within its latency limit. As a result, the ego-vehicle selects new resources at slot s_{r_5}, $s_{r_5} < s_{lim}$, and leaves the reservation at slot s_{r_4} idle.

Size and latency reselections cause vehicles to perform additional resource reselections before the depletion of the reselection counter. Therefore, they increase the total number

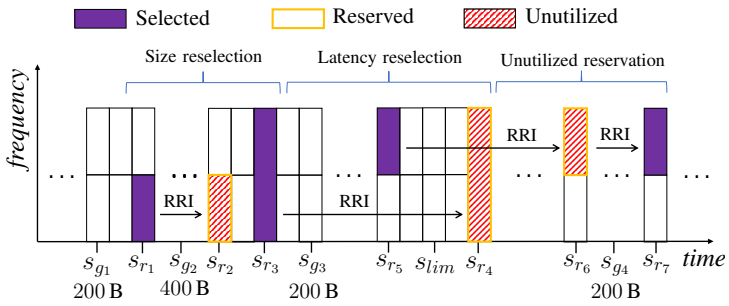

Figure 2.12: Impact of different traffic types on NR-V2X Mode 2.

of resource reselections and have a negative impact on the SPS performance, being responsible for an increase in the collision probability. Neither size nor latency reselections occur in the DS strategy, as the ego-vehicle performs a counter reselection for each newly generated TB. Table 2.4 summarizes the occurrence conditions of the different reselection types when both the SPS and the DS scheme are considered.

When investigating the SPS behavior, a further detrimental phenomenon to consider is that of *unutilized reservations*, which occur when previously reserved resources are not employed by the ego-vehicle for transmitting any TB. Unutilized reservations always occur after a size or latency reselection: as exemplified in Fig. 2.12, the resources initially reserved at slots s_{r_2} and s_{r_4} do not accommodate any packet transmission. Fig. 2.12 additionally reveals that reserved resources can be left idle even when no size or latency reselections take place. In this case, unutilized reservations occur when the employed RRI is lower than the PDB of the generated packets, RRI $<$ PDB. After employing the selected resources at slot s_{r_5} for transmitting the TB and the associated SCI, the ego-vehicle reserves the same resources after RRI slots, at slot s_{r_6}. However, the ego-vehicle does not generate any packet until slot s_{g_4}, with $s_{g_4} > s_{r_6}$. Therefore, the ego-vehicle leaves the reservation at slot s_{r_6} idle and transmits the generated packet at slot s_{r_7}. Note that the unutilized reservation at slot s_{r_6} does not allow the ego-vehicle to broadcast the corresponding SCI and announce the reservation of the same resources at slot s_{r_7}, thus increasing the collision probability. In this case, unutilized reservations are responsible for the missed advertisement of the next reserved resources and for the transmission of the TB on selected, rather than reserved resources. Moreover, observe that unutilized reservations cannot be included by neighboring vehicles in their list of candidate resources, which have a reduced pool of candidates wherefrom they select the resources for their transmission. Such a waste of system capacity translates into a higher collision probability.

Table 2.4: Occurrence conditions of different reselection types.

	Counter reselection	Size reselection	Latency reselection
SPS	$C_{resel} = 0$, depends on P	TB size larger than reservation size	$RRI > PDB$
DS	Every TB	Never	Never

2.4.2 Impact of the PDB

As reported by 3GPP in [55], future eV2X applications will be characterized by a wide spectrum of latency requirements. For instance, the maximum tolerated end-to-end latency ranges from 3 ms for emergency trajectory alignment applications to 500 ms for lower priority platooning-related reporting. In this Subsection, we concentrate on the impact that different latency requirements have on the operation of the SPS and DS strategies. In NR-V2X Mode 2, the latency requirement of the packet is mapped into the PDB, which identifies the upper bound to the latency that the packet can experience. Also recall that the width of the selection window is limited by the T_2 parameter, $T_2 \leq$ PDB, as it was evidenced in Section 1.4. Unless otherwise stated, we assume that all vehicles employ the same PDB and that $T_2 =$ PDB.

1) *Dynamic Scheduling:* with DS, vehicles trigger a counter reselection for every generated message and do not reserve any resources. As a result, vehicles performing a resource reselection blindly assume that all the subchannels included in their selection window are available, and randomly select a sufficient number of subchannels able to accommodate the message transmission. If we model the selection window as an $N_f \times N_t$ grid, where N_f indicates the number of subchannels on the frequency axis and N_t the number of slots in W, we can conclude that the DS scheme behaves as a multichannel slotted Aloha access strategy. When a new message is generated, it is transmitted on one (or more) randomly selected subchannels out of N_f, during a single time slot.

   Note that the same intuition was reported in Subsection 2.3.1.2, where we analyzed the performance of a one-shot solution to disseminate aperiodic packets in LTE-V2X Mode 4. Based on the Aloha-like assumption, this Section extends the analysis of Subsection 2.3.1.2 and analytically derives the PDB impact on the DS performance.

   Let us assume that vehicles generate traffic with an overall average rate λ messages/s, no matter what traffic type is considered. Accordingly, the average number of potentially colliding messages that are generated during a slot of duration t_s is approximately given by

$$G \simeq \lambda \cdot t_s. \tag{2.5}$$

   Due to the random selection of resources that characterizes the DS scheme, the transmissions of all the messages generated during a generic slot are uniformly distributed within the selection window W. Therefore, the average number of potentially colliding messages that are transmitted on a single slot is

$$G_i \simeq \frac{G}{N_t}. \tag{2.6}$$

   However, this term adds to $N_t - 1$ analogous contributions originated during the previous $N_t - 1$ slots. As a result, the average number of potentially colliding messages

that are transmitted on every selection window slot is

$$G_i + \sum_{i=1}^{N_t-1} G_i = \sum_{i=1}^{N_t} G_i \simeq N_t \cdot \frac{G}{N_t} = G. \tag{2.7}$$

This reasoning reveals that the number of potentially colliding messages transmitted on every slot does not depend on the width N_t of the selection window. In other words, the collision probability that characterizes the DS scheme is independent of the PDB, as it will be numerically demonstrated in the following sections.

2) *Semi-Persistent Scheduling:* with SPS, the resource reselection process is not random as in the DS case, but strives to avoid the selection of subchannels already occupied by other vehicles. When the SPS scheme is examined, more stringent latency requirements (i.e., a smaller PDB) and the associated reduction of the selection window width have a three-fold effect. First, a smaller PDB lowers the number of available subchannels in W. As a result, vehicles that trigger a resource reselection during the same time slot have a larger probability of selecting the same subchannels and generate a collision. Second, a shorter selection window reduces the number of potentially interfering vehicles (i.e., the number of vehicles whose selection windows overlap), accordingly reducing the collision probability. As it will be demonstrated by simulation results, these two effects cancel out.

Third, when the RRI does not match the inter-arrival time between messages, a smaller PDB reduces the probability that reserved subchannels satisfy the latency requirements of the generated messages. Hence, more latency reselections may occur and the collision probability increases.

2.4.3 Analysis of the Re-Evaluation Mechanism

The re-evaluation mechanism is an important novelty introduced in NR-V2X Mode 2 to guarantee a more flexible and effective scheduling of transmissions. This Subsection analyzes the operation of the re-evaluation mechanism and discusses the impact that the most relevant NR-V2X Mode 2 parameters have on the effectiveness of the re-evaluation check, the re-evaluation detection, and the resource replacement phase.

2.4.3.1 Re-Evaluation Check

Vehicles use the re-evaluation check to assess whether selected resources are still available or not right before transmitting the TB. The objective is to detect and avoid potential collisions. 3GPP standards establish that re-evaluation checks are only possible on selected (and not reserved) resources. Recall from Subsection 1.4.2.2 that we call selected resources the subchannel(s) that have not been reserved by vehicles through SCI broadcasting. Instead, we term reserved resources the subchannel(s) whose reservations have been announced by the vehicles' SCI before being employed for the TB transmission.

Accordingly, re-evaluation checks are performed before the transmission of all TBs when using the DS scheme, since this strategy selects new resources for each TB. When the SPS scheme is considered, re-evaluation checks affect a smaller number of TBs since SPS only selects new resources as a result of a counter reselection, by design. Once new resources are selected, the remaining TBs are transmitted on reserved resources. If we assume, for example, $P = 0$ and RRI ≥ 100 ms, only 1 TB out of 10 triggers a counter reselection (the average reselection counter value is 10 in this case); hence, only 10% of the generated TBs are transmitted on selected resources that are eligible for a re-evaluation check. However, we should note that latency reselections, size reselections, and unutilized reservations (see Subsection 2.4.1) increase the fraction of TBs that are transmitted on selected resources in SPS, and thus increases the number of re-evaluation checks. Regardless of the scheduling scheme, the fraction of TBs that triggers a re-evaluation check is also affected by the value of T_2, i.e., by the width of the selection window (see Fig. 1.1). Let us assume that a vehicle V_A generates a new TB at slot s_{g_1} and performs a resource reselection. The selection window is defined by the range of slots $[s_{g_1} + T_1, s_{g_1} + T_2]$ where $T_1 \leq T_{proc,1}$. In principle, any selected resource included within the selection window shall be eligible for a re-evaluation check. However, a re-evaluation check can be performed only if the vehicle has sufficient processing capabilities to run the entire re-evaluation mechanism before transmitting the TB. According to the standard [54], a re-evaluation check can only be performed if the selected resource is included in the $(s_{g_1} + T_3, s_{g_1} + T_2]$ interval, where T_3 is strictly equal to $T_{proc,1}$. Therefore, the candidate resources included from slot $s_{g_1} + T_1$ to slot $s_{g_1} + T_3$ are not eligible for a re-evaluation check. Depending on T_2, the number of resources included in $[s_{g_1} + T_1, s_{g_1} + T_3]$ can be a significant fraction of the total number of resources within the selection window. For example, let us assume that $T_2 = \text{PDB} = \text{RRI}$, and that $\mu = 0$, $T_1 = 1$ slot and $T_3 = 5$ slots. In this case, the percentage of selection window resources that are not eligible for a re-evaluation check is equal to $\{5, 25, 50\}$ % when RRI $= \{100, 20, 10\}$ ms.

2.4.3.2 Re-Evaluation Detection

A re-evaluation detection is triggered after a re-evaluation check when the initially selected resources are no longer available, i.e., are no longer included in list L_1. Typically, a re-evaluation detection occurs when the initially selected resources have also been reserved by a neighboring vehicle, and a potential collision is detected. This Subsection sheds light on the circumstances under which a potential collision does and does not trigger a re-evaluation detection. To do so, we separately analyze the re-evaluation detection phase when each TB is transmitted once ($N = 1$) and when it is transmitted twice ($N = 2$, with one blind retransmission). This Subsection ends with a discussion about the effectiveness of the re-evaluation detection phase.

1) *Single transmission per TB ($N = 1$)*: we first consider the case where a collision occurs on selected resources. This type of collision cannot be detected by a re-

evaluation detection since vehicles do not announce their selection before transmitting on selected resources. This is illustrated in Fig. 2.13 where V_A and V_B select new resources to transmit their TBs generated at slots s_{g_1} and s_{g_2}, respectively. If their selection windows (SW_A and SW_B) overlap, the two vehicles may select the same resources at slot s_{r_1}. V_A performs a re-evaluation check at slot $s_{r_1} - T_3$, but step 1 does not exclude the resources at slot s_{r_1} since V_B has not yet announced its reservation. This is the case because also V_B has performed a reselection after generating the TB at slot s_{g_2} and its transmission at slot s_{r_1} occurs on selected resources. The re-evaluation mechanism is not capable to detect and avoid the collision at s_{r_1}. The same situation occurs when V_B executes its re-evaluation check. If V_A and V_B use the SPS strategy, they will reserve the same resources for transmitting their next TB, at slot s_{r_2}, and they will persistently collide until a (counter, latency, or size) reselection occurs if they employ the same RRI. Such persistent collision is illustrated in Fig. 2.13. The persistent collision cannot be avoided by the re-evaluation mechanism because it is only executed over selected and not reserved resources. If vehicles use the DS strategy, they select new resources for every TB when $N = 1$. Therefore, collisions occur only on selected resources and the re-evaluation mechanism is not able to detect and avoid any collision. As demonstrated in the remainder of this Subsection, only collisions between selected and reserved resources can be identified by the re-evaluation detection. Since vehicles using DS do not reserve resources, the following discussions apply only to the SPS scheme.

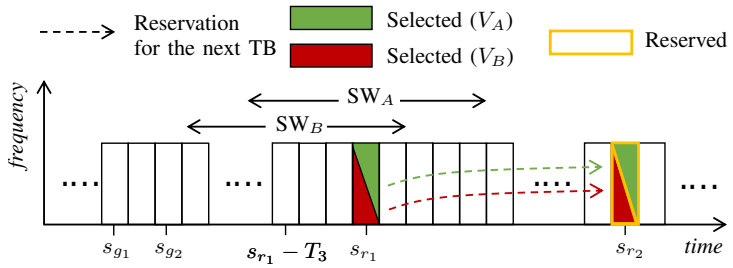

Figure 2.13: Persistent collisions not detected by re-evaluation when using SPS.

Depending on the RRIs used by vehicles and the type of generated traffic, we can identify four different cases in which a re-evaluation detects a potential collision. The first case is illustrated in Fig. 2.14(a) and corresponds to the scenario where vehicles V_A and V_B transmit periodic TBs of fixed size and use the same RRI. V_B selects new resources to transmit the TB generated at slot s_{g_1} and selects the resources reserved by V_A at slot s_{r_1}. V_B does not exclude the resources reserved by V_A from its selection window (SW_B) during the resource reselection process because V_A announced its reservation in the range of slots $[s_{g_1} - T_{proc,0}, s_{g_1}]$, i.e., just after the end of V_B's sensing window. However, V_B can avoid the collision thanks to the re-evaluation check executed at slot $s_{r_1} - T_3$. At this time, V_B defines a new sensing window that includes the reservation announced by V_A. Then, the re-evaluation

detection identifies the potential collision and V_B excludes the resources reserved by V_A from its new selection window. It is worth pointing out that the probability of this type of re-evaluation detection is very low since the width of the $[s_{g_1} - T_{proc,0}, s_{g_1}]$ interval is equal to 2 slots for an SCS of 15 or 30 kHz, and equal to 3 slots for a 60 kHz SCS.

The second case is illustrated in Fig. 2.14(b) and corresponds to the scenario where V_A and V_B are involved in a persistent collision and an unutilized reservation occurs. In this figure, V_A and V_B initially collide at slot s_{r_1} without triggering a re-evaluation detection, like in Fig. 2.13, and start to persistently collide since they use the same RRI and re-evaluation cannot be applied to reserved resources. Let us suppose that V_B does not generate a TB and leaves the reserved resources at slot $s_{r_2} = s_{r_1} + \text{RRI}$ unutilized. In this case, V_B will not be able to announce the reserved resources at slot $s_{r_3} = s_{r_2} + \text{RRI}$, and will transmit its next TB in selected, rather than reserved, resources. As a result, V_B runs a re-evaluation check at slot $s_{r_3} - T_3$, right before transmitting its TB. During the re-evaluation check, the new sensing window of V_B includes the reservation announced by V_A at slot s_{r_2}, and V_B excludes the resources reserved by V_A at slot s_{r_3} from its new selection window (SW'_B). This triggers a re-evaluation detection that resolves the persistent collision between the two vehicles.

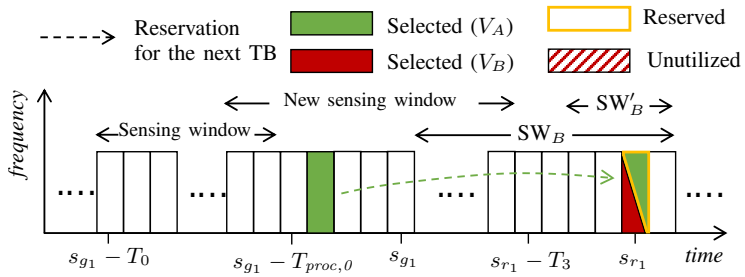

(a) Periodic TBs of fixed size.

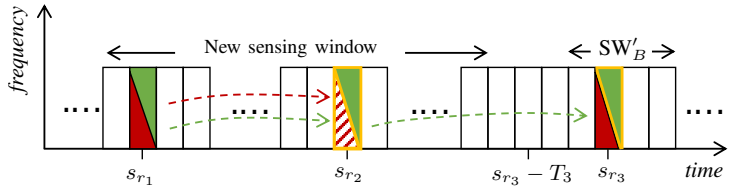

(b) Unutilized reservations.

Figure 2.14: Re-evaluation detection when vehicles use the same RRI.

The third and fourth cases where re-evaluation detection successfully detects potential collisions occur when the two vehicles involved use different RRIs. Without loss of generality, we consider two different RRI values, RRI_1 and RRI_2, with $\text{RRI}_1 < \text{RRI}_2$. The third case occurs when vehicles transmit periodic TBs of fixed size, and the two following conditions are satisfied: (i) the resources selected and reserved by a vehicle using the smaller RRI_1 are included within the selection win-

dow of a vehicle configured with RRI_2, and (ii) the vehicle using RRI_2 selects the resources reserved by the vehicle using RRI_1. This situation is illustrated in Fig. 2.15(a) where V_A generates a new TB at slot s_{g_1}, transmits it on the selected resources at slot s_{r_1}, and reserves the same resources at slot $s_{r_2} = s_{r_1} + \text{RRI}_1$. When V_B generates its new TB at slot s_{g_2}, it cannot be aware of the reservation announced by V_A at slot s_{r_1} due to the overlap between their selection windows (SW_A and SW_B). Let us then suppose that V_B selects the same resources at slot s_{r_2} and generates a collision. Note that SW_B is wider than SW_A because V_B uses the largest RRI_2 value and $T_2 = \text{PDB} = \text{RRI}_2$. V_B can avoid the collision at slot s_{r_2} by executing a re-evaluation check at slot $s_{r_2} - T_3$. The new sensing window of V_B will now include the reservation announced by V_A at slot s_{r_1} since $s_{r_2} - T_3 > s_{r_1}$. Then, V_B excludes the resources at slot s_{r_2} from its new selection window (SW'_B) and the re-evaluation detection triggers the process to select new resources.

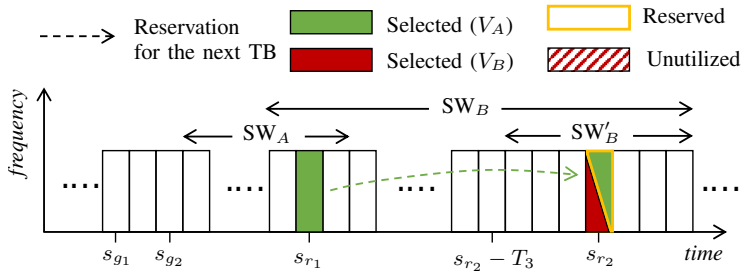

(a) Periodic TBs of fixed size.

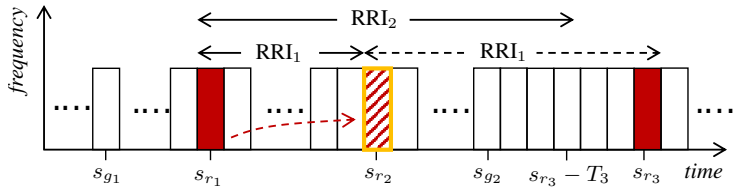

(b) Unutilized reservations.

Figure 2.15: Re-evaluation detection when vehicles use different RRIs.

The fourth case where re-evaluation detection successfully detects a collision occurs when the two RRI values are multiples of each other (e.g., $\text{RRI}_2 = 2 \cdot \text{RRI}_1$) and a vehicle leaves one of its reservations unutilized. This is illustrated in Fig. 2.15(b), where V_B generates a TB at slot s_{g_1}, selects the resources for its transmission at slot s_{r_1}, and periodically reserves them at slots s_{r_2} and s_{r_3} using the smallest RRI value, i.e., RRI_1. Let us suppose that V_B leaves the resources at slot s_{r_2} unutilized because it has no TB ready to be transmitted. As a result, V_B cannot reserve the resources at slot s_{r_3}, and it will run a re-evaluation check at slot $s_{r_3} - T_3$. During the re-execution of step 1, V_B will remove the resources at slot s_{r_3} from its selection window due to its half-duplex limitations, as it could not sense the reservations announced from neighboring users at slot s_{r_1}, therefore triggering a re-evaluation detection. We should recall from Section 1.4 that step 1 excludes from the selection

window all the slots in which V_A was previously transmitting, considering the entire list of allowed RRI values. Since $s_{r_3} = s_{r_1} + \text{RRI}_2$ and V_B was transmitting at slot s_{r_1}, it excludes slot s_{r_3} from its selection window.

2) *Two transmissions per TB ($N = 2$)*: without loss of generality, this Subsection analyzes the impact of retransmissions on the re-evaluation detection considering one blind retransmission per TB, i.e., $N = 2$. When $N > 1$, the 1st-stage SCI associated with the TB's initial transmission can reserve the resources used for the retransmission of the same TB if the distance between selected resources is smaller than 32 slots. In this case, the number of reservations announced by the SCI is indicated with $N_{SCI} = 2$. If the distance between selected resources is larger than 32 slots, the SCI is not able to announce reservations for the retransmission of the same TB and $N_{SCI} = 1$. In the $N_{SCI} = 1$ case, the initial transmission and the retransmission of the TB behave as two completely independent events, and no additional collision between selected and reserved resources can occur with respect to the $N = 1$ analysis. For this reason, we assume $N_{SCI} = 2$ in the rest of this Subsection. We should also note that vehicles using the DS are allowed to reserve resources for the retransmission of a TB. As a result, the discussion in this Subsection applies to both the SPS and the DS. Like for the $N = 1$ case, re-evaluation cannot detect potential collisions between the initial transmissions of TBs on selected resources when $N > 1$. This is the case because vehicles transmitting on selected resources have not yet announced their selection, and do not allow the re-evaluation mechanism to detect the collision. In addition to the four cases described when $N = 1$, there are two additional cases

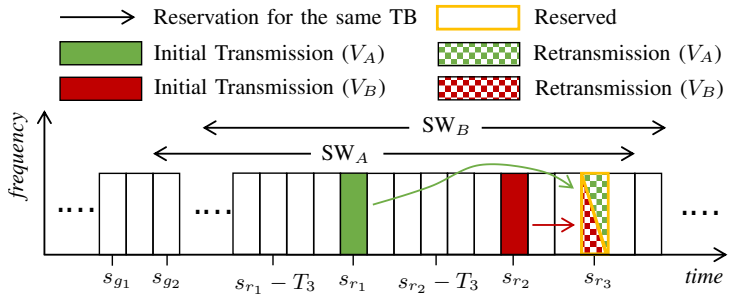

(a) Collision between retransmissions.

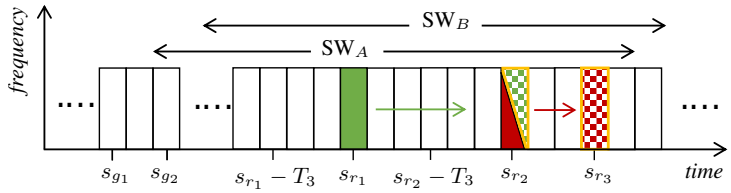

(b) Collision between initial transmission and retransmission.

Figure 2.16: Impact of retransmissions on the re-evaluation detection.

when $N = 2$ where the re-evaluation detection can successfully detect a potential collision. These two additional cases originate from potential collisions that involve

resources reserved for the retransmission of a TB, and therefore do not depend on the employed RRI values. The first case is illustrated in Fig. 2.16(a) where a potential collision between the retransmissions of two TBs is considered. In this figure, the initial transmission of V_A and V_B is performed on collision-free resources at slots s_{r_1} and s_{r_2}, respectively. Due to the overlap between the selection windows of V_A and V_B, let us now assume that the retransmission of both TBs is scheduled on the same resources at slot s_{r_3}, potentially leading to a collision. Before transmitting at s_{r_2}, V_B runs a re-evaluation check at slot $s_{r_2} - T_3$ and senses the reservation announced by V_A for the retransmission of the same TB; this reservation is announced by the SCI associated with the TB's initial transmission. V_B triggers then a re-evaluation detection to select new resources for the retransmission. Note that also V_A runs a re-evaluation check at slot $s_{r_1} - T_3$, but it cannot sense the reservation announced by V_B because $s_{r_2} > s_{r_1} - T_3$. The second case occurs when there is a potential collision between the retransmission and the initial transmission of TBs, and is illustrated in Fig. 2.16(b). In the figure, V_A selects resources at slots s_{r_1} and s_{r_2} for the initial transmission and the retransmission of a TB, while V_B selects resources at slots s_{r_2} and s_{r_3} for the initial transmission and the retransmission of a TB. V_B runs a re-evaluation check at slot $s_{r_2} - T_3$ and senses the reservation announced by V_A at slot s_{r_1}. This reservation included the resources initially selected (and now reserved, since $s_{r_2} - T_3 > s_{r_1}$) at s_{r_2} for the retransmission of the TB by V_A. V_B detects the possible collision between its initial transmission and the retransmission of V_A, excludes the resources initially selected at s_{r_2} from its new selection window, and triggers a re-evaluation detection.

3) *Effectiveness of re-evaluation detections:* this Subsection has identified and analyzed carefully all the circumstances under which a collision can (and cannot) trigger a re-evaluation detection. However, a re-evaluation detection is not always effective in avoiding collisions. An ineffective re-evaluation detection occurs if the reservations that triggered a re-evaluation detection are not finally used for transmitting a TB. To further clarify the notion of effective re-evaluation detection, let us consider the scenario illustrated in Fig. 2.15(a). In this figure, V_B triggers a re-evaluation detection because it detected the imminent collision with V_A at slot s_{r_2}. Then, V_B completes the re-evaluation process to select new resources and avoid the collision. If V_A eventually transmits its next TB using the reserved resources at slot s_{r_2}, then the re-evaluation detection triggered by V_B was effective in avoiding the collision with V_A. Conversely, let us now suppose that the next TB of V_A does not fit in the resources reserved at s_{r_2}, and V_A must perform a size reselection to reserve new resources able to accommodate the size of the new TB. In this case, the resources at slot s_{r_2} are unutilized since both V_A and V_B selected new resources. As a result, the re-evaluation detection has been ineffective since it did not avoid any collision between V_A and V_B. Re-evaluation detections would also be ineffective if V_A performs a latency reselection

or leaves unutilized the resources that it has reserved at s_{r_2}. It is important to point out that vehicles cannot determine in advance if a re-evaluation detection will be ultimately effective or not, except when it is triggered by a reservation for the retransmission of the same TB. Reservations for the retransmission of the same TB always satisfy the size and latency requirements of the generated TB, and they are not subject to latency reselections, size reselections, or unutilized reservations. Therefore, a re-evaluation detection triggered by a retransmission of the same TB is always effective.

2.4.3.3 Resource Replacement

During a re-evaluation, if a vehicle detects a potential collision it triggers the re-execution of step 2 of the resource reselection mechanism as part of the resource replacement phase. The objective is to select new collision-free resources and avoid the identified collision; however, the selection of collision-free resources cannot be fully guaranteed. During the resource replacement phase, a vehicle might select resources that are already occupied by neighboring vehicles and experience a collision on selected resources that cannot be detected by the re-evaluation mechanism. Therefore, the selection of collision-free resources during the resource replacement phase is instrumental to the effectiveness of the re-evaluation mechanism. Since such collision-free selection cannot be always guaranteed, it is necessary to evaluate the actual effectiveness of the re-evaluation mechanism.

2.5 NR-V2X Mode 2: Numerical Results

In this Section, the impact of different traffic types, PDB requirements, and the latest MAC sublayer features on the NR-V2X Mode 2 performance is numerically quantified, leveraging the MoReV2X simulator and following the 3GPP evaluation guidelines. The obtained results corroborate the qualitative analysis of NR-V2X Mode 2 provided in Section 2.4 and thoroughly assess the achievable performance of NR-V2X Mode 2.

2.5.1 Simulation Environment

The examined scenario consists of a 5 km long highway segment with 3 lanes per driving direction. To avoid border effects, performance metrics are collected only in the central 2 kilometers. The highway trunk is populated considering three different vehicular density values: 25, 50, and 100 vehicles/km. The speed of the vehicles is 70 km/h. In all simulations, the OFDM numerology is $\mu = 1$. Accordingly, SCS = 30 kHz and $t_s = 0.5$ ms. NR-V2X radios are configured to operate on a 20 MHz channel in the 5.9 GHz ITS band and employ 12 RBs long subchannels. As a result, the total number of available subchannels in every time slot is 4. TBs are transmitted with a 16QAM-0.5 MCS. The transmission power is set to 23 dBm and the receiver sensitivity to -103.5 dBm, according to the prototype data in [68].

For both the SPS and the DS scheme, we set the processing delay times $T_{proc,0}$, T_0 and T_3 equal to 1 slot, 1100 ms (equivalent to 2200 slots with a subcarrier spacing of 30 kHz) and 5 slots respectively. The limits of the selection window, T_1 and T_2, are set equal to 2 slots and to the PDB, respectively. At MAC sublayer, $\beta = 20\%$ and the initial RSRP threshold is set to its minimum value, i.e., -128 dBm, according to the results obtained in [28]. In the SPS strategy, the keep probability is $P = 0$. The most relevant simulation parameters are summarized in Table 2.5.

Table 2.5: Simulation parameters

Parameter	Values
Traffic density	25, 50, 100 vehicles/km
Highway length	5 km
Number of lanes	6 (3 per driving direction)
Vehicles' speed	70 km/h
OFDM numerology μ	1
SCS	30 kHz
Time slot duration t_s	0.5 ms
Channel bandwidth	20 MHz
Subchannel size	12 RBs
Available subchannels	4
MCS	16QAM-0.5
Transmission power	23 dBm
Receiver sensitivity	-103.5 dBm
RSRP threshold	-128 dBm
Keep probability P	0

2.5.1.1 Traffic Models

In line with the traffic models recommended by 3GPP and reported in Subsection 1.5.1, the performance of the SPS and DS strategies is analyzed considering two different types of traffic:

- Periodic traffic, Fixed size (PF): when this traffic model is adopted, vehicles periodically generate constant size packets every T ms. The packet size is 190 bytes. By default, the PDB of PF traffic coincides with the generation period, i.e., PDB = T.
- Aperiodic traffic, Fixed size (AF): in this case, the inter-arrival time between packets, τ, follows the definition reported in Chapter 1, Subsection 1.5.1. The packet size is fixed and equal to 200 bytes. Unless otherwise stated, the PDB of AF traffic is set to c.
- Aperiodic traffic, Variable size (AV): this type of traffic is characterized by the same description of the inter-arrival time as the AF model. The packet size is uniformly distributed in the [200:200:1200] bytes range. Like in the AF case, the PDB of AV traffic is set to c.

According to the adopted MCS, the set of packet size values considered in this study, i.e., $[190, 200, 400, 600, 800, 1000, 1200]$ bytes, is accommodated over $[1, 1, 2, 3, 3, 4, 4]$ subchannels, respectively. This allows to conclude that PF and AF traffic occupy an average number of subchannels equal to 1, whereas AV traffic occupies an average number of 2.83 subchannels.

In this Section, the PF, AF, and AV traffic models are employed to analyze the performance of the SPS strategy in two different scenarios: single traffic and mixed traffic. In the single traffic scenario, all vehicles generate packets with an average inter-arrival time of 100 ms. Accordingly, $T = 100$ ms in the PF traffic case and $c = \bar{r} = 50$ ms for AV traffic. In the mixed traffic scenario, the 80% of vehicles generates traffic with an average inter-arrival time of 100 ms, whereas the remaining 20% generates traffic with an average inter-arrival time of 20 ms. In the latter case, $T = 20$ ms and $c = \bar{r} = 10$ ms for the PF and AV traffic models, respectively. Following 3GPP terminology, the 100 ms and 20 ms average inter-arrival time settings are termed low intensity and high intensity traffic profiles, respectively.

As mentioned in Subsection 2.4.1, the RRI selection can have a significant impact on the number of latency reselections and unutilized reservations which characterize the SPS strategy, ultimately affecting its collision probability. For this reason, we examine two different strategies for the selection of the RRI:

- Average RRI: the RRI is set equal to the average packet inter-arrival time.
- Minimum RRI: the RRI is set equal to the minimum packet inter-arrival time. This strategy seeks to avoid latency reselections.

Note that the two RRI strategies are characterized by the same RRI value when PF traffic is considered, since the inter-arrival time between packets is constant. Conversely, vehicles generating AV traffic set RRI $= c+\bar{r}$ and RRI $= c$ when the average RRI and the minimum RRI strategies are employed, respectively.

2.5.1.2 Performance Metrics

In addition to the performance metrics recommended by 3GPP and reported in Subsection 1.5.3, we define the following set of additional metrics to analyze the operation of the SPS strategy and of the re-evaluation mechanism in NR-V2X Mode 2:

- Counter Reselections Ratio (CRR): fraction of TBs that triggered a counter reselection with respect to the total number of transmitted TBs.
- Size Reselections Ratio (SRR): fraction of TBs that triggered a size reselection with respect to the total number of transmitted TBs.
- Latency Reselections Ratio (LRR): fraction of TBs that triggered a latency reselection with respect to the total number of transmitted TBs.

- Unutilized Reservations Ratio (URR): ratio between the number of reservations that are left unutilized and the total number of reservations. The URR metric does not take into account the reservations that are left unutilized after a size or latency reselection.

- Re-evaluation Check Ratio (ReCR): fraction of TBs that have been checked for re-evaluation at least once with respect to the total number of transmitted TBs.

- Re-evaluation Detection Ratio (ReDR): fraction of TBs that experienced at least one re-evaluation detection with respect to the total number of transmitted TBs.

- PRR-Re-evaluation: PRR of the TBs for which at least one re-evaluation has been detected. In the case of retransmissions, this metric is separately obtained for each of the N TB transmissions.

- Ineffective Re-evaluation Detection Ratio (IReDR): fraction of TBs that experienced an ineffective re-evaluation detection. As mentioned in Subsection 2.4.3.2, a re-evaluation detection is ineffective when the reservations that triggered the re-evaluation detection are not finally utilized for transmitting a TB.

2.5.2 Impact of Re-Evaluations

This Subsection analyzes the operation of the re-evaluation mechanism, assessing its impact on NR-V2X Mode 2 performance when both the SPS and the DS strategy are examined [11]. The re-evaluation mechanism is carefully analyzed considering both periodic and (fixed or variable size) aperiodic traffic models, relying on the MoReV2X simulator. This Subsection also evaluates the impact of retransmissions on the effectiveness of re-evaluation.

2.5.2.1 SPS Strategy Without Retransmissions ($N = 1$)

This Subsection analyzes the impact of re-evaluations on the operation and performance of SPS when $N = 1$, i.e., when each TB is transmitted once with no retransmissions. First, we focus on the mixed traffic scenario with vehicles transmitting AV traffic. This is a key target scenario since most V2X services to be supported by NR-V2X generate this type of traffic. As shown in Subsection 2.4.1, the generation of packets with a variable inter-arrival time and with variable size can create instability in the operation of SPS. This instability is due to frequent unutilized reservations, size reselections, and latency reselections which increase the probability of packet collisions. Re-evaluation was introduced to mitigate the impact of traffic variability on the SPS strategy.

Due to the variability introduced by AV traffic, more than 50% of the generated packets are transmitted in selected resources; hence, they are eligible for a re-evaluation check. This is visible in Table 2.6, which reports the different metrics for the two RRI selection strategies and all traffic densities. We should note that the ReCR, SRR, LRR and URR

metrics do not vary with the vehicular density. These metrics exclusively depend on the traffic type and on the reservation strategy employed by each vehicle. The table shows that the ReCR is larger than 50% for both the average RRI and the minimum RRI strategies. Vehicles execute a large number of re-evaluation checks because they transmit the majority of packets in selected resources. This is due to the large number of size reselections, latency reselections, and unutilized reservations that affects the SPS strategy operation when AV traffic is examined (see SRR, LRR, URR in Table 2.6). Note that the average RRI strategy reduces the URR, but augments the SRR and LRR metrics, while the minimum RRI strategy minimizes the SRR and the LRR at the cost of increasing the URR. Table 2.6 also shows that re-evaluation is able to detect a larger number of packet collisions as the vehicular density increases. In this regard, note that the ReDR is larger than 16% in the 100 veh/km, average RRI setting. The ReDR increases to over 44% with the minimum RRI selection strategy.

Table 2.6: SPS strategy, AV traffic, mixed traffic scenario: performance metrics wen $N = 1$.

Scenario	ReCR	SRR	LRR	URR	25 veh/km		50 veh/km		100 veh/km	
					ReDR	IReDR	ReDR	IReDR	ReDR	IReDR
Average RRI	60.9	27	57	4	10.7	7.7	14.2	10.2	16.2	11.7
Minimum RRI	57.6	4	3	55	37.3	23.5	41.9	26.4	44.6	28.1

Despite Table 2.6 shows that the re-evaluation mechanism is able to detect a large number of potential packet collisions (see the ReDR metric), Fig. 2.17 reveals that re-evaluation is not fully effective in avoiding collisions and in improving the PRR; this is true regardless of the RRI selection strategy. Fig. 2.17 compares the SPS strategy performance when re-evaluation is implemented and when it is not: Fig. 2.17(a) and Fig. 2.17(b) report the PRR as a function of the transmitter-receiver distance for two different traffic densities (25 veh/km and 100 veh/km), whereas Fig. 2.17(c) concentrates on the CLR in the 100 veh/km case. All figures show that the performance is nearly identical when utilizing re-evaluation and when not. The reason why re-evaluation is not effective in avoiding packet collisions and improving the PRR with AV traffic is two-fold. First, re-evaluation cannot detect collisions that occur between selected resources, as illustrated in Subsection 2.4.3.2. Second, traffic variability triggers size or latency reselections, eventually increasing the probability of transmitting a TB on selected, rather than reserved, resources. Since re-evaluation cannot detect collisions that occur between selected resources, traffic variability increases the probability of having collisions that cannot be detected by the re-evaluation mechanism. In addition, we should note that re-evaluations may not be effective if the reservations that triggered a re-evaluation detection are not finally used for transmitting a TB. In this case, vehicles select new resources to avoid a collision that never happened, and the resource replacement phase cannot guarantee that an undetectable collision will not happen in the newly selected resources. In our analysis, 72% (average RRI) and 63% (minimum RRI) of the resource reservations that triggered a re-evaluation detection were not finally used for transmitting a TB under all evaluated vehicle densities. The

ineffectiveness of the re-evaluation mechanism is reflected in the IReDR metric reported in Table 2.6, and negatively impacts the PRR-Re-evaluation metric, as shown in Fig. 2.17(d). Fig. 2.17(d) shows that the PRR of the TBs that perform a resource replacement after a re-evaluation detection deteriorates compared to the PRR obtained when the re-evaluation mechanism is not implemented.

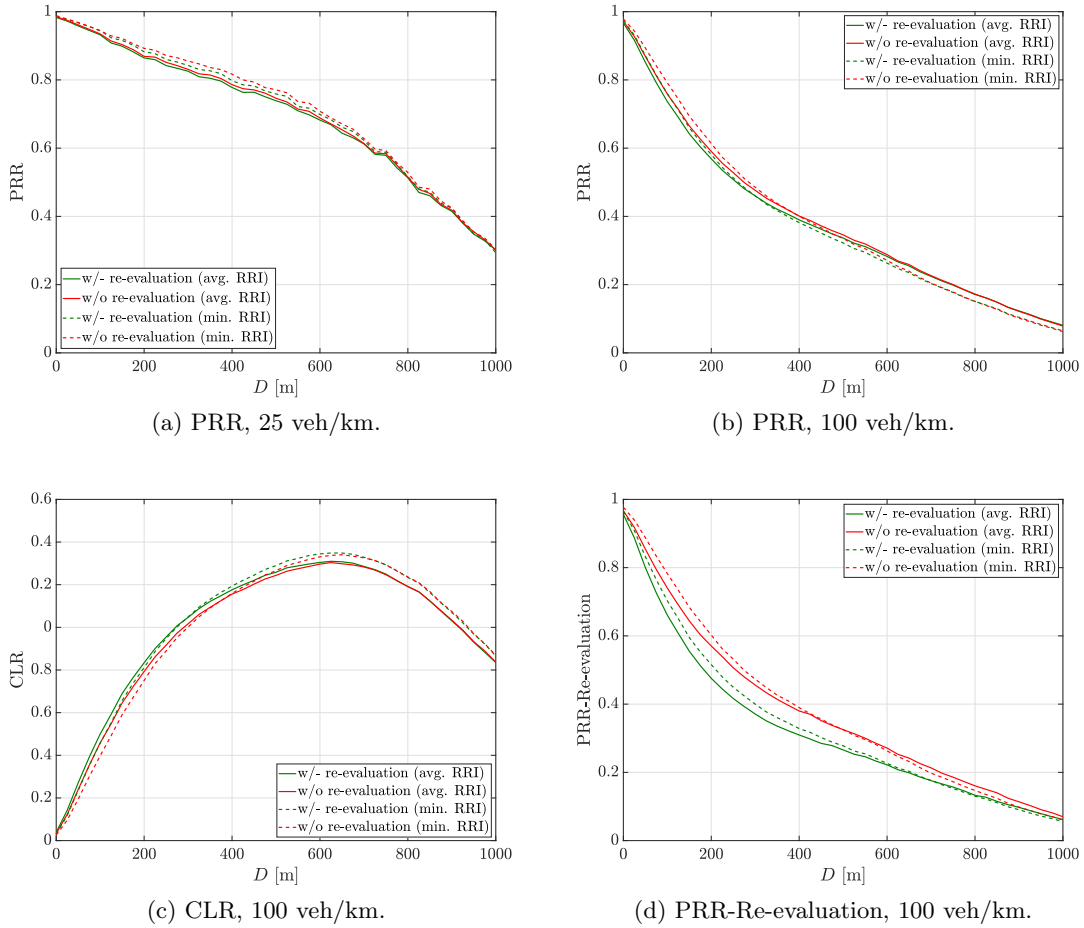

(a) PRR, 25 veh/km.

(b) PRR, 100 veh/km.

(c) CLR, 100 veh/km.

(d) PRR-Re-evaluation, 100 veh/km.

Figure 2.17: AV traffic, $N = 1$: SPS performance in the mixed traffic scenario.

We analyze now the impact of the re-evaluation mechanism on the SPS strategy when vehicles generate PF traffic. Differently from AV traffic, PF traffic does not generate any size reselection, latency reselection, or unutilized reservation. As a result, PF traffic is characterized by a smaller number of collisions between selected resources that cannot be detected by the re-evaluation mechanism. The impact of these undetected collisions that are not avoided by re-evaluation can be visualized in Fig. 2.18, which compares the PRR with PF and AV traffic for the same vehicular density when re-evaluation is implemented. This figure clearly shows how these undetected collisions deteriorate the PRR in the AV traffic case, and their impact increases with the vehicular density.

PF traffic can be affected by persistent collisions when two or more vehicles select

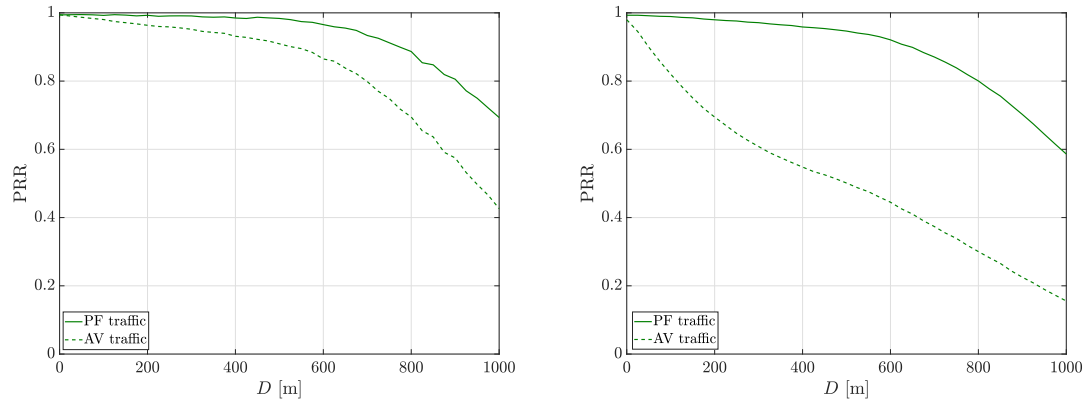

(a) 25 veh/km. (b) 100 veh/km.

Figure 2.18: SPS strategy, single traffic scenario, $N = 1$: PRR curves for PF and AV traffic (minimum RRI strategy). Similar trends are observed in the mixed traffic scenario and with the average RRI strategy.

the same resources within overlapping selection windows[1]. These persistent collisions cannot be detected by the re-evaluation mechanism, since re-evaluation cannot detect collisions between selected resources. With PF traffic, collisions persist until one of the vehicles depletes its reselection counter and executes a resource reselection. We should note that only TBs transmitted after the reselection counter depletes are eligible for a re-evaluation check since they are transmitted on selected resources. With RRI = 100 ms, C_{resel} is in the $[5, 15]$ range, and the ReCR is on average equal to 10% for the single traffic scenario, as reported in Table 2.7; similar trends are observed for the mixed traffic

Table 2.7: SPS strategy, PF traffic: performance metrics (in %) when $N = 1$.

Scenario	ReCR	25 veh/km		50 veh/km		100 veh/km	
		ReDR	IReDR	ReDR	IReDR	ReDR	IReDR
Single traffic	10.3	0.006	0	0.01	0	0.03	0
Mixed traffic	5.5	0.1	0	0.2	0	0.5	0

scenario. Out of the limited set of TBs that are eligible for a re-evaluation check, a vehicle can only use re-evaluations to detect a collision under the conditions illustrated in Fig. 2.14. These conditions require that the reservation that causes the collision is made by a vehicle in a 2-slot time interval just before the generation of the TB. Such an unlikely condition leads to the extremely small ReDR values reported in Table 2.7 and to the small impact of re-evaluation on the CLR illustrated in Fig. 2.19(a). Nevertheless, the vehicles that did execute re-evaluation avoided the persistent packet collisions generated by an initial collision between a selected and a reserved resource. The avoided persistent packet collisions affected on average the transmission of 5.65 consecutive TBs (100 veh/km, RRI = 100 ms). Fig. 2.19(b) reports the PRR-Re-evaluation metrics and shows significant

[1]Using [59], we can estimate that around 30% of packets that trigger a resource reselection experience persistent collisions in the 100 veh/km setting.

gains compared to the performance obtained if re-evaluations were not implemented. In the PF traffic case, re-evaluations were effective to avoid the fairly small number of detected collisions.

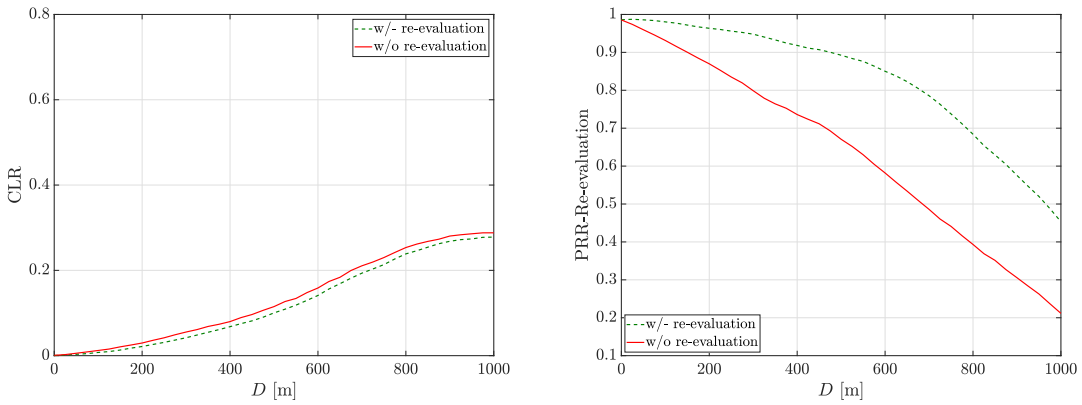

(a) CLR. (b) PRR-Re-evaluation.

Figure 2.19: SPS strategy, PF traffic, mixed traffic scenario, $N = 1$: CLR and PRR-Re-evaluation curves obtained in the 100 veh/km setting. Similar trends are observed in the single traffic scenario.

2.5.2.2 SPS Strategy With Retransmissions ($N = 2$)

This Subsection evaluates the impact of re-evaluations on SPS considering that each TB is transmitted twice ($N = 2$): an initial transmission and a blind retransmission. When N = 2, SPS selects 2 candidate resources that are separated by less than 32 slots for the initial transmission and the retransmission of its TB. In this case, the 1st-stage SCI transmitted with the initial transmission of the TB announces the resources reserved for the retransmission of the same TB, and for the initial transmission and retransmission of the next TB. As discussed in Subsection 2.4.3.2, this results in additional situations in which re-evaluation can detect collisions with respect to the case without retransmissions ($N = 1$). This includes possible collisions between retransmissions, and between initial transmissions and retransmissions.

Table 2.8: SPS strategy, AV traffic: performance metrics (in %) when $N = 2$ (average RRI strategy).

Scenario	ReCR	SRR	LRR	URR	25 veh/km		50 veh/km		100 veh/km	
					ReDR	IReDR	ReDR	IReDR	ReDR	IReDR
Single traffic	74	29	64	3	25	1.1	39.3	2.3	44	4.2
Mixed traffic	58.3	29	67	3	28.8	15.2	31.1	15.7	29.8	14.9

Table 2.8 reports the performance metrics when $N = 2$ and AV traffic is examined. The table shows that the ReDR increases to more than 25% in both the single and mixed traffic scenarios, compared to the 10.7% observed when $N = 1$ (see Table 2.6). Traffic variability

can still impact the initial transmission of TBs when $N = 2$. However, retransmissions do not generate any size reselection, latency reselection, or unutilized reservation, as the resources reserved for the retransmission of the same TB always satisfy the size and latency requirements of the generated packet. This brings some stability to the operation of the SPS strategy and mitigates the impact of ineffective re-evaluations. Re-evaluation detections triggered by reservations for the retransmission of the same TB are always effective, since they always avoid an imminent collision. The numerical results reveal that, in the single traffic scenario, more than 88% of the re-evaluation detections are triggered by reservations made for the retransmission of the same TB.

Therefore, the re-evaluation mechanism is able to improve the performance of the TBs that experienced at least a re-evaluation detection, as shown by the PRR-Re-evaluation curves reported in Fig. 2.20(a); this was not the case without retransmissions ($N = 1$). Fig. 2.20(a) shows that re-evaluations improve the PRR-Re-evaluation for both initial transmissions and retransmissions when the single traffic scenario is considered. In the mixed traffic scenario, reported in Fig. 2.20(b), less than 37% of the detected re-evaluations are triggered by reservations for the retransmission of the same TB (compared to more than 88% in the single traffic scenario). The remaining 63% of re-evaluation detections is triggered by reservations for the next TB. Reservations for the next TB do not always hold a transmission in the reserved resources and deteriorate the effectiveness of the re-evaluation mechanism. This explains the higher IReDR values measured in the mixed traffic scenario reported in Table 2.8, as well as the limited impact of re-evaluation in Fig. 2.20(b).

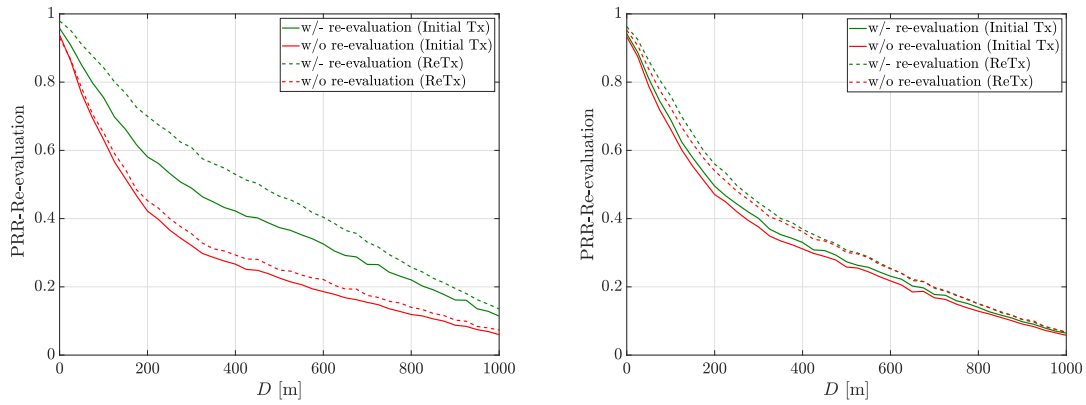

(a) Single traffic scenario.

(b) Mixed traffic scenario.

Figure 2.20: SPS strategy (average RRI strategy), AV traffic, $N = 2$: PRR-Re-evaluation as a function of D in the 50 veh/km setting. Similar trends have been obtained for other densities.

When $N = 2$ transmissions per TB are considered, this Subsection revealed that reservations for the retransmission of the same TB can have a positive impact on the functioning of the re-evaluation mechanism. However, we should note that the re-evaluation mechanism can improve the PRR only if the following conditions are met: 1) both the initial

transmission and the retransmission of a TB experience a collision (a packet is correctly received if just one of the two transmissions is correctly received); 2) re-evaluation can detect at least one of the two collisions; 3) the resource replacement phase is able to select new collision-free resources. In the single traffic scenario, 20% (50 veh/km) and 26% (100 veh/km) of TBs experienced a collision on their initial transmission and retransmission, and the re-evaluation mechanism detected at least one of them. Despite these non-negligible percentages, Fig. 2.21(a) shows that re-evaluation does not significantly improve the PRR. This is the case since the resource replacement phase does not guarantee the selection of collision-free resources. Recall that, during a resource replacement, two or more vehicles might overlap their selection windows and select the same resources, generating a collision between selected resources that cannot be avoided by the re-evaluation mechanism.

When PF traffic is considered, the re-evaluation mechanism is able to effectively avoid the detected collisions, like in the $N = 1$ case. However, only a small fraction of collisions can trigger a re-evaluation detection also when $N = 2$. The attained ReDR levels are smaller than 2% and the impact of re-evaluations on the PRR is negligible.

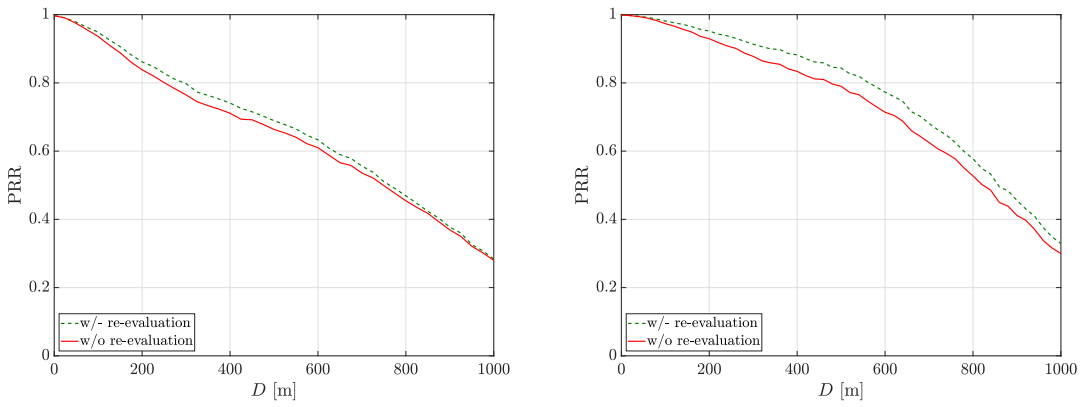

(a) SPS. (b) DS.

Figure 2.21: AV traffic, $N = 2$, 50 veh/km: PRR performance of the SPS (average RRI strategy) and DS schemes in the single traffic scenario. Similar trends have been obtained for other densities.

2.5.2.3 DS Strategy

Vehicles using the DS scheme always transmit the generated TBs on selected resources when only one transmission per TB is considered ($N = 1$). As reported in Subsection 2.4.3.2, collisions that occur between selected (not reserved) resources do not trigger any re-evaluation detection. Re-evaluation has therefore no impact or benefit when using the DS strategy with $N = 1$. Therefore, this Subsection concentrates on the impact of re-evaluations on the DS strategy when $N = 2$. Recall from Subsection 2.4.3.2 that retransmissions occur on reserved resources and can trigger a re-evaluation detection. Although we analyze the impact of re-evaluation on the DS strategy considering AV traffic,

we should note that the performance of the DS scheme does not depend on the generated traffic pattern. As the DS strategy forces the selection of new resources for the initial transmission and the retransmission of every TB, it does not experience any size reselection, latency reselection, or unutilized reservation.

Table 2.9: DS strategy, AV traffic: performance metrics (in %) when $N = 2$.

Scenario	ReCR	25 veh/km		50 veh/km		100 veh/km	
		ReDR	IReDR	ReDR	IReDR	ReDR	IReDR
Single traffic	96.6	22.1	0	37.7	0	54.7	0
Mixed traffic	84.2	21.4	0	34.3	0	47.1	0

Table 2.9 reports the ReCR and ReDR metrics that characterize the DS strategy in the single and mixed traffic scenarios. Table 2.9 shows that the ReCR is equal to 96.6% in the single traffic scenario, i.e., a much larger value with respect to its SPS counterpart in Table 2.8. Such an increase in the ReCR occurs because almost every TB is transmitted on selected resources and is therefore eligible for a re-evaluation check when the DS strategy is examined. This was not the case with SPS because TBs are transmitted on selected resources only after an unutilized reservation or a (counter, size, latency) reselection. Table 2.9 also shows that DS is characterized by fairly large ReDR values in both the single and mixed traffic scenarios. The ReDR values increase with the vehicular density, since a larger density increases the number of potential collisions. With respect to the single traffic case, the ReCR decreases in the mixed traffic scenario. Note that the same trend has also been reported for the SPS scheme, in Table 2.8. During a reselection, vehicles with a smaller RRI have a larger probability of selecting resources that are not eligible for a re-evaluation check (see Subsection 2.4.3.1). With DS, the mixed traffic scenario does not experience additional re-evaluation detection opportunities compared to SPS. Therefore, a smaller ReCR implies a reduction in the measured ReDR levels with respect to the single traffic scenario. As shown in Table 2.9, such ReDR reduction is more evident at larger densities.

Fig. 2.22 depicts the impact of re-evaluations on the CLR when using the DS strategy. In this figure, the CLR is separately measured for the initial transmission of a TB and its retransmission. Fig. 2.22(a) shows that re-evaluation can improve the CLR of both initial transmissions and retransmissions when the channel is lightly loaded. Moreover, Fig. 2.22(a) reveals that re-evaluations are more effective in reducing the CLR experienced by retransmissions since initial transmissions of a TB are accommodated over selected resources and are more prone to experience undetected collisions. This effect is more visible in Fig. 2.22(b), which corresponds to the highest vehicular density, i.e., 100 veh/km. This figure shows that re-evaluations have a negative impact on the CLR of initial transmissions under high channel loads, whereas it improves the CLR of retransmissions. With DS, re-evaluation detection is always effective, and re-evaluation improves the PRR for the TBs for which at least a re-evaluation has been detected. However, the impact of re-evaluations on the PRR is limited also in the DS case, as shown in Fig. 2.21(b): like for the SPS strategy, the re-evaluation mechanism can improve the PRR only if both the

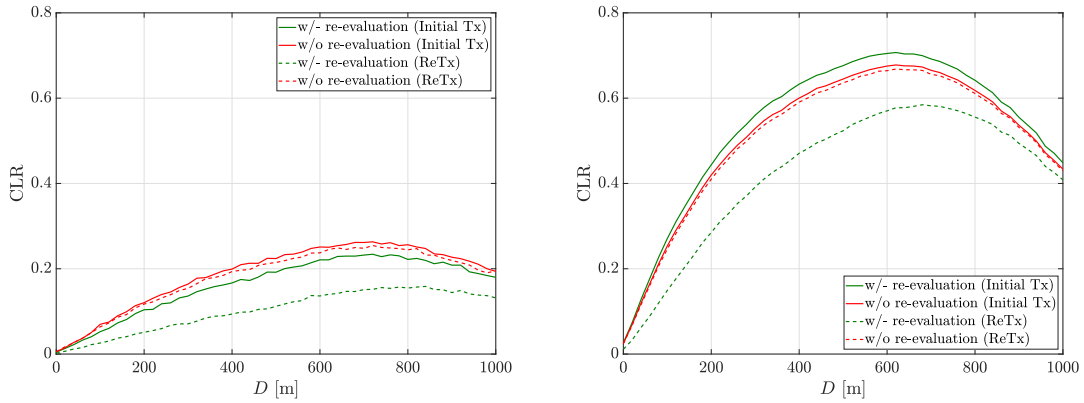

(a) 25 veh/km.

(b) 100 veh/km.

Figure 2.22: DS strategy, AV traffic, single traffic scenario, $N = 2$: CLR as a function of D for two different vehicular densities.

initial transmission of a TB and its retransmission experience a collision, and re-evaluation can detect at least one of them. Despite the large ReDR values reported in Table 2.9, this occurs for only the 0.35%, 2.8% and 9.8% of the TBs in the 25 veh/km, 50 veh/km and 100 veh/km settings, respectively. In addition, the impact of re-evaluation on the PRR is limited by the accuracy of the resource replacement phase. As illustrated in Fig. 2.22, the selection of collision-free resources during the resource replacement phase is not guaranteed (especially when the channel load is large) and vehicles are prone to experience potentially undetected collisions after the resource replacement.

2.5.3 Comparative Analysis of the SPS and DS Strategies

The previous results revealed that the re-evaluation mechanism is not able to improve the system performance in NR-V2X Mode 2, and that the SPS strategy is affected by additional (size or latency) reselections and unutilized reservations when aperiodic traffic is examined, like in LTE-V2X Mode 4. We now conclude the analysis of NR-V2X SL communications, providing an exhaustive comparison of the SPS and DS strategies when the 3GPP-defined traffic models and different latency constraints are considered. A specific attention is paid to analyzing the impact which different PDB choices have on the SPS and the DS strategy performance.

Simulations have been performed following the same assumptions presented in Subsection 2.5.1 and summarized in Table 2.5, with a minor modification: the considered vehicular densities are 50, 120, and 300 veh/km. Here we consider only the low traffic intensity profile and, based on the results presented in [10], vehicles adopting the SPS scheme always reserve resources using the minimum RRI strategy. As a result, vehicles generating PF traffic set RRI $= T = 100$ ms, and vehicles generating AF or AV traffic set RRI $= 50$ ms. Moreover, the PF, AF and AV traffic models are employed to analyze and compare the performance of the SPS and DS strategies in two different settings: single

model and mixed models. In the single model scenario, all vehicles employ the same traffic model (PF, AF or AV). In the mixed models scenario, a percentage Δ of vehicles generates PF traffic, whereas the remaining $(100 - \Delta)\%$ generates AV traffic.

2.5.3.1 Impact of the PDB on the SPS and DS Strategies

Before delving in the comparison between the performance attained by the SPS and DS strategies, this Subsection elaborates on the impact that the PDB has on the collision probability which characterizes the two access schemes.

Fig. 2.23 reports the CLR of the SPS scheme as a function of the average transmitter-receiver distance D in the single model scenario, when the PF, AF, and AV traffic models are examined. The vehicular density is set to 120 veh/km and three representative PDB values are taken into consideration: a PDB value which coincides with the RRI of the SPS scheme, i.e., PDB = RRI, PDB = 25 ms, and PDB = 10 ms. In [55], 25 ms and 10 ms represent the maximum end-to-end latency required by the lowest and the highest degree of automation in cooperative driving applications. In the PDB = RRI case, the PDB is set to 100 ms when PF traffic is examined, whereas it is equal to 50 ms for the AF and AV traffic models. Note that the PDB = RRI choice has been employed in many existing studies, e.g., [9, 10, 15].

Fig. 2.23 reveals that the CLR curves perfectly overlap for the three different PDB values when considering PF traffic, showing that the collision probability is not affected by increasingly stringent PDB requirements. It was qualitatively observed in Subsection 2.4.2 that a smaller PDB reduces the amount of available subchannels included in the selection window, but at the same time decreases the number of vehicles whose selection windows overlap. Here, we numerically demonstrate that these two effects cancel out for PF traffic, since the CLR is not affected by the PDB.

When the AF traffic model is considered, Fig. 2.23 shows that the CLR performance deteriorates as tighter PDB requirements are considered. The same trend is observed with

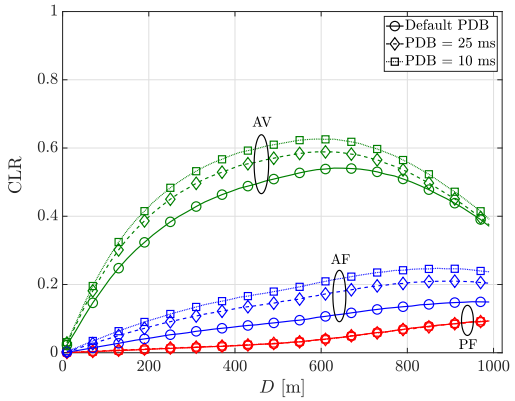

Figure 2.23: SPS strategy: CLR as a function of D.

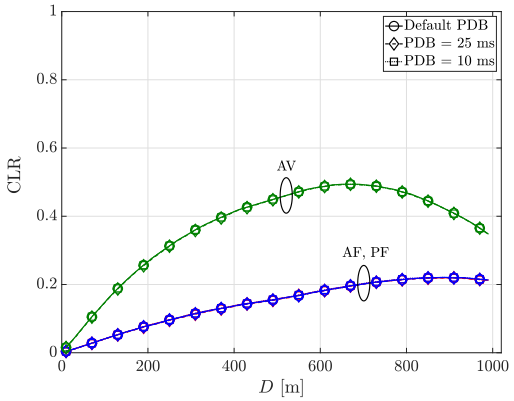

Figure 2.24: DS strategy: CLR as a function of D.

Table 2.10: CBR Values, Single Model Scenario

	50 veh/km	120 veh/km	300 veh/km
PF	0.09	0.22	0.47
AF	0.09	0.22	0.45
AV	0.23	0.46	0.79

AV traffic. For a given vehicular density, we should note that AV traffic is characterized by a larger CBR compared to PF and AF traffic models, since the average number of subchannels occupied by AV traffic is 2.83 times larger. Table 2.10 reports the CBR values computed for the different vehicular densities and traffic models. As anticipated in Subsection 2.4.2, tighter PDB requirements increase the number of latency reselections and accordingly augment the collision probability when the RRI does not match the inter-arrival time between messages. Table 2.11 reports the values of the LRR metric for the SPS scheme. From Table 2.11, the upsurging of latency reselections is manifest when the AF and AV traffic types are examined. Note that the LRR increases as more stringent PDB requirements are considered, justifying the CLR deterioration observed in Fig. 2.23. We should also note that, as the distribution of the inter-arrival time between messages of AF and AV traffic is the same, the LRR metric is coincident in the last two rows of Table 2.11.

Table 2.11: SPS Scheme, Single Model Scenario - LRR.

	PDB = RRI	PDB = 25 ms	PDB = 10 ms
PF	0	0	0
AF	0	0.48	0.77
AV	0	0.48	0.77

Next, Fig. 2.24 quantifies the impact of the PDB on the CLR performance of the DS scheme. This figure corroborates the conclusions provided by the MAC sublayer analysis of Subsection 2.4.2 with simulation results. For each traffic type, more stringent PDB requirements do not affect the collision probability of the DS scheme since the CLR curves referring to the three different PDB choices perfectly overlap.

The independence of DS on the PDB has important implications. According to it, the DS scheme can be used to support V2X applications with a wide range of latency requirements without suffering any performance degradation, regardless of the generated traffic type. On the other hand, the SPS scheme guarantees a similar robustness to PDB variations only when PF traffic is considered. Last, it is worth pointing out that the impact of the PDB on the SPS and DS schemes exclusively depends on the type of generated traffic and not on the vehicular density.

2.5.3.2 SPS and DS Performance Comparison: Single Model

This Subsection compares the performance attained by the SPS and DS schemes in the single model scenario when the PF, AF, and AV traffic models are separately examined. For each traffic model, the impact of different PDB requirements and channel load levels on the PRR is quantified. The evaluation helps identify the conditions under which the two scheduling schemes should be utilized.

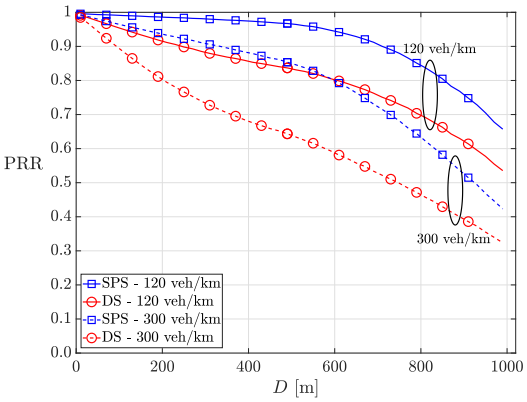

Figure 2.25: Single model scenario, PF traffic, Default PDB: PRR as a function of D.

Fig. 2.25 compares the PRR performance of the SPS and DS schemes as a function of the transmitter-receiver distance D when PF traffic is considered. The PDB is set to 100 ms (PDB = RRI), and two vehicular densities are examined, 120 and 300 veh/km. The former corresponds to CBR = 0.22 and the latter to CBR = 0.47, as reported in Table 2.10. The figure shows that SPS outperforms the DS scheme at both vehicular densities, and that the performance gap becomes larger as the channel load increases. As a matter of fact, the periodic reservation of resources which characterizes the SPS scheme perfectly suits the dissemination of PF traffic and does not generate any latency reselections (LRR = 0 in Table 2.11), size reselections, or unutilized reservations. The SRR and URR metrics are reported in Table 2.12. Since Section 2.5.3.1 demonstrated that the PDB does not have any impact on the SPS and DS performance with PF traffic, Figure 2.25 shows the results for only the PDB = RRI choice. To conclude, the SPS scheme is the best approach in NR-V2X mode 2 to support periodic fixed size traffic, no matter how stringent the PDB requirements are.

Next, Fig. 2.26 compares SPS and DS in the presence of AF traffic for the same

Table 2.12: SPS Scheme, Single Model Scenario - SRR, and URR.

	SRR			URR		
	PDB = RRI	PDB = 25 ms	PDB = 10 ms	PDB = RRI	PDB = 25 ms	PDB = 10 ms
PF	0	0	0	0	0	0
AF	0	0	0	0.47	0.55	0.58
AV	0.07	0.24	0.32	0.47	0.55	0.58

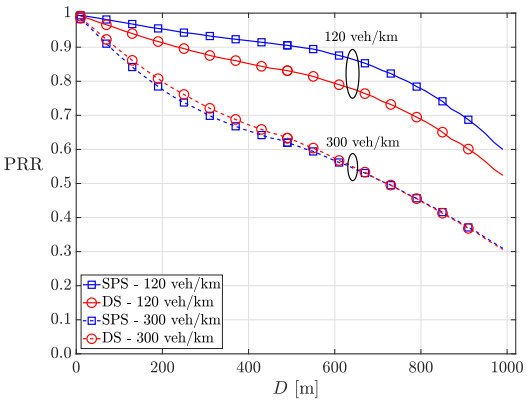

(a) PDB = RRI.

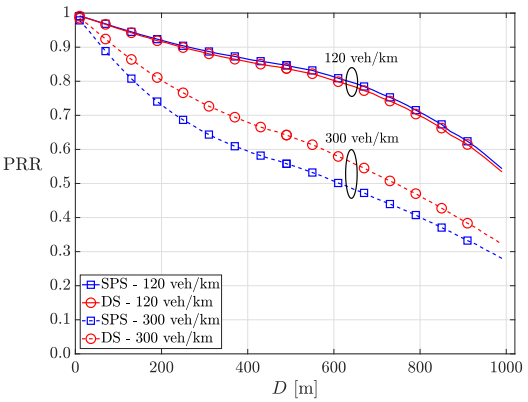

(b) PDB = 25 ms.

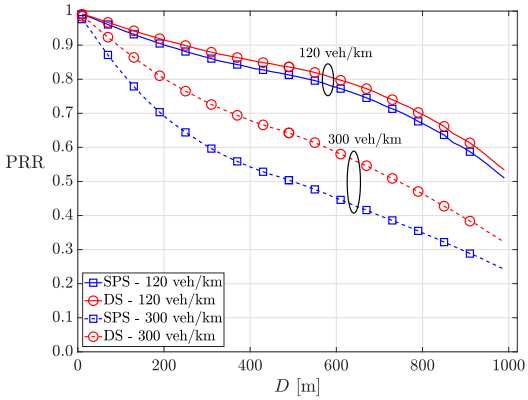

(c) PDB = 10 ms.

Figure 2.26: Single model scenario, AF traffic: PRR as a function of D.

vehicular densities of Fig. 2.25. In Fig. 2.26(a), the PDB = RRI choice is examined. According to it, PDB = 50 ms for AF traffic. Fig. 2.26(a) shows that the SPS scheme achieves a better performance than DS only when the vehicular density is 120 km/h, i.e., when the radio channel is lightly loaded. When the vehicular density increases to 300 km/h and the channel load is no longer negligible, the PRR performance of the two schemes becomes comparable, with DS attaining slightly better values.

In this setting, the degradation of the SPS performance for increasing channel loads is due to unutilized reservations (URR) only, as both the LRR and SRR metrics are equal to zero (see Tables 2.11 and 2.12). When the radio channel is lightly loaded, the fraction of unutilized reservations is not sufficiently large to affect the SPS performance. As the channel load increases, the waste of system capacity associated with unutilized reservations deteriorates the performance of the SPS scheme, which becomes slightly worse than that of the DS scheme. On the other hand, the DS scheme does not experience any latency reselection, size reselection, or unutilized reservation (i.e., LRR = 0, SRR = 0, and URR = 0), since new subchannels are selected for every generated message.

Figs. 2.26(b) and 2.26(c) compare the SPS and DS performance considering smaller

PDB values, i.e., PDB = 25 ms and PDB = 10 ms. For AF traffic, the performance of the SPS scheme is further deteriorated when more stringent PDB requirements are considered, due to the presence of latency reselections (see Table 2.11). Recall from Section 2.4.1 that latency reselections occur as soon as the PDB becomes smaller than the employed RRI. Conversely, the DS scheme operation is not affected by PDB variations, and the corresponding PRR curves do not modify when moving from Fig. 2.26(a) to Figs. 2.26(b)-(c). These results show that DS always attains better or comparable PRR levels than the SPS scheme when vehicles generate AF traffic, independently of the channel load.

In the presence of AV traffic, Fig. 2.27 compares SPS and DS for the PDB = RRI, PDB = 25 ms, and PDB = 10 ms choices. For AV traffic, the SPS and DS performance is compared reducing the vehicular densities to 50 veh/km and 120 veh/km in order to consider the same CBR values analyzed so far. This is necessary since AV traffic occupies an average number of subchannels which is 2.83 times larger with respect to PF and AF traffic, as discussed in Section 2.5.3.1. When the vehicular density is 50 veh/km (CBR = 0.23), Fig. 2.27(a) shows that the gap between the SPS and DS curves is greatly reduced with respect to Fig. 2.26(a). This is the case since the operation of SPS is also

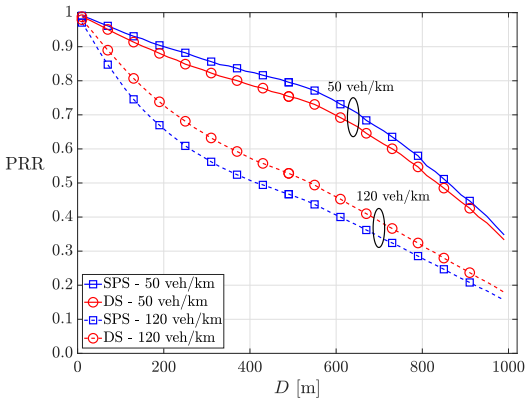

(a) PDB = RRI.

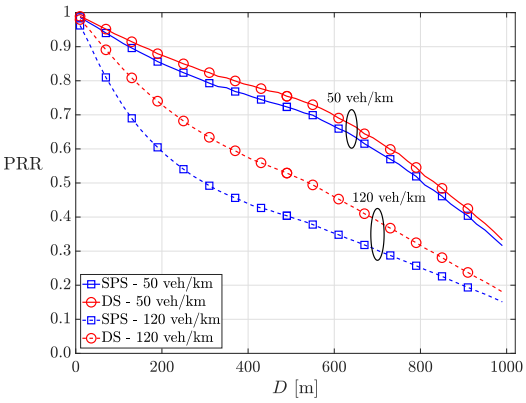

(b) PDB = 25 ms.

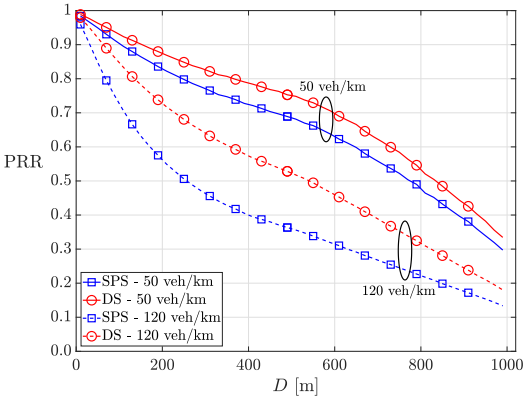

(c) PDB = 10 ms.

Figure 2.27: Single model scenario, AV traffic: PRR as a function of D.

affected by size reselections when AV traffic is considered, as shown in Table 2.12. Size reselections increase the collision probability (see Section III-A) and further deteriorate the PRR performance of the SPS scheme. Accordingly, DS outperforms its SPS counterpart with a larger margin with respect to the AF case when the vehicular density increases to 120 veh/km and CBR = 0.46. The superiority of DS over the SPS scheme becomes even more evident in Figs. 2.27(b) and 2.27(c), where the more stringent PDB = 25 ms and 10 ms values are examined. As in the AF traffic case, latency reselections further penalize the performance of the SPS scheme, which attains the smallest PRR values observed so far and is outperformed by the DS solution for all the considered vehicular densities.

The PRR of the DS scheme is independent of the PDB requirements and exclusively depends on the channel load, thus representing the best solution for serving AV traffic in NR-V2X Mode 2.

2.5.3.3 SPS and DS Performance Comparison: Mixed Models and Adaptive Scheduling

Based on the findings of Subsection 2.5.3.2, this Subsection introduces and evaluates an adaptive scheduling strategy that allows vehicles to select the scheduling scheme that best suits their generated traffic.

This Section compares the SPS and DS strategies in the mixed models scenario, where $\Delta\%$ of the vehicles generates PF traffic and the remaining $(100-\Delta)\%$ generates AV traffic. In this context, we also propose a novel AS strategy that allows vehicles to dynamically select the scheduling scheme that best suits their generated traffic type. Based on the results and conclusions drawn in the previous Subsection, the AS strategy works as follows: if a vehicle generates PF traffic with periodicity T, it employs the SPS scheme with RRI $= T$; if a vehicle generates AV traffic, it uses the DS scheme. Note that the standard does not provide any indication about the circumstances under which the SPS or DS scheme should be used [54], thus allowing the implementation of the adaptive scheduling strategy.

Fig. 2.28(a) compares the performance of SPS, DS, and AS strategy in the 120 veh/km, $\Delta = 10\%$ setting for two different PDB choices, i.e., PDB = RRI and PDB = 10 ms. This figure shows that the best PRR performance is attained by the AS strategy, which allows vehicles to employ the most appropriate scheduling scheme based on the type of traffic they generate. We should also note that the AS strategy performance does not depend on the PDB. With AS, vehicles generating PF and AV traffic use the SPS and DS schemes, respectively. Therefore, the AS strategy does not experience any latency reselections, size reselections, or unutilized reservations, and its collision probability is not affected by more stringent PDB requirements. Figures 2.28(b) to 2.28(d) show that this trend is maintained for increasing values of Δ and that the AS strategy always achieves the best performance, no matter what value of Δ and PDB is examined.

When we concentrate on SPS and DS, which force vehicles to employ a pre-determined

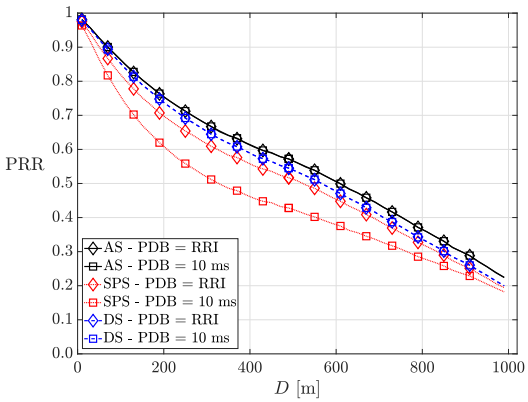

(a) $\Delta = 10\%$.

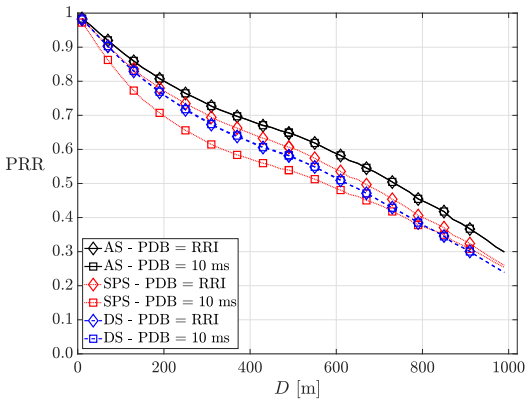

(b) $\Delta = 25\%$.

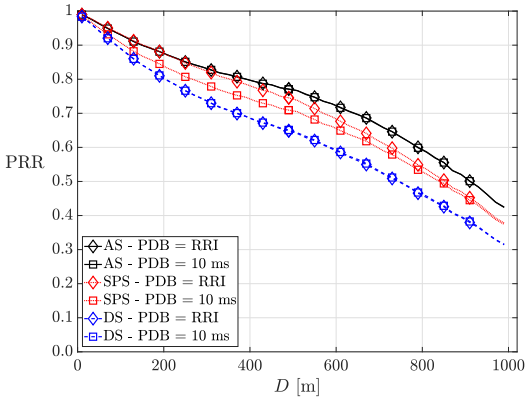

(c) $\Delta = 50\%$.

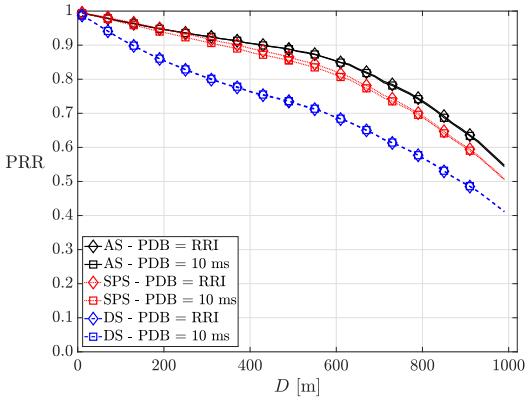

(d) $\Delta = 75\%$.

Figure 2.28: Mixed models, 120 veh/km, $\Delta = 10, 25, 50$, and 75%: PRR as a function of D.

scheduling scheme regardless of the generated traffic type, Fig. 2.28(a) shows that the DS scheme outperforms its SPS counterpart regardless of the PDB. With $\Delta = 10\%$, the majority of the vehicles (90%) generates AV traffic. As a result, the SPS scheme experiences a large number of size reselections and unutilized reservations. In the PDB = 10 ms case, the PRR of the SPS scheme is further deteriorated by the occurrence of latency reselections. The values of the LRR, SRR, and URR metrics which characterize the operation of the SPS scheme are reported in Table 2.13. As already highlighted in Section 2.5.3.2, variations in the message size and in the inter-arrival time between messages can significantly deteriorate the SPS performance, whereas they have no impact on the DS scheme. Indeed, the operation of the DS scheme is not affected by the PDB requirements also in the mixed traffic scenario for any value of Δ, as highlighted from Fig. 2.28(a) to Fig. 2.28(d).

As Δ increases, the percentage of vehicles generating AV traffic reduces and so does the number of latency reselections, size reselections, and unutilized reservations experienced by the SPS scheme. Accordingly, SPS can attain better PRR levels with respect to Fig. 2.28(a). When $\Delta = 25\%$ (Fig. 2.28(b)), the SPS scheme is able to outperform its DS

Table 2.13: SPS Strategy, Mixed Models Scenario - LRR, SRR, and URR

	LRR		SRR		URR	
	PDB = RRI	PDB = 10 ms	PDB = RRI	PDB = 10 ms	PDB = RRI	PDB = 10 ms
$\Delta = 10\%$	0	0.69	0.06	0.28	0.41	0.53
$\Delta = 25\%$	0	0.58	0.05	0.24	0.35	0.44
$\Delta = 50\%$	0	0.39	0.03	0.16	0.23	0.29
$\Delta = 75\%$	0	0.19	0.02	0.08	0.12	0.15

counterpart in the PDB = RRI case. With the more stringent PDB = 10 ms, DS is still superior to the SPS scheme. When $\Delta = 50\%$ (Fig. 2.28(c)) and $\Delta = 75\%$ (Fig. 2.28(d)), PF traffic becomes dominant, and the number of latency reselections, size reselections, and unutilized reservations that characterize the SPS scheme operation significantly reduces (see Table 2.13). As a result, the SPS scheme outperforms the DS scheme for any PDB choice. The impact of the PDB on the SPS scheme becomes less relevant as the percentage of vehicles generating PF traffic increases. In this regard, recall from Section 2.5.3.1 that the PDB does not affect the SPS scheme operation when all vehicles generate PF traffic (i.e., $\Delta = 100\%$).

The mixed scenario has revealed the importance of adapting the scheduling scheme to the type of generated traffic. In this regard, Figs. 2.28(a) through 2.28(d) showed that the AS strategy achieves the best PRR performance regardless of the considered Δ and PDB values. When a pre-determined scheduling scheme is employed, DS is superior to its SPS counterpart only when the majority of vehicles generates AV traffic, due to the impact of unutilized reservations, size reselections and latency reselections on the SPS operation.

2.6 Conclusions

This Chapter has provided a fair and exhaustive analysis of C-V2X SL communications, concentrating on the LTE-V2X Mode 4 and NR-V2X Mode 2 distributed resource allocation modes.

On the LTE-V2X Mode 4 rim, this Chapter has studied the coexistence of periodic and aperiodic traffic sources. A reservation-less variant of the SPS strategy has been put forth and its behavior has been analytically modeled in the limiting condition where aperiodic traffic only is present. Leveraging the MoReV2X simulator, the impact of different percentages and arrival rates of aperiodic flows, size of aperiodic packets and vehicular densities on LTE-V2X Mode 4 performance has been quantified through extensive simulations. The obtained results demonstrate that LTE-V2X Mode 4 is not able to effectively disseminate aperiodic traffic, and that the PRR of aperiodic traffic is insensitive to different latency requirements.

Then, this Chapter has concentrated on NR-V2X Mode 2 and has focused on the new re-evaluation mechanism. Our thorough study has demonstrated that the re-evaluation mechanism is not able to overcome the limitations associated with the dissemination of

aperiodic traffic. In detail, simulation results have shown that the impact of re-evaluations on NR-V2X Mode 2 performance depends on the generated traffic patterns and the system configuration. Without retransmissions, the re-evaluation mechanism is able to detect collisions only with the SPS scheme. However, re-evaluation cannot improve the system performance due to the occurrence of size and latency reselections when aperiodic traffic is considered. Additional reselections reduce the effectiveness of re-evaluations and increase the probability of selecting already occupied resources after the resource replacement phase. In the periodic traffic case, re-evaluation is detected on a very limited amount of collisions and its overall impact on the SPS performance is negligible. With retransmissions, the re-evaluation mechanism is able to avoid a larger number of collisions both in the SPS and DS case. Retransmissions always satisfy the size and latency requirements of the generated packets, and do not trigger any additional reselection. Yet, the impact of re-evaluations on the performance of the SPS and DS strategies is limited, since the benefit of retransmissions prevails over the gains due to re-evaluations. As the implementation of the re-evaluation mechanism implies a significant computational cost, the obtained results question its current design and whether re-evaluation should be mandatory in NR-V2X Mode 2.

Lastly, this Chapter has offered a comparison between the SPS and DS strategies. The obtained results have demonstrated that different PDB requirements do not affect the DS strategy operation, whereas they can significantly deteriorate the performance of the SPS strategy. It has also been shown that the SPS strategy represents the optimal solution for serving fixed size periodic traffic. In the (fixed or variable size) aperiodic traffic case, the best performance is attained by the DS strategy, especially when more stringent PDB requirements are considered. The SPS and DS strategies have been also compared when fixed size periodic traffic and variable size aperiodic traffic sources coexist in the same NR-V2X system. In this case, a novel adaptive scheduling strategy is proposed that allows vehicles to select the scheduling scheme that best suits the type of traffic they generate. Our evaluation demonstrates that this adaptive strategy yields the best performance for all traffic types and PDB requirements.

Chapter 3

Machine Learning for Disseminating Awareness Messages

According to the CAR 2 CAR Communication Consortium (C2C-CC) [69], V2X applications can be organized into three subsequent phases: Day 1, Day 2, and Day 3+.

Day 1 safety applications are already present on vehicles to increase space awareness and grant the car and its driver more time to react to unexpected situations. During Day 1, vehicles are expected to exchange information regarding their status as well as information related to unexpected and potentially dangerous events. Most Day 1 applications currently rely on the DENMs and CAMs [16] specified by ETSI. These messages include basic information such as the position, speed, acceleration, and heading of the transmitting vehicle. The exchange of DENMs and CAMs among vehicles is at the basis of several V2V applications like emergency vehicle warning, stationary vehicle warning, and traffic jam warning.

In the Day 2 phase, V2X communications will enable the exchange of information about objects detected by vehicles using their local sensors, e.g., cameras, LIDARs, and RADARs. As a result, Day 2 applications will allow receiving vehicles to be aware of objects beyond their LOS that could not otherwise be detected. Day 2 applications will mitigate the limitations associated with the coexistence between V2X-enabled vehicles and non-connected road users. At the same time, they will allow the implementation of enhanced safety applications, e.g., VRU protection, advanced intersection collision warning, and Cooperative Adaptive Cruise Control (C-ACC) functionalities.

Beyond the Day 1 and the Day 2 phases, Day 3+ V2X applications will rely on the exchange of intention and coordination data to support the implementation of cooperative driving services. During the Day 3+ phase, Connected and Autonomous Vehicles (CAVs) will leverage V2X communications to share their driving intentions and enhance their automated driving capabilities. Moreover, the Day 3+ phase foresees an active participation of VRUs within the V2X ecosystem. This is not the case in Day 1 and Day 2 phases,

where VRUs have just a "passive" role, i.e., they can be detected by vehicles or the infrastructure, but are not able to exchange messages and notify their presence to other road users.

3.1 Research Motivation and Main Contribution

In Day 1, 2, and 3+ V2X applications, vehicles will disseminate data collected using their on-board sensors to share their environmental perception and driving intentions. As a result, the information locally generated by each vehicle will depend on the surrounding environment and on the transmitting vehicle dynamics, and will likely be characterized by aperiodic patterns and variable size packets [19]. As demonstrated in the previous Chapter, the reliable dissemination of variable size aperiodic traffic is a critical concern in LTE-V2X SL and NR-V2X SL communications.

This Chapter intends to take a fresh look at the problem of aperiodic safety message dissemination in LTE-V2X Mode 4, concentrating on the main traffic type that LTE-V2X was designed to deliver, namely, Day 1 CAMs, and it proposes to harness ML to broadcast such messages. As a matter of fact, ML has stirred an unprecedented interest and consensus in numerous wireless settings. This major branch of artificial intelligence is often seen as the appropriate tool to pick the lock of complex problems encountered in, e.g., radio resource allocation and optimization; with no ambition for completeness, [70–72] represent captivating examples in the field. The survey in [73] offers an excellent portrayal of the recent ML applications to the specific domain of vehicular networks. However, to the authors' knowledge, none of the studies in the field have scouted the adoption of ML in V2V safety communications.

Overall, this Chapter paves a new research path on LTE-V2X Mode 4, as it addresses the question of whether ML can help in serving aperiodic traffic employing the SPS strategy in Mode 4, given the latter is a recognized benchmark for safety communications in a vehicular environment. The obtained results provide a largely positive answer. Mode 4 is therefore enhanced with a ML-based algorithm to predict when CAMs will be generated and when to reserve radio resources on the LTE time-frequency grid. In detail, the current study proposes to interpret the aperiodic CAM sequence as a series of sub-sequences that are periodic over a short time scale, and to rely on ML to forecast the sub-sequence length and periodicity. Then, the idea is to tailor the SPS radio resource reservations so as to fit the period and length of the single sub-sequence forecast by ML, reducing the amount of additional (size or latency) reselections and unutilized reservations. Additionally, the identification of free radio resources performed by Mode 4 is modified and made more effective. The main outcomes are the following:

- ML achieves excellent accuracy in predicting the temporal patterns of CAMs;
- the new, ML-enhanced scheduling of resources outperforms LTE-V2X Mode 4 under all points of view, warranting higher rates of packet delivery, fewer collisions and

better channel utilization.

At the time of this study, the standardization of NR-V2X SL had not been finalized yet. Nevertheless, the obtained results disclose the potential of ML for disseminating aperiodic traffic also in NR-V2X Mode 2, as its SPS strategy has been inherited from LTE-V2X Mode 4 with minor modifications.

The rest of this Chapter is organized as follows. Section 3.2 lists the most significant related works and Section 3.3 discusses the generation rules of CAMs. In Section 3.4, the new ML-based policy to accommodate aperiodic CAM traffic is illustrated in detail. Section 3.5 presents the simulation results and the conclusions are drawn in Section 3.6.

3.2 Related Work

When the literature on CAMs is explored, it is worth recalling [74], which determined an accurate message broadcasting threshold and broadcasting interval to avoid high rate of unnecessary messages. In [75], a novel CAM triggering condition on winding roads was also proposed. Moreover, the authors of [76] explored a simple mechanism to confine the queueing delays suffered by CAMs which coexist with different traffic types. In all of the above papers, the authors modified the generation pattern of CAM traffic. As opposed to such contributions, the current study adheres to the ETSI standard for CAM generation [16], and forecasts when the next CAMs will be produced; it then reserves resources accordingly, modifying LTE-V2X Mode 4 scheduling in a very effective manner.

In analogy to the present study, [77] and [78] investigated the issue of radio resource assignment in vehicular networks. In [77], CAMs were compressed to reduce the channel load; yet, the authors themselves evidenced that compressing and decompressing is time-consuming. In [78], the SPS strategy was tuned to serve different packet sizes. However, it has to be observed that the size of CAMs is not known *a priori* and can vastly vary, which in practice prohibits the adoption of the solution in [78].

On a different rim, [79] and [80] considered message delivery for cooperative awareness, but focused on Carrier Sense Multiple Access/Collision Avoidance (CSMA/CA), the access strategy adopted in the MAC sublayer of 802.11p. Namely, [79] considered a simplified, periodic model for CAM traffic and leveraged on full-duplex transceivers; [80] highlighted the impact of realistic mobility patterns on 802.11p operation.

Additional references are [81–84]: the authors of [81] examined LTE-V2X Mode 3, hence the scenario where the eNB controls the allocation of resources to V2V communications; the authors of [82] investigated a centralized multicast/broadcast approach too. Conversely, the ML-based solution proposed in this Chapter is totally decentralized, as LTE-V2X Mode 4 mandates. The study in [83] faced the design of V2V communications and employed the sub-6 GHz band exclusively for the control plane, whereas the data plane was positioned at mmWave frequencies. Similarly to [83], [84] considered mmWave communications and allowed for multi-hop transmissions among vehicles. On the contrary, the

current investigation is sub-6 GHz centered and concentrates on single-hop transmissions, adhering to the standard for cellular V2V communications.

3.3 ETSI-Generated CAM Sequences

When initially investigating LTE-V2X performance, CAM occurrences were modeled as periodic packets, a choice that perfectly suits the SPS strategy in both LTE-V2X Mode 4 and NR-V2X Mode 2. However, the standard algorithm for the generation of CAMs released by ETSI [16] indicates that the inter-arrival time between consecutive messages, T_{CAM}, is variable. Its duration strongly depends on the vehicle dynamics: if the vehicle modifies its trajectory, if its speed or acceleration/deceleration are sufficiently high, then T_{CAM} shortens and CAMs become more frequent. Further details about the ETSI algorithm that rules the CAMs generation have been reported in Subsection 1.5.1.

Moreover, ETSI details the mandatory and optional fields in a CAM, allowing for variable size messages [16]. The rules of the standard therefore lead to CAM traffic which in most of the cases exhibits aperiodic inter-arrival times and variable message sizes. The last remark is well documented by the numerical results that we presented in [13] and by the experimental survey in [14], which offers an analysis of CAMs collected during actual test-drives. The study reveals that T_{CAM} often changes from one message to the next, that its distribution is very diverse, and heavily dependent on the drive scenario (urban, suburban, or highway). Similar conclusions hold for the size variability of CAMs. The correlation between CAM inter-arrival times and the vehicle behavior is exemplified in Fig. 3.1, where the temporal sequence of T_{CAM} values, i.e., a CAM trace, and the vehicle speed are reported as a function of the T_{CAM} sample index. These patterns refer to a vehicle moving along a straight trajectory, that initially decelerates until a complete stop at T_{CAM} index equal to 4, and accelerates again from T_{CAM} index equal to 23, causing T_{CAM} to accordingly vary. Fig. 3.1 shows that when the vehicle decelerates (accelerates)

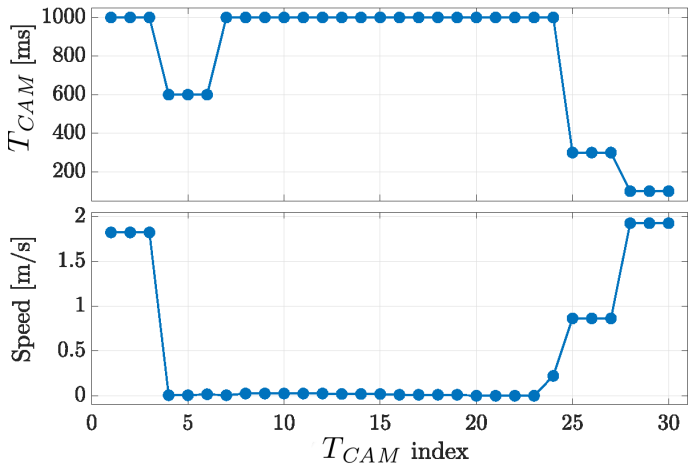

Figure 3.1: Correlation between vehicular speed and T_{CAM}.

in the first (last) portion of the CAM trace, CAMs are more frequently issued. On the other hand, T_{CAM} settles at 1000 ms when the vehicle stops, in the central portion of the trace. Here, variations in heading, position, or speed are not sufficiently large for generating a CAM before the timeout condition, $T_{GenCamMax} = 1000$ ms, occurs. Such a simple, yet exemplary instance is extracted from a wider measurement campaign we performed in different settings [13].

The strong correlation between CAM inter-arrival times and vehicle dynamics suggests that the adoption of ML can be beneficial to predict the temporal evolution of CAM sequences and in turn, lead to an effective reservation policy of radio resources. Indeed, a carefully chosen set of input features, that the vehicle locally retrieves, can be used to feed an ML algorithm, producing the desired outcome, i.e., when next CAMs are likely to occur. The following Section will therefore illustrate a novel approach to deliver aperiodic CAMs, removing the intrinsic inefficiencies that plague the original SPS strategy.

3.4 The Predictive Reservation Framework

Subsection 2.4.1 highlighted the mismatch between aperiodic traffic and the periodic reservation of resources employed by the SPS strategy, causing the undesired phenomena of latency reselections and unutilized reservations. Moreover, Section 3.3 dwelled on the aperiodicity of actual CAM sequences, suggesting that their temporal evolution can be successfully predicted.

The key proposal of this paper is therefore the following: adopt ML to forecast what the next T_{CAM} value will be, and how many occurrences of it will appear. Next, exploit ML prediction to set: (i) the resource reservation interval RRI; (ii) the reselection counter C_{resel}, whose value is no longer randomly chosen, rather, it exactly matches the number of occurrences forecast by ML.

Additionally, the current study significantly intervenes in the list creation phase of the original SPS strategy. As better explained in the next subsection, it builds a more reliable list of available candidate resources than the one produced by the legacy SPS.

3.4.1 Modified SPS Implementation

In our proposed solution, resources are drawn from list L_1, as opposed to list L_2. As a matter of fact, in the presence of aperiodic traffic, L_2 is not as meaningful as when vehicles periodically generate packets. It is not a case that NR-V2X Mode 2 no longer uses L_2 (see Subsection 1.4.2.2). Moreover, our proposal sets the selection window width to 100 ms, the minimum CAM inter-arrival time, to avoid broadcasting out-of-sequence messages. To better understand the last statement, recall that CAM inter-arrival times can take on any value in the $[100, 200, \ldots, 1000]$ ms set; hence, if the selection window W is wider than 100 ms, the $(j+1)$-th CAM might be transmitted before the j-th, an event that has to be prevented.

An additional and meaningful modification concerns the list creation phase of the original SPS strategy. Given the CAM selection window W is 100 ms wide and that the RRI is dynamically determined via ML, RRI $\in [100, 200, \ldots, 1000]$, observe that not all ongoing reservations fall within W and are spotted by the ego-vehicle. It follows that the original list creation mechanism loses effectiveness, increasing the risk of packet collisions. For this reason, we propose a new *look-ahead* version of the SPS strategy. This SPS reworking requires that the SCI also includes the current C_{resel} value, in addition to the RRI. It is a minimal modification with respect to the legacy SCI format, necessitating very few bits. Yet, it remarkably extends the collision-avoidance capability of the original SPS strategy, as the numerical results will show.

As a matter of fact, if the ego-vehicle exploits the C_{resel} knowledge, it can detect potential collisions for any possible combination of the RRIs used by itself and by its neighbors. To further clarify such enhanced capability of identifying collisions, Figs. 3.2(a)-(c) exemplify the SPS strategy operation in three different scenarios. In these figures, the candidate resources examined by the ego-vehicle are represented in green, the subchannels in use by the generic neighboring vehicle are indicated in red, and the selection window in yellow. Furthermore, RRI_{TX} indicates the RRI adopted by the ego-vehicle, whereas RRI_{RX} represents the RRI of the generic nearby vehicle, heard by the ego-vehicle in the SCI it receives. In Fig. 3.2(a), the candidate resource examined by the ego-vehicle is immediately excluded, as it coincides with the reservation placed by the nearby vehicle inside the selection window. The collision is avoided in the case exemplified in Fig. 3.2(b) too, as the ego-vehicle also verifies if any of its future reservations outside of the selection window coincides with the very next resource reserved by the nearby vehicle. Fig. 3.2(c) portrays an instance where the reservation heard by the ego-vehicle is not included within

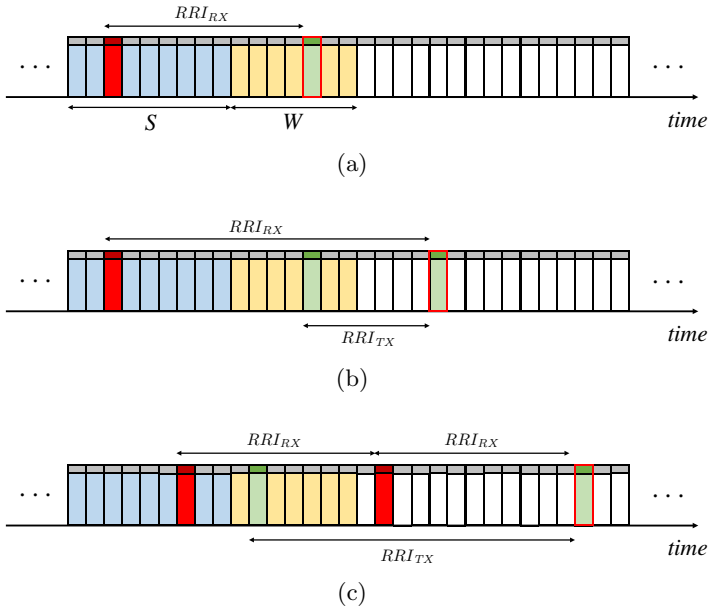

Figure 3.2: SPS detection of potential collisions for different RRI_{TX} and RRI_{RX} settings.

its selection window and RRI_{RX} is lower than RRI_{TX}. In this case, the original SPS strategy cannot detect the future collision, as the ego-vehicle is exclusively aware of the first reservation from the neighboring vehicle, after RRI_{RX} ms, and therefore does not exclude the examined resource. Here, the future collision would be spotted only if the ego-vehicle additionally knew the remaining number of ongoing reservations, i.e., the current C_{resel} value of the nearby vehicle, in addition to the periodicity of ongoing transmissions. Our *look-ahead* version of the SPS strategy proposes to exploit the C_{resel} knowledge and performs this further check. Therefore, it creates a smaller, yet more reliable L_1 list, detecting and avoiding all the potential collisions exemplified in Figs. 3.2(a)-(c).

3.4.2 Machine Learning to Predict CAM Sequences

When the proposed ML-based strategy is employed, the first step that the resource re-selection mechanism accomplishes is to forecast through ML the very next T_{CAM} value, as well as the length of the next sequence of identical T_{CAM} inter-arrival times. To do so, ML explores a large set of CAM traces to identify correlation patterns between a set of dedicated input features and the CAM traces. Then, such knowledge is leveraged to anticipate future CAM inter-arrival times [85]. In this study, the set of input features taken into account are:

- trajectory, current speed, and position of the ego-vehicle;
- current speed and position of the vehicle immediately preceding the ego-vehicle.

We choose to predict the next CAM inter-arrival time through the k-Nearest Neighbors (KNN) ML algorithm, an instance-based learning technique used for both regression and

Algorithm 1: predictive reservation

Input : KNN input features
Output: RRI, C_{resel}

$i = 1$;
T_{CAM_i} = `Predict` (Input features, i);
RRI = T_{CAM_i};
$C_{resel} = 1$;
while $i \leq N$ **do**
  $i = i + 1$;
  T_{CAM_i} = `Predict` (Input features, i);
  if $T_{CAM_i} = RRI$ **then**
    $C_{resel} = C_{resel} + 1$;
  else
    break;
  end
end
if $C_{resel} > 3$ **then**
  $C_{resel} = random[3, C_{resel}]$
end

classification problems. KNN simply stores the training data without attempting to infer a general structure out of them. Moreover, KNN is inherently designed for multi-class problems and its classification consists in assigning the input features the most common label, i.e., the next predicted T_{CAM} value, among the k nearest neighbors. The second action of the ego-vehicle is to dynamically set the (RRI, C_{resel}) pair employed by the SPS strategy in accordance to the ML forecast.

In greater detail, whenever the SPS strategy triggers a resource reselection, Algorithm 1 is invoked. The algorithm exploits KNN to predict the very next T_{CAM} value, T_{CAM_1}, and sets RRI equal to it, i.e., RRI $= T_{CAM_1}$, while C_{resel} is initially set equal to 1. Then, as long as the next predicted inter-arrival time $T_{CAM_{i+1}}$ coincides with the previous T_{CAM_i}, the algorithm keeps incrementing the estimate of the reselection counter C_{resel}. Furthermore, when KNN outcome indicates that more than 3 consecutive CAM inter-arrival times will display the same value, the actual reselection counter value is randomly selected within the $[cntr_{min}, cntr_{max}]$ interval, where $cntr_{min} = 3$ and $cntr_{max}$ is the current C_{resel} estimate. This expedient avoids repeated packet collisions on resources reserved by different vehicles, a circumstance that might occur when vehicles generate CAMs with the same periodicity, e.g., in a congested intersection. The output of Algorithm 1 is finally used to set RRI and the reselection counter C_{resel} that indicates how many times the selected resource is reserved. Note that there is a maximum allowed value for C_{resel}, indicated by N. Moreover, observe that inequality $C_{resel} \geq 1$ reveals that at least one reservation has to be placed.

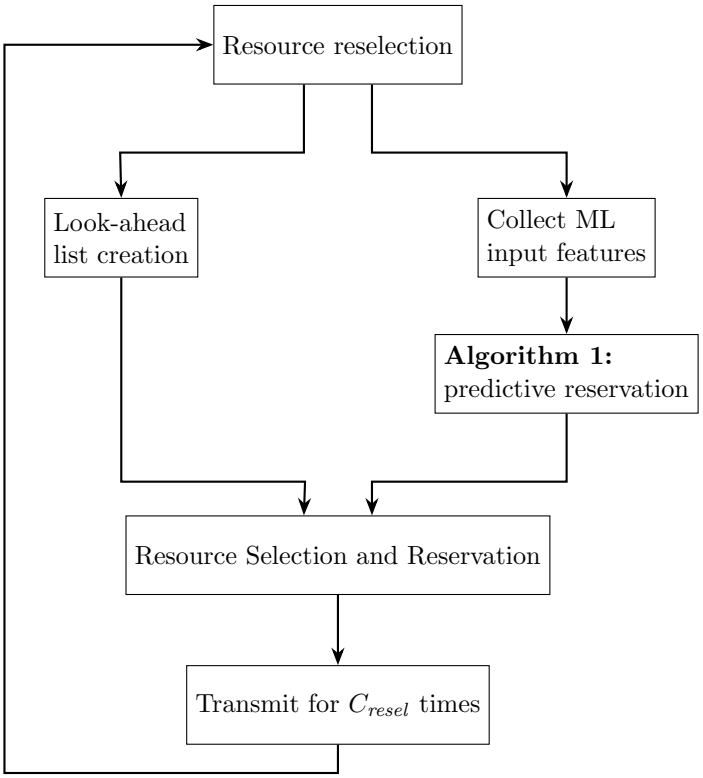

Figure 3.3: Flowchart showing the proposed KNN-look ahead solution.

The overall flowchart of the proposed solution, termed KNN-look ahead from now onward, is reported in Fig. 3.3.

3.5 Numerical Results

3.5.1 PHY and MAC Layer Configuration

As regards the PHY layer and the MAC sublayer, this study relies on MoReV2X simulator. Vehicles have been configured to transmit their messages over a 10 MHz wide channel located in the 5.9 GHz ITS band, with 15 kHz subcarrier spacing. The 10 MHz channel is divided into 4 subchannels that consist of 12 RBs each, assuming adjacent transmission of the TB and of its associated SCI. The size of CAMs, indicated by X, is fixed to either 190 or 470 bytes, which are the smallest and the largest statistically relevant sizes reported in [14]. Vehicles transmit their packets using QPSK modulation with 0.7 code rate, therefore mapping the 190 and 470 byte-long packets into 1 and 2 subchannels, respectively. The transmission power is set to 23 dBm and the receiver sensitivity to -90.4 dBm. As in [56], the RSRP threshold is -140 dBm.

3.5.2 Outcomes

3.5.2.1 Suburban setting

The first set of results refer to some outskirts of the Italian city of Modena, that we classify as a suburban setting example. Here, microscopic vehicular mobility has been simulated through SUMO [20]. The examined road topology is reported in Fig. 3.4 and it has been imported in SUMO using Open Street Map (OSM) [86]. The area is approximately 2.5 km wide and 3 km long. Vehicles have been randomly generated at the area edges and have been assigned random trajectories that traverse the entire topology. The average

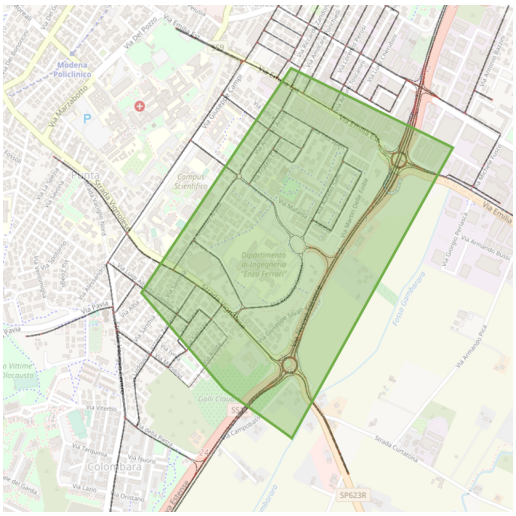

Figure 3.4: The examined suburban road topology.

vehicular density is 42 vehicles/km and the vehicles speed varies in the $[50, 100]$ km/h interval, depending on traffic conditions and on the vehicle *speedFactor*. The *speedFactor* is a SUMO parameter that defines the maximum velocity of each vehicle as a function of the lane speed limits.

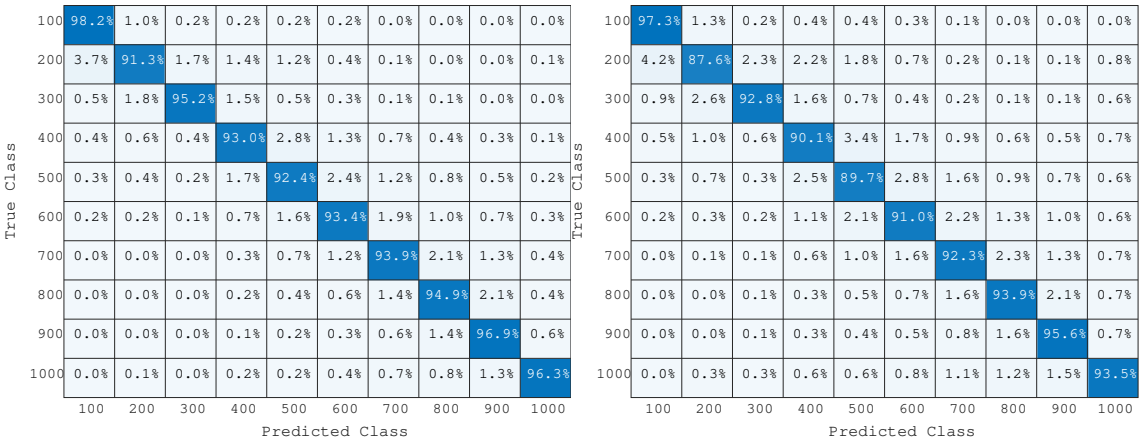

(a) T_{CAM_1} (b) T_{CAM_5}

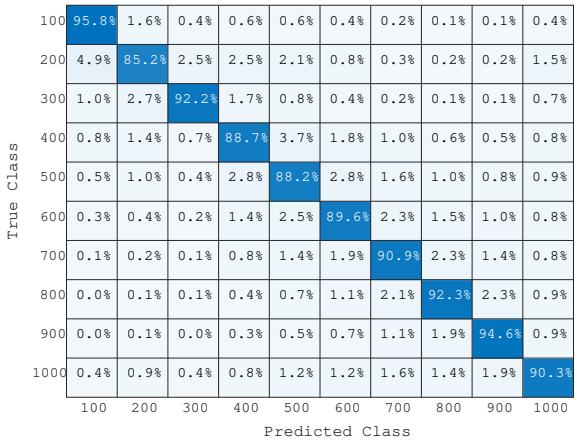

(c) $T_{CAM_{10}}$

Figure 3.5: Confusion matrix for the prediction of the T_{CAM_i} values.

We have additionally developed a set of custom tools based on the SUMO Traffic Control Interface (TraCI) [87], to extract the elements that characterize the behavior of every vehicle, namely, heading, position, and speed; the periodicity for their collection was coincident with $T_{GenCamMin} = 100$ ms. They have allowed us to generate CAM messages in accordance with the rules set by the ETSI algorithm recalled in Sec. 3.3. For every car, we also recorded the position and speed of the preceding vehicle, to complete the set of input features used by ML. As requested by Algorithm 1, these features fed a real-time implementation of the KNN algorithm, to predict the longest sequence of T_{CAM} values with the same periodicity. The number k of KNN nearest neighbors was taken equal to 3.

The dataset of CAM traces was collected from a total of 6800 vehicles during 20 minutes of SUMO simulation. The least represented T_{CAM} values in the dataset were

oversampled using the Synthetic Minority Oversampling TEchnique (SMOTE) [88]. All input features were further normalized using min-max normalization, i.e., their range of values was re-scaled between 0 and 1. The training set and the test set were generated employing a 70-30 split ratio.

First, Figs. 3.5(a)-(c) delve into the ability of the KNN algorithm to predict the upcoming sequence of CAM messages, reporting the confusion matrix for three different values of the T_{CAM} index i, $i = 1, 5$, and 10. The confusion matrix is a two-dimensional matrix indexed with the true and predicted class labels, and it is commonly used to visualize the performance of an algorithm. Fig. 3.5(a) reports the confusion matrix of KNN for $i = 1$, that is, when KNN forecasts next T_{CAM} value. Figs. 3.5(b) and (c) show the confusion matrix when KNN predicts the fifth ($i = 5$) and the tenth ($i = 10$) T_{CAM} value, respectively. These figures reveal that KNN is able to accurately forecast T_{CAM_1} value, and that the degradation in predicting T_{CAM_5} and $T_{CAM_{10}}$ values is modest.

The quality of the prediction outcomes is further highlighted by Fig. 3.6, which reports the accuracy and the macro-F1 score of KNN as a function of i, $i = 1, 2, \ldots, N$, with $N = 10$. For a given T_{CAM} index i, the accuracy measures the fraction of correctly predicted T_{CAM} values over the total number of samples, and the macro-F1 score is evaluated as the mean of class-wise F1-scores. The F1-score is a common ML metric that combines precision and recall [85]. As could have been expected, KNN performance deteriorates for larger values of i, as the CAM inter-arrival time to be forecast is increasingly distant in time. However, both indicators settle on fairly high levels, greater than 0.9, even for $i = 10$.

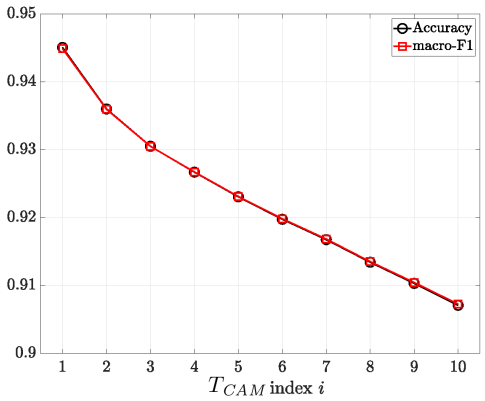

Figure 3.6: Accuracy and macro-F1 score as a function of the T_{CAM} index.

In the next set of figures, the focus shifts to the performance of the proposed KNN-look ahead solution. First, Fig. 3.7 shows the PLR as a function of the distance D between the transmitting and the receiving vehicle. Solid lines refer to $X = 470$, dashed lines to $X = 190$ bytes. Recall that the PLR measures the fraction of packets that were lost because of scarce propagation conditions over the total; as a matter of fact, these curves do not depend on the resource assignment strategy but are exclusively determined by the PHY layer choices and by the CAM size. So, when the radio propagation environment

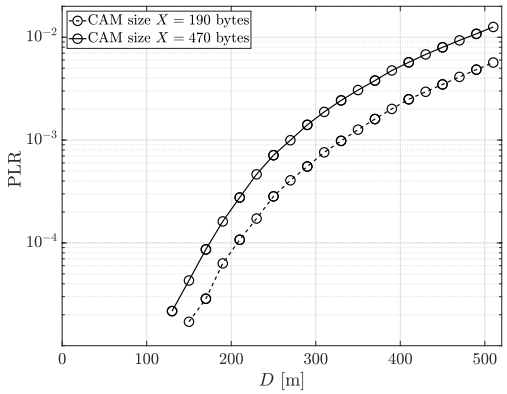

Figure 3.7: PLR as a function of the Tx-Rx distance D.

is more hostile (e.g., greater D values) and the CAM size is longer, the PLR increases. Recall that larger CAMs are transmitted over a wider bandwidth and experience a larger amount of noise at the receiver. For these curves, as well as for the results shown next, a proper number of simulations has been executed to obtain sufficiently tight 95% confidence intervals. To avoid border effects, the results have been collected only from the central area of the setting; this corresponds to the green-shaded area in Fig. 3.4.

In the following figures, the proposed approach is confronted against the original SPS strategy with RRI $= T_{GenCamMin} = 100$ ms; the latter is a convenient setting, as it guarantees that CAMs gain access to the channel without generating any latency reselections (minimum RRI strategy). Adhering to the standard, our SPS implementation randomly chooses the actual C_{resel} value in the $[5, 15]$ range. In accordance with [59], we set $P = 0$. According to it, vehicles select new transmission resources every time the reselection counter expires.

To quantify how effective the KNN choice is within the ML domain, the Ground Truth (GT) benchmark is considered: this benchmark exploits the *a priori* knowledge of the

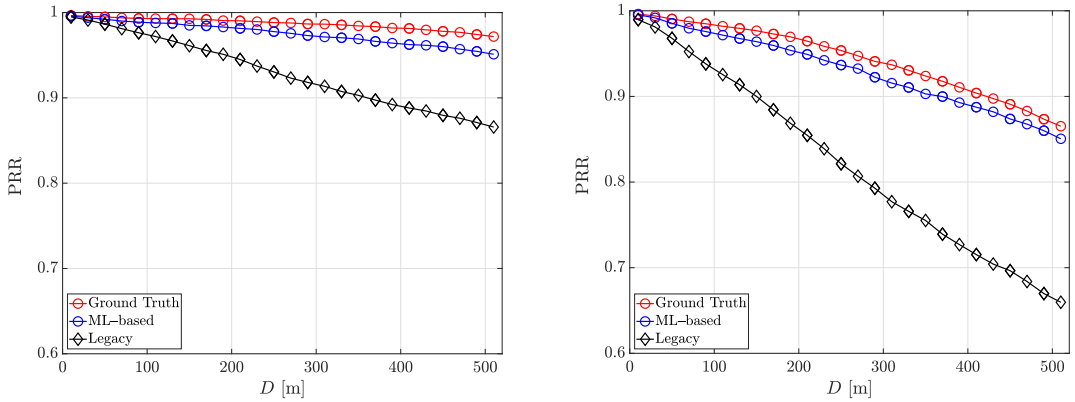

(a) CAM size $X = 190$ bytes.

(b) CAM size $X = 470$ bytes.

Figure 3.8: PRR as a function of the Tx-Rx distance D, suburban scenario.

CAM sequences to assign radio resources and place reservations that perfectly match the actual CAM sequences.

Figs. 3.8(a) and (b) report the PRR curves for the original SPS strategy with RRI = 100 ms, the curves obtained when the KNN-look ahead proposal is adopted, as well as the PRR values corresponding to the ideal GT benchmark. When the CAM size is 190 bytes, Fig. 3.8(a) indicates that our proposal guarantees an attractive improvement, and Fig. 3.8(b) reveals that the gain becomes significant when a larger size ($X = 470$ bytes) is considered, that is, when the load on the radio channel increases. Both figures also reveal that the KNN-look ahead approach attains a performance that is very close to the GT benchmark, i.e., to the ideal performance.

Figs. 3.9(a) and (b) quantify the effects that different mechanisms for the creation of the L_1 list have on the KNN resource selection process, hence on PRR. In both figures, the lowest PRR curve refers to the solution that relies on the original L_1 list; the intermediate curve refers to the alternative where the L_1 creation additionally concentrates on perspective collisions that might occur outside of the selection window W without exploiting the knowledge of C_{resel}; the third, upper curve, to the proposed KNN look-ahead solution. These two figures indicate that the design choices summoned in our proposal consistently lead to the best-performing approach.

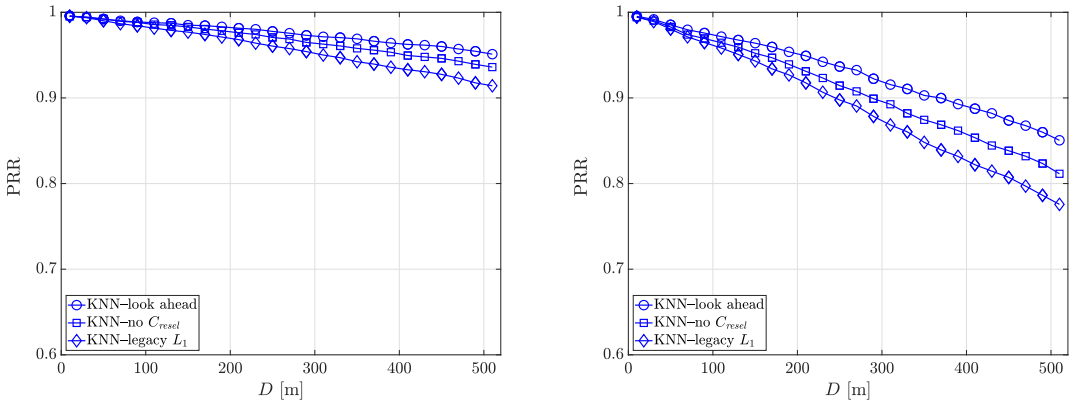

(a) CAM size $X = 190$ bytes.

(b) CAM size $X = 470$ bytes.

Figure 3.9: PRR comparison for different L_1 lists, suburban scenario.

Figs. 3.10(a) and (b) offer a further insight, displaying the CLR curves for the same choice of parameters as in Figs. 3.8(a) and (b). Coherently, the proposed KNN-look ahead strategy displays the lowest CLR values. These figures additionally reveal that the CLR values of the GT benchmark are not zero for all D values. The existence of a CLR "floor" is justified by observing that, even if every vehicle were able to perfectly forecast its CAM transmission requirements over time and select resources accordingly, its selection could nevertheless coincide with the choice performed by other vehicles due to the overlap between selection windows. As highlighted in Subsection 2.4.1, this phenomenon is intrinsic to the distributed nature of the channel access mechanism and cannot be further

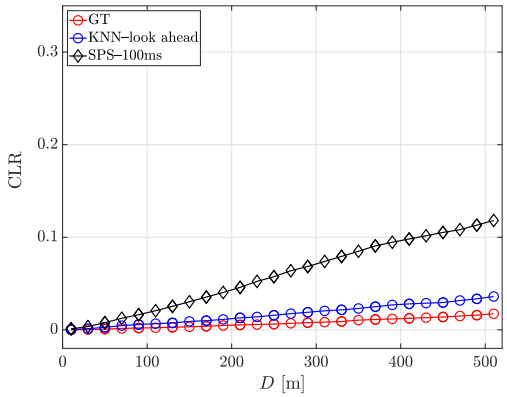

(a) CAM size $X = 190$ bytes.

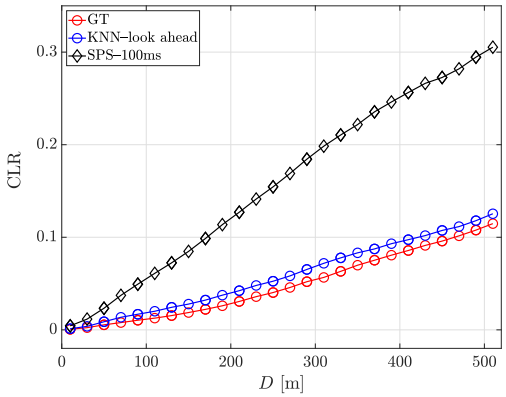

(b) CAM size $X = 470$ bytes.

Figure 3.10: CLR as a function of D, suburban scenario.

Table 3.1: Suburban scenario: URR, LRR and CRR metrics.

	URR	LRR	CRR
SPS, RRI = 100 ms	0.61	0.0	0.18
KNN-look ahead	0.12	0.10	0.22
Ground Truth	0.0	0.0	0.34

reduced, unless a total redesign of the radio access technique is undertaken.

The effectiveness of the KNN-look ahead approach is further evidenced by the values provided in Table 3.1, where the LRR and the URR of the proposed solution are compared against the values of these ratios for the original SPS strategy with RRI = 100 ms and for the GT benchmark.

The table shows that SPS with RRI = 100 ms guarantees no latency reselections, as it respects the most stringent delay requirement, but the URR climbs to 0.61. At the other end of the scale, the GT benchmark perfectly eliminates unutilized reservations and reselections. The proposed solution lies in between, being able to significantly reduce the URR from 0.61 to 0.12. However, this improvement is achieved at the expense of a non-zero fraction of latency reselections. It is worth noting that, as long as they are not prevalent, latency reselections do not have the same negative impact on communication reliability as unutilized reservations. For the sake of completeness, the CRR is also reported in the last column of the table: as expected, its value increases for the proposed solution and even more for the GT benchmark, as reselections become more frequent to track T_{CAM} variability.

Next, Fig. 3.11 reports the Probability Mass Function (PMF) of the T_{CAM} values observed in the suburban scenario. It is interesting to note that the PMF mainly concentrates around two values, 200 ms and 300 ms. As they are not integer multiples, the SPS strategy with RRI set equal to 100 ms is not very effective in detecting potential collisions. This explains why we observed fairly low PRR values for it.

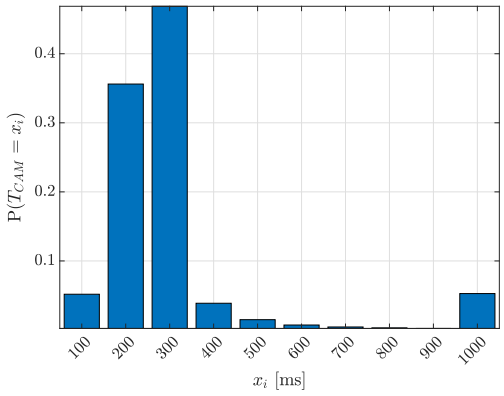

Figure 3.11: T_{CAM} PMF, suburban scenario.

Given the *a posteriori* knowledge that the PMF reported in Fig. 3.11 provides, Fig. 3.12 shows the PRR attained by the legacy SPS strategy when its reservation periodicity RRI is set so as to match the first or the second most frequently observed T_{CAM} value; that is, RRI = 300 ms and RRI = 200 ms. Such PRR values are further confronted against the baseline performance provided by SPS with RRI = 100 ms, and against the performance that the proposed KNN-based look ahead solution attains. Tuning the reservation periodicity improves the PRR of the legacy SPS strategy: unfortunately, the most proper RRI selection would be possible only if the T_{CAM} PMF were *a priori* known. Instead, KNN – or any alternative ML choice – does not necessitate such knowledge and yet, provides far higher PRR levels.

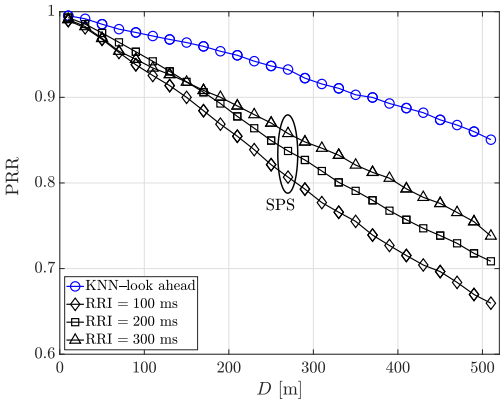

Figure 3.12: PRR comparison of the proposed KNN-look ahead approach against SPS with three different RRI settings, suburban scenario.

To further complete the assessment picture, Table 3.2 reports the CBR levels observed in the suburban scenario. The CBR of the generic vehicle has been computed every 0.2 seconds, the values have been time-averaged over the central portion of the simulation time and finally averaged over all vehicles. The RSSI threshold to discriminate between a busy and an idle channel is set to a value 0.5 dB greater than the receiver sensitivity level, therefore to -89.9 dBm. The CBR values reported in Table 3.2 reveal the magnitude

of the channel load increase due to a larger packet size. Moreover, the CBR is not only useful for assessing the amount of traffic insisting on the communication channel. Given a specific setting, the CBR also reflects the effectiveness of the adopted access strategy: a more accurate scheduling mechanism maximizes the use of the available transmission resources, resulting in larger CBR values. This is the case encountered here, where the KNN-look ahead approach achieves higher CBR values than the SPS strategy with 100 ms.

Table 3.2: Suburban scenario: CBR values.

	$X = 190$	$X = 470$
SPS, RRI = 100 ms	0.24	0.4
KNN-look ahead	0.26	0.46
Ground Truth	0.27	0.47

Figs. 3.13(a) and (b) show the PIR CDF when $D = 520$ m, for the GT benchmark, for the proposed KNN-look ahead solution, and for the legacy SPS. When $X = 470$ bytes, Fig. 3.13(b) indicates that the probability of observing PIR values lower than 200 ms and 300 ms, the two most frequent T_{CAM} values, is 0.74 and 0.90 for the KNN-look ahead solution. This is an improvement with respect to the values of the original SPS strategy, equal to 0.70 and 0.84, respectively. Moreover, the discrete nature of the T_{CAM} values in the $[100, 1000]$ ms range reflects in the step behavior of the PIR CDF.

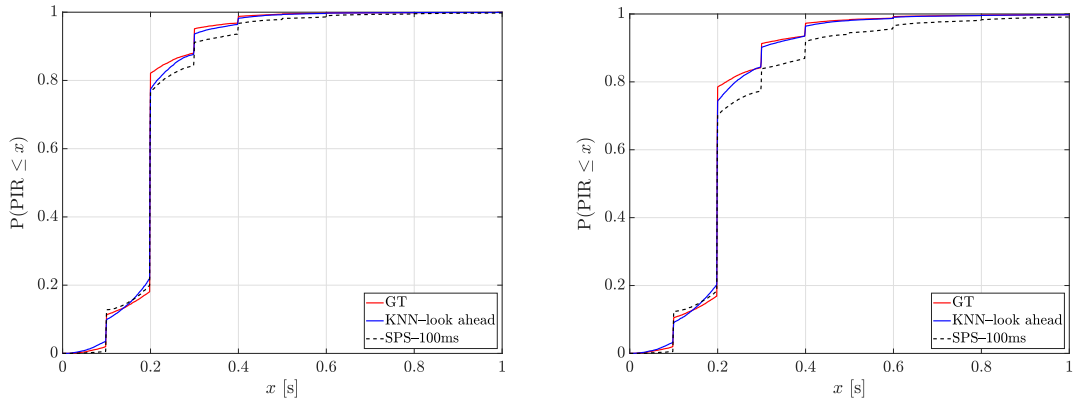

(a) CAM size $X = 190$ bytes.

(b) CAM size $X = 470$ bytes.

Figure 3.13: PIR CDF, suburban scenario.

3.5.2.2 Highway setting

We also considered a second setting, termed *highway*, represented by a 5 km-long highway trunk, where six 4-meter wide lanes are deployed. Adhering to the specifications in [5], the vehicles' speed is 70 km/h and the vehicular density is 120 vehicles/km. For these numerical choices, Fig. 3.14 compares the PRR of the proposed KNN-look ahead solution to the PRR of the SPS strategy with RRI = 100 ms, and to the GT upper bound, for the

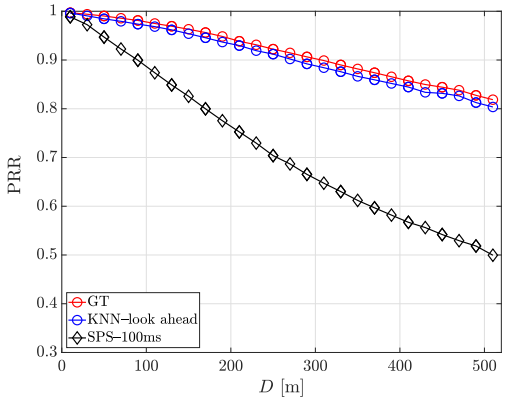

Figure 3.14: PRR as a function of D, highway scenario, CAM size $X = 470$ bytes.

most demanding CAM size $X = 470$ bytes. The figure shows that the KNN-look ahead approach leads to a remarkable improvement in the PRR performance with respect to the original SPS scheme, achieving PRR levels very close to the GT benchmark. It is however known that SUMO reveals some limits in the highway set-up: the constant speed and the nearly straight vehicular trajectories lead to an almost constant CAM inter-arrival time, $T_{CAM} = 300$ ms. The same behavior was observed when the vehicular speed varies within the $[70, 140]$ km/h range: here too, T_{CAM} is nearly constant and equal to 200 ms. We have overcome this simulation hurdle by employing one of the empirical models for the generation of CAM messages that were proposed in [47]. These models are derived from real-world traces of CAM traffic collected on a highway trunk [14], for different implementations of the ETSI algorithm by two Original Equipment Manufacturers (OEMs), Volkswagen and Renault. They consist of m-th order Markov sources that model: (i) CAM size and T_{CAM} variability; (ii) CAM size variability only; (iii) T_{CAM} variability only. We

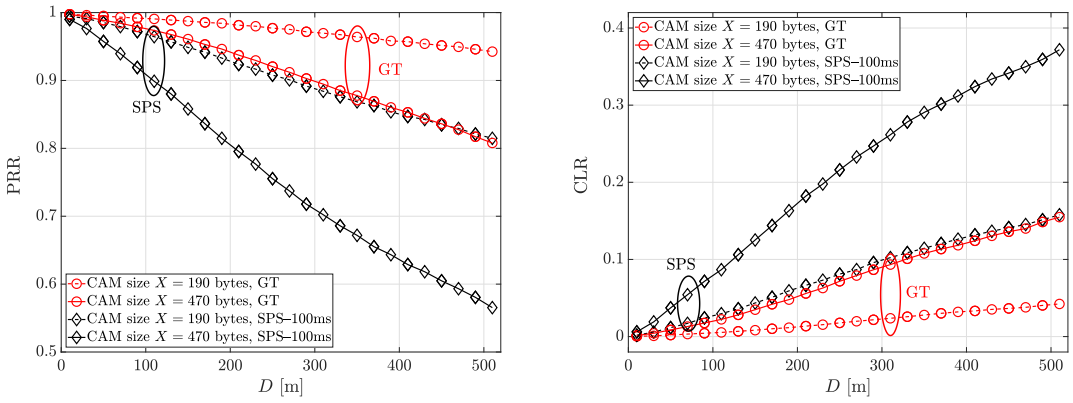

(a) PRR. (b) CLR.

Figure 3.15: PRR and CLR as a function of D, highway scenario, CAM trace Markov model.

adopted the model that seizes CAM temporal variability, drawn from the Volkswagen CAM traces, setting $m = 5$. For this model, the average T_{CAM} value is 330 ms, close to

Table 3.3: Highway scenario: CBR values.

	$X = 190$	$X = 470$
SPS, RRI = 100 ms	0.34	0.49
Ground Truth	0.39	0.61

the constant T_{CAM} value characterizing the SUMO implementation at 70 km/h constant speed. With the help of this analytical tool, we associated every vehicle with a specific CAM trace. Unfortunately, such empirical models have no notion of vehicle dynamics, so they do not provide the input features the KNN algorithm requires. Nonetheless, the reproduction of highway CAM traces allows to determine the GT performance, and therefore to assess the maximum improvement that ML can potentially achieve. In this respect, Fig. 3.15(a) concentrates on the PRR performance considering two different packet sizes, $X = 190$ bytes and $X = 470$ bytes. Adopting the same choice of colors and markers of Figs. 3.8(a)-(b), the black curves correspond to the original SPS implementation with RRI = 100 ms, whereas the red curves refer to the GT benchmark, identifying the PRR upper bound. The significant improvement achieved by the GT solution with respect to the legacy SPS strategy is evident and becomes remarkable when $X = 470$ is considered. The original SPS performance drops below 0.6 when $D \geq 450$ m, whereas the GT sets at PRR = 0.85. Fig. 3.15(b) is the counterpart of Fig. 3.15(a) on the (CLR, D) plane. This figure further highlights the enhanced collision-avoidance capability of the ML-based strategy with respect to the standard-compliant solution, that increases for increasing packet sizes.

Its superiority is substantiated by the CBR values reported in Table 3.3. The first column of the Table refers to $X = 190$ bytes: the CBR increases from 0.34 to 0.39 when moving from the SPS strategy with RRI = 100 ms to the GT benchmark. Likewise, the CBR rises from 0.49 to 0.61 in the second column that refers to $X = 470$ bytes, once more highlighting the significant impact of T_{CAM} predictions on the selection of collision-free resources.

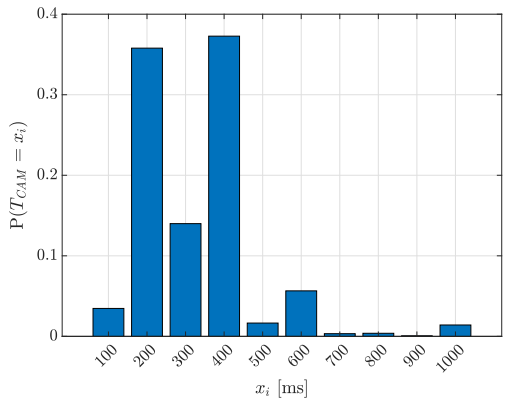

Figure 3.16: T_{CAM} PMF, highway scenario.

The PMF of the T_{CAM} samples generated in the highway scenario is shown in Fig. 3.16. As in the suburban setting, the PMF mainly condenses around two values, 200 ms and 400 ms. Finally, Fig. 3.17 reports the PIR CDF when $D = 520$ m and $X = 470$ bytes. Here too, the GT benchmark provides an upper bound to the PIR achievable performance, highlighting the maximum amount of improvement with respect to the original SPS reservation strategy.

It is worth observing that the implementation of the proposed approach on an actual vehicle is feasible, as the input features that KNN employs can be easily retrieved. The ego-vehicle position can be obtained via the GNSS, its speed can be measured by in-vehicle sensors, and the use of onboard LIDARs and RADARs can offer accurate estimates of the position and speed of the preceding vehicle. The vehicle trajectory prediction is a widely investigated topic in the industrial and the academic world, and algorithms like the one reported in [89] can estimate the ego-vehicle trajectory in an accurate manner.

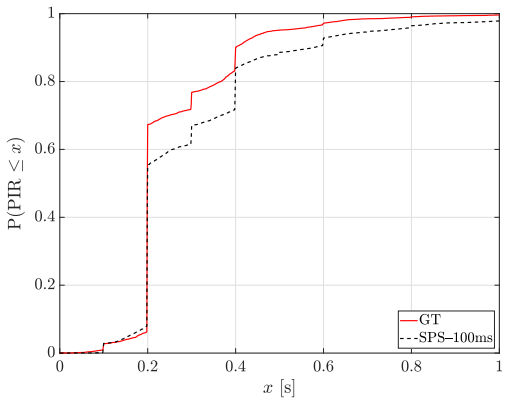

Figure 3.17: PIR CDF, CAM size $X = 470$ bytes, highway scenario.

3.6 Conclusions and Future Work

This Chapter has put forth an ML-based solution termed KNN-look ahead to distribute aperiodic CAMs to vehicles. The proposed approach relies on a limited set of features to forecast future CAM inter-arrival times and allows vehicles to optimize the selection and the reservation of the available radio resources. As regards the introduction of ML, we showed that a simple technique such as KNN is able to significantly mitigate the impact of size reselections, latency reselections, and unutilized reservations on the SPS strategy operation. Simulation results have demonstrated that the proposed ML-based solution achieves an excellent performance, consistently outperforming the legacy LTE-V2X Mode 4 approach.

Although the obtained results concentrate on LTE-V2X Mode 4, they disclose the potential of ML for disseminating realistic, and aperiodic, traffic patterns employing the SPS strategy also in NR-V2X Mode 2. Moreover, we believe that AI techniques can be leveraged by vehicles to adapt the MAC sublayer parameters not only to CAM traffic, but

also to beyond Day 1 traffic sources like Cooperative Perception Messages (CPMs) and Maneuver Coordination Messages (MCMs).

Chapter 4

Awareness Messages by VRUs and Vehicles: Field Tests via LTE-V2X

As of today, significant research efforts have been devoted to the development and enhancement of direct V2V communications to support the deployment of new use cases able to increase traffic safety and efficiency while reducing emissions and fuel consumption.

However, vehicles are not the only type of road occupants. A variety of users swarm on modern roads, experiencing heterogeneous safety levels. Occupants of vehicles travel in a more protected manner; pedestrians, bikes, and more generally, VRUs are by far more exposed to road accidents. Although the most advanced societies witness an increased sensitivity to sustainable mobility and healthy lifestyles, thus encouraging walking and cycling, the attention to VRUs and the efficacy of the solutions to protect them has not proceeded with the same impetus.

To this end, ETSI has recently formalized the dissemination of awareness messages also for VRUs [17]. Like CAMs, VRU Awareness Messages (VAMs) include basic information such as the position, speed, heading, and type of the transmitting VRU. Awareness messages are broadcasted by VRUs to grant vehicles more time to react to unexpected situations, thus representing the fundamental elements to build safety-related applications.

4.1 Research Motivation and Main Contribution

This Chapter presents the findings of a measurement campaign that involved bicycles, e-scooters, motorbikes, and cars, conducted to provide a broad view of the messages generated by different road user categories, assuming they are all connected and therefore able to transmit and receive. LTE-V2X SL modules for direct communications, requiring no support from the cellular infrastructure, were employed. The ETSI algorithms formalized in [16] and [17] were implemented in the firmware of the boards, to broadcast CAMs and VAMs with the proper timing.

First, VRUs in an urban environment were considered: the PMF of the time between the generation of consecutive VAMs was experimentally determined, and the various causes

that trigger the transmission of such messages were identified for bicycles and e-scooters. The investigation was performed employing the VAM generation rules recommended by ETSI in [17]. The analysis revealed that many VAMs were originated under non-relevant conditions; this urged for the proposal of an alternative. The intent was to avoid generating too many VAMs, that excessively clutter the radio channel, without missing relevant information about the VRU movements. For motorbikes, the same characterization was offered about CAMs, determining the temporal features of the messages broadcasted by this category of road users.

Next, field tests were performed for cars traveling in urban, suburban, and highway scenarios, and here too, the generation times of CAMs were critically analyzed. Lastly, the PRR of a car receiving VAMs broadcasted by a bicycle was evaluated in an urban area, as well as the PRR experienced by V2V communications in all of the above mentioned settings.

The most relevant contributions this Chapter provides can be summarized as follows:

- for the first time, awareness messages generated by VRUs (bicycle, e-scooter, motorbike) were analyzed in an urban environment, highlighting their peculiarities as opposed to CAMs broadcasted by vehicles;
- an adjustment to reduce the unnecessarily high generation rate of VAMs was proposed and proved successful;
- the PRR of VAMs received by a vehicle was determined, employing LTE-V2X prototype boards;
- for CAMs, the different causes of CAMs issued by vehicles were disclosed in the urban, suburban, and highway environments;
- the PRR of CAMs received by cars with LTE-V2X boards was evaluated in the above scenarios, revealing more pessimistic results than in controlled field tests.

The rest of the Chapter is organized as follows: Section 4.2 positions the current work in literature. Section 4.3 introduces the definition of VRU and illustrates the rules to follow when VRUs generate awareness messages. The commonalities and differences from CAMs are also highlighted. Section 4.4 details the measurement campaign that broadcasted and collected VAMs and CAMs in various settings, and critically analyzes the data provided by the field tests. Section 4.5 draws the conclusions.

Differently with respect to the simulation studies presented in Chapter 2, the present study undertakes an experimental approach, which examines the transmission of awareness messages via LTE-V2X SL. The investigation zooms into the generation process of VAMs and CAMs, and dissects their causes, demonstrating that they largely vary, depending on the type of road user (bicycle, e-scooter, motorbike, vehicle) and the examined scenario. Furthermore, the work evaluates the PRR performance of LTE-V2X SL when employed for direct communications between a bicycle transmitting VAMs and a car in

the urban setting, as well as its performance when broadcasting CAMs via direct V2V communications in urban, suburban, and highway settings. Finally, the entire dataset collected during the field tests this thesis is made available to the scientific community; unlike car manufacturers, which reluctantly provide their measurements and only disclose their final outcomes, we offer the scientific community both.

4.2 Related Work

Some of the papers that recently focused on VRUs are [90] through [91]. The authors of [90] offered a map-based and data-driven analysis of accident black spots, identified the circumstances that led to accidents, and listed the most effective countermeasures. Similarly, [92], [93] and [94] dealt with the motion prediction of VRUs; while [92] offered a survey on the main techniques available for this task, [93] and [94] focused on artificial intelligence techniques to predict trajectories and detect VRUs, respectively. The study in [95] presented a methodology that enables the inclusion of connected VRUs in the VEINS simulator. Considering contributions based on real tests, on one hand, [96] analyzed the specific problem of a heavy goods vehicle equipped with ultrasonic sensors to measure the motion of a cyclist driving close to it; on the other hand, [97] proposed a warning alarm system based on the standard sensors available in the VRUs drivers' smartphones. Differently from the current study, all of these papers had no notion of VAMs. In [98], it was demonstrated that VAM adoption enables the early triggering of the Forward Collision Warning (FCW) system in the vehicle, and a simulation study was developed to quantify its beneficial effects in a pre-built Euro NCAP scenario. The focus was on the FCW improvement, whereas VAM details and dynamics were neglected. Cooperative perception was explored in [99] and [100], which investigated how to improve the safety of connected and automated vehicles, increasing their awareness of VRUs. The former work considered vehicles equipped with IEEE 802.11p-based devices, communicating with intelligent roadside units. VRUs and the ITS-G5 standard were examined in [91] too; by simulation, the paper found that not-connected vehicles could only perceive a small number of relevant persons, while cooperative perception was able to significantly increase awareness and reduce the detection delay of nearby pedestrians. These works considered neither VRUs able to issue VAMs, nor direct communications among vehicles and VRUs. To the authors' knowledge, the generation of VAMs and their broadcasting, therefore, appear as vastly uncharted territory.

Among CAM studies, [101] explored the practical limits of cooperative awareness in vehicular communications. Relying upon field tests, this work gave a significant contribution; the PRR was computed employing ITS-G5 devices, considering both V2V and Vehicle-to-Infrastructure (V2I) transmissions. Yet, due to the time of its writing, the LTE-V2X SL technology was not considered. More recently, [102] performed an experimental comparison between ITS-G5 and LTE-V2X. In this case, an assessment was performed in terms of PRR and latency, but the focus was exclusively on highway communications. Moreover,

in [101] and [102] CAMs were periodically issued, which is not what the ETSI standard dictates [16]. The study published in [14] highlighted the very diverse and non-persistent nature of the CAM messages of two different car manufacturers, both in size and transmit rate. Unlike the current investigation, [14] did not explore the causes behind CAM generation or made the collected traces publicly available. Furthermore, the breath of the data collection was very modest, being confined to a few tens of minutes. An interesting paper on benchmark testing of V2X technology is [103], which provided an analysis based on cable and field tests, the latter being performed in a confined context, i.e., a field track. LOS and NLOS range tests were accomplished; the coexistence with 802.11p and Wi-Fi networks was also tested. The aim of [103] was to fairly compare 802.11p and LTE-V2X in a controlled environment. On the contrary, the present study intends to experiment with LTE-V2X SL in a real setting and to put this technology at the service of VRUs.

4.3 Vulnerable Road Users

4.3.1 VRU definition, features, and requirements

Although the most diffused road occupants are vehicles, a growing interest is currently being reserved for those road users that either do not use a mechanical device for their trips or utilize alternative means of transportation, less invasive than cars. Such "road inhabitants" are classified as VRUs by the ETSI standard document [104], which the current subsection mostly draws from. The following categories are identified:

- pedestrians;
- road workers;
- wheelchair users and prams;
- skaters;
- skateboards and segways;
- bikes and e-bikes;
- scooters;
- motorcycles;
- animals such as dogs, horses, and wild animals which present a safety risk to other road users.

In [104], VRUs are further categorized in three groups, each with its traits: VRU profile 1 mainly refers to pedestrians, whose behavior is often unpredictable and whose speed range is limited; VRU profile 2 includes light vehicles, that may be equipped with an electric engine: they move at a relatively low speed and their behavior can be more

easily predicted than for pedestrians, yet it is still subject to random movements; VRU profile 3 includes motorbikes, whose speed is similar to cars and that exactly like vehicles can send CAMs when properly equipped. Every profile exhibits its own challenges: VRUs belonging to profiles 2 and 3 are often hard to perceive from other vehicles; moreover, profiles 1 and 2 VRUs sometimes travel in groups and do not follow road rules.

In the upcoming years, VRU safety will go through the widespread adoption of wireless connectivity: when connected, the VRU can either have only a transmitter that broadcasts awareness messages or a receiver for messages from other road users and roadside units, or both.

The transmission of VRU standard messages, the so-called VAMs, is needed in the majority of cases. To reduce the amount of load that VAMs would generate on the radio channel at specific occurrences, for instance, at a pedestrian crossing in a metropolitan area or at a location where a major event takes place, VRUs are grouped in a new logical entity, termed cluster [105]. Users belonging to the same cluster may belong to the same profile or exhibit different profiles; however, they move with similar speed or direction and within a bounding box. Importantly, other road users are informed of the cluster presence through a single VAM, rather than by a VAM for each VRU. It is up to the cluster head to issue the VAM and indicate whether the cluster is homogeneous or heterogeneous; the latter difference is particularly useful, as it offers information about the trajectory and behavior prediction once the cluster disperses.

4.3.2 VRU Awareness Messages

The current Section explains the generation rules and format of VAMs. It then shortly elaborates on CAMs, which are much more frequently encountered in literature than VAMs, citing the related works.

4.3.2.1 VAM Generation

The ETSI standard document [17] dictates that the minimum time elapsed between consecutive VAM generation events has to be equal to or larger than T_{GenVam}, where T_{GenVam} falls in the interval $[T_{GenVamMin}, T_{GenVamMax}] = [100, 5000]$ ms. Furthermore, an individual VAM is generated whenever:

(i) the time elapsed since the last VAM transmission exceeds $T_{GenVamMax}$;

(ii) the absolute distance between the current VRU position and the one included in the previous VAM exceeds $\Delta_d = 4$ m;

(iii) the absolute difference between the current VRU heading and the heading included in the previous VAM exceeds $\Delta_h = 4°$;

(iv) the absolute difference between the VRU current speed and the speed included in the previous VAM exceeds $\Delta_s = 0.5$ m/s;

(v) the difference between the currently estimated trajectory interception probability with vehicles(s) or other VRU(s) and the trajectory interception probability with vehicle(s) or other VRU(s) lastly reported in a VAM exceeds a threshold;

(vi) the VRU decides to join a cluster;

(vii) the VRU has determined that one or more new vehicles or other VRUs are coming closer than the minimum safe lateral distance or the minimum safe longitudinal distance or the minimum safe vertical distance.

The conditions for triggering the VAM generation shall be checked every $T_{CheckVamGen}$, where $T_{CheckVamGen} \leq T_{GenVamMin}$ [17].

If the VRU is exclusively equipped with a transmitter, conditions (v) and (vii) do not apply. Last, condition (vi) comes into play when a cluster of VRUs is considered. Finally, redundancy mitigation techniques are enforced, to reduce the communication load in settings where VRUs can be very numerous, respecting the constraint that the VRU safety and awareness is not diminished. Here, the underlying ideas are: (i) to decrease VAM frequency by a given factor when a peer system (e.g., another VRU) has just issued a VAM while being in very similar conditions as the reference VRU in terms of location, speed, and orientation; (ii) to skip VAM transmission in a non-drivable or low-risk geographical area.

4.3.2.2 VAM Format

The VAM format includes the ITS Protocol Data Unit (PDU) header, generation time, and multiple data containers, as illustrated in Fig. 4.1. The ITS PDU header adheres to a common format that is employed for application and facility layer messages: it contains data elements such as the identifier of the transmitting VRU and the message ID. The message ID is used to indicate the message type, and it is set equal to 14 for VAMs.

After the header and the timestamp, the Basic Container and the High-Frequency Container represent the first two components of the payload. According to the standard, these are mandatory containers that shall be included in every transmitted VAM. In detail, the Basic Container provides fundamental information about the VRU generating the awareness message, such as the VRU type (e.g., pedestrian, cyclist) and its latest geographical position. Next to the Basic Container, the High-Frequency Container has been designed to provide fast-changing information, such as the latest heading, speed, and acceleration values that the originating VRU has retrieved. Moreover, this container may include optional information used only by specific VRU profiles. In addition to the mandatory containers, a VAM shall also include the Low-Frequency Container when the time elapsed since the last VAM transmission exceeds 2000 ms. As opposed to the High-Frequency Container, the Low-Frequency Container holds slow-changing information, like the VRU profile and size. In this case, too, the container's content comprises both mandatory and optional data elements.

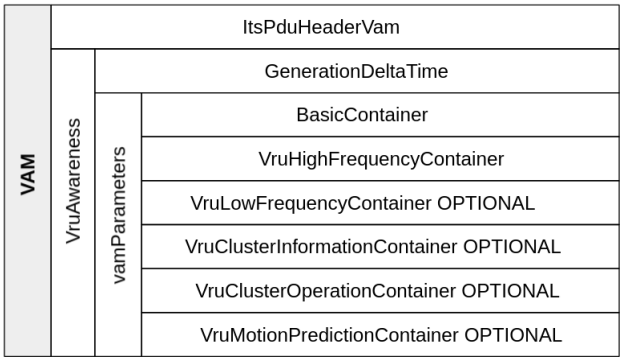

Figure 4.1: VAM structure.

VAMs may also include three new containers, that are new with respect to CAMs: the Cluster Information Container, the Cluster Operation Container, and the Motion Prediction Container. When two or more users operate within a cluster, the cluster leader shall include the Cluster Information Container in its transmitted VAMs. Specifically, this container is used to disseminate basic information about the VRU cluster, such as the cluster ID, the cluster bounding box, and the profiles of the VRUs forming the cluster. On the other hand, the Cluster Operation Container can be included by either the cluster leader or a cluster member. The cluster leader exploits this container to announce the cluster breakup and indicate the conditions that triggered such an event. Conversely, cluster members include the Cluster Operation Container in the transmitted VAMs when they perform cluster joining or leaving procedures. Last, the Motion Prediction Container is an optional payload component that provides information about the past and future movements of the originating VRU. This container can include up to 40 data entries that report the most recent activity of the VRU or predictions about its future trajectory.

As the VAM payload contains both mandatory and optional information, its size can significantly change depending on the considered VRU profile and the specific implementation. For instance, the size of the VRU Motion Prediction Container can vary from 45 to 1180 bytes depending on the number of data entries. The ETSI standard indicates that the VAM size lies between 40 and 1260 bytes if no or all optional containers are included, respectively. Furthermore, VAMs are tentatively harmonized with the Personal Safety Message (PSM) defined by SAE in [106].

Differently from VAMs, CAMs are often mentioned in scientific and technical literature, albeit the case of messages issued by motorbikes is never taken into consideration. The role of CAMs in enhancing road safety has been explored in various settings, e.g., see [107] through [46].

Their generation rules and format are given in the ETSI standard document [16] and appear in several papers, [47] [108]. For a CAM, the value of the message ID that indicates the message type is 2. Moreover, the CAM generation times are ruled by the conditions (i)-(iv) previously listed, whereas conditions (v) through (vii) apply to VAMs only. Furthermore, $T_{GenCamMax}$ replaces $T_{GenVamMax}$, and $T_{GenCamMax} = 1$ s. It is worth

pointing out that the Δ_d, Δ_h, and Δ_s thresholds are fixed in the case of CAMs, and can be adjusted when VAMs dissemination is considered. The CAM size ranges from 45 to 400 bytes, as highlighted in [77], and a further increase of some hundred bytes is due to the introduction of security content like signatures and certificates [14]. Overall, CAMs are less flexible in length and content than VAMs.

Instead of tackling any arbitrary CAM and VAM implementation profile, this study concentrates on their temporal patterns and considers messages that carry basic information, namely: the sequence number, the spatial coordinates of the originating road user, its speed, and heading at the time of the message generation.

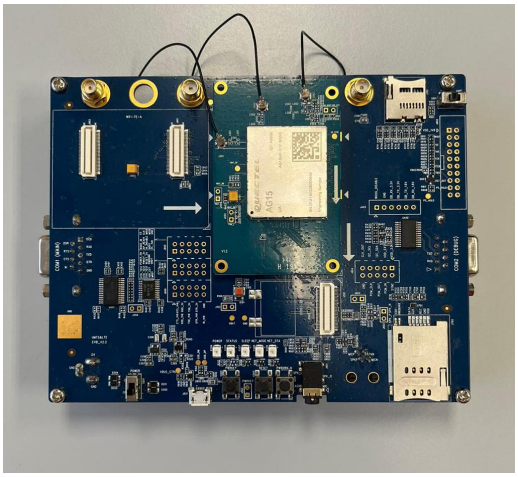

Figure 4.2: Quectel evaluation board used for the field measurements.

4.4 Field Tests

4.4.1 LTE-V2X Boards

The experimental results presented in this Chapter have been obtained employing the AG15 module, manufactured by Quectel [109]. The module is built around the Qualcomm 9150 C-V2X chipset, which is compliant with LTE-V2X Sidelink specifications and respects the quality constraints for automotive products set by the IATF 16949 standard. It supports Mode 4 vehicular communications in the n47 band (5855-5925 MHz). In accordance with the standard, every broadcasted VAM (and equivalently every CAM) is encapsulated in a TB transmitted in the PSSCH. The module also features a GNSS multi-constellation receiver that allows for positioning, speed, and heading measurements. The supported constellations are GPS, GLONASS, BeiDou, and Galileo. The AG15 module is mounted on the Quectel evaluation board shown in Fig. 4.2.

The board is equipped with 2 omnidirectional antennas, characterized by a 5 dBi gain. When installed on the rooftop of cars, the antennas were connected to the board using 50 Ω RG316 coaxial cables; when used on bicycles, they were directly inserted on the board, via SMA coaxial connectors. The transmit power and the receiver sensitivity were set to

23 dBm and -93 dBm, respectively. Because of the attenuation due to the coaxial cables and connectors, the emitted power was always lower than 28 dBm. Therefore, it fulfilled the requirement set by the EU regulation [110], which states that the maximum RF output power shall not exceed 33 dBm Equivalent Isotropically Radiated Power (EIRP) in the n47 frequency band. The adopted MCS was 16QAM-0.5. This is a typical setting for LTE-V2X SL communications, also employed in simulative works (see Section 2.5.1).

We custom-developed the firmware to generate CAMs and VAMs, implementing the ETSI standard rules which determine the generation times of the awareness messages [17]; in this respect, the availability of the high sensitivity GNSS receiver was fundamental. As indicated at the end of Subsection 4.3.2, the messages carried the essential pieces of information needed to trigger their creation. Their size was constant and equal to 300 bytes.

4.4.2 Bicycle and E-scooter VAMs

In what follows, the VAM generation times of two different VRU types, a bicycle and an e-scooter, are analyzed. They were obtained by a VRU equipped with the Quectel board that recorded its GNSS output with a 10 Hz sampling frequency. Based on such output, the VAM traces were *a posteriori* generated, through custom Python scripts that allow for the setting of different threshold values in conditions (ii) through (iv). In accordance with the GNSS sampling period, the VAM triggering conditions of Subsection 4.3.2 were checked with a 10 Hz frequency, i.e., every $T_{CheckVamGen} = 100$ ms. Note that no receiving board is needed when measuring the time interval between consecutive CAMs and VAMs. We only need to monitor the dynamics of the vehicle/VRU that generates the messages. No other traffic participants broadcasting awareness messages were considered, so the case of a VRU-Tx only equipment, according to ETSI [104], was recreated. It followed that the VAM triggering conditions (v) and (vii) introduced in Subsection 4.3.2.1 could not be encountered. The VAM triggering condition (vi) could not be met either, as all the experiments involved a single VRU. Therefore, VAMs were caused either by a timeout event, i.e., by the occurrence of condition (i), or by events related to the VRU dynamics, i.e., by conditions (ii) through (iv).

The bicycle VAMs are analyzed first. They refer to an urban environment, namely, a residential suburb and the downtown area of the city of Modena, Italy. The rides took place in a propagation environment featuring no harsh urban canyons. For this setting, the

Table 4.1: Number of collected CAMs and VAMs.

Vehicle type	Urban	Suburban	Highway
Bicycle	44543	-	-
E-scooter	14437	-	-
Motorbike	40308	-	-
Car	18218	31222	16106

Table 4.2: Field tests region, duration and traveled distance.

latitude in [44.575593,44.643645]		longitude in [10.859186,10.989253]	
Vehicle type	Urban	Suburban	Highway
Bicycle	3 hours, 42 km	-	-
E-scooter	2 hours, 26 km	-	-
Motorbike	4.5 hours, 128 km	-	-
Car	2 hours, 82 km	2 hours, 131 km	1 hour, 127 km

dataset consists of 44543 VAMs, as reported in Table 4.1. The duration of the experiments and the traveled distance are reported in Table 4.2, which also details the latitude and longitude intervals of the area where all the tests took place.

Fig. 4.3 shows an exemplary portion of a bicycle VAM trace: T_{VAM}, the time between the generation of consecutive VAMs, is reported as a function of the VAM index. The trace was obtained by setting the thresholds that trigger the VAM to $\Delta_d = 4$ m, $\Delta_h = 4°$, and $\Delta_s = 0.5$ m/s [17]. As expected, the minimum observed value of T_{VAM} is 100 ms and is coincident with $T_{CheckVamGen}$. The figure reveals that the VAM trace is fairly irregular: T_{VAM} rarely stays constant between consecutive samples unless specific conditions are met. An example of this is provided in the last portion of the trace, where the sample index falls in the 610-620 interval, and T_{VAM} stabilizes at 700-800 ms. Such behavior is typically observed when the bicycle follows a straight trajectory at an approximately constant speed.

Next, Fig. 4.4(a) reports the percentages of VAMs classified on the basis of the triggering conditions (i) through (iv) at the end of Subsection 4.3.2, for the same choice of thresholds as in Fig. 4.3. An additional type, labeled as "mixed", is considered; it indicates that a VAM was generated when more than one triggering condition was satisfied, an event not excluded by the standard. Interestingly, VAMs due to heading variations are the most frequent, accounting for 74% of the total. On the other hand, only 6.8%, 11%, 8%, and 0.2% of VAMs are generated by a speed, distance, mixed, and timeout trigger, respectively. The 98% of VAMs classified as "mixed" also satisfies the heading condition.

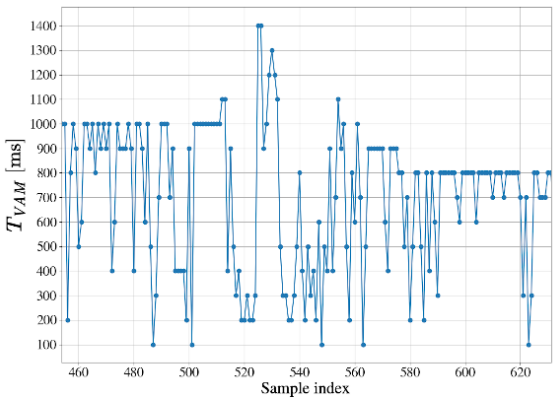

Figure 4.3: Bicycle: VAM trace.

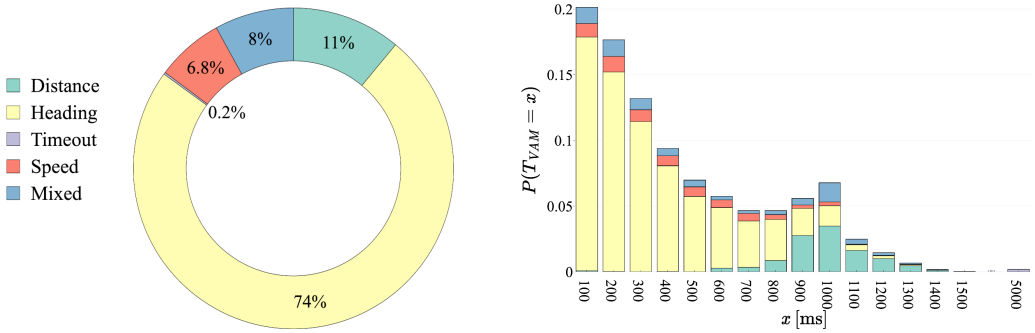

(a) Triggers distribution. (b) PMF.

Figure 4.4: Bicycle: VAM triggers and T_{VAM} PMF for $\Delta_h = 4°$.

Fig. 4.4(b) completes the data analysis, displaying the PMF of T_{VAM}, the time between two consecutive VAMs. In this figure, the contributions to the PMF due to a variation in the VRU speed, heading, and distance, as well as to a timeout, are separately displayed. The figure reveals that VAMs triggered by heading variations greater than $\Delta_h = 4°$ are mainly observed for low T_{VAM} values, whereas VAMs due to distance variations greater than 4 m become predominant when T_{VAM} falls within the 800-1300 ms range. As a matter of fact, the average and maximum bicycle speeds that were recorded are 12.9 and 27.7 km/h; under the assumption of uniform linear motion, the first value allows to conclude that it takes more than 1100 ms to cover a distance coincident with $\Delta_d = 4$ m. Fig. 4.4(b) also shows that the probability of observing T_{VAM} values in the [1400,5000] ms range is practically zero. $P(T_{VAM} = 5000$ ms) is insignificant too: this value is exclusively due to timeouts that occur if the bicycle stands still, as might be the case for a red traffic light. Such a situation was seldom encountered during an average bicycle ride, which explains the negligible fraction of VAMs generated by this condition.

The second type of VRU that was analyzed is the e-scooter. As for the bicycle, we begin our analysis considering the VAMs generated when $\Delta_d = 4$ m, $\Delta_h = 4°$, and $\Delta_s = 0.5$ m/s. The collected dataset led to 14437 VAMs and refers to routes leading to the downtown area of the city of Modena and to the city center. Fig. 4.5(a) shows that

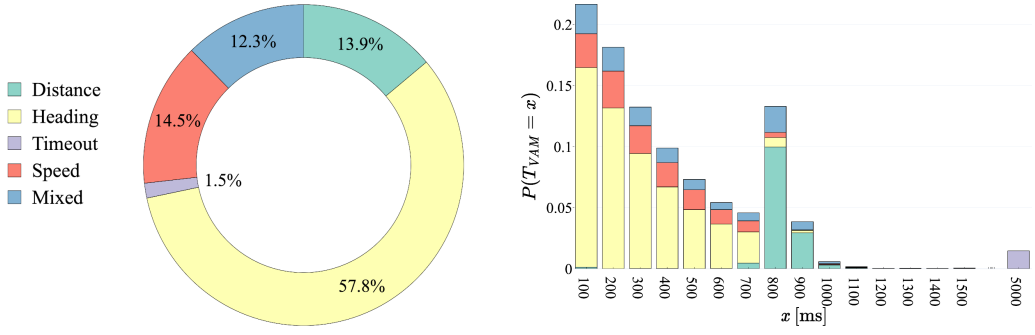

(a) Triggers distribution. (b) PMF.

Figure 4.5: E-scooter: VAM triggers and T_{VAM} PMF for $\Delta_h = 4°$.

57.8%, 13.9%, and 14.5% of the VAMs were generated by a heading, distance, and speed variation, respectively. This figure also indicates that mixed triggers represent 12.3% of the total and that the percentage of timeout triggers is negligible. A significant percentage of mixed triggers, 94.7%, satisfies the heading condition. The PMF of T_{VAM} is reported in Fig. 4.5(b). As in Fig. 4.4(b), heading-triggered VAMs are mostly observed for small T_{VAM} values. Differently from the case of the bicycle, VAMs caused by distance variations are mostly concentrated at $T_{VAM} = 800$ and 900 ms, rather than being distributed over a wider range. This has to be ascribed to the speed-limited e-scooter employed for the tests, whose speed was more stable than that of the bicycle. In this setting, the recorded average speed was 14 km/h; the maximum was 25.1 km/h. Both figures indicate that the VAMs due to heading variations are predominant, similarly to what was observed for the bicycle.

The significant number of VAMs due to heading variations that we observed raised the following questions: does the $\Delta_h = 4°$ threshold represent the most proper choice for the examined VRUs? Or should a more carefully selected value be employed?

To help identify the correct answers, the next set of figures considers an exemplary

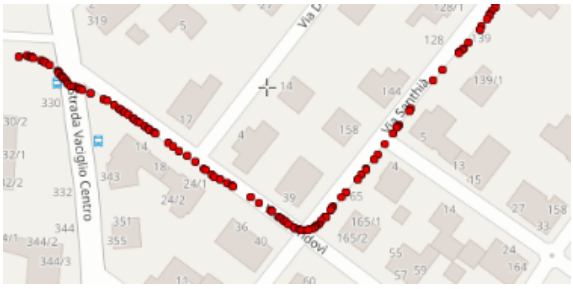

(a) $\Delta_h = 4°$.

(b) $\Delta_h = 7°$.

(c) $\Delta_h = 10°$.

Figure 4.6: Bicycle: GNSS coordinates when heading-triggered VAMs are generated.

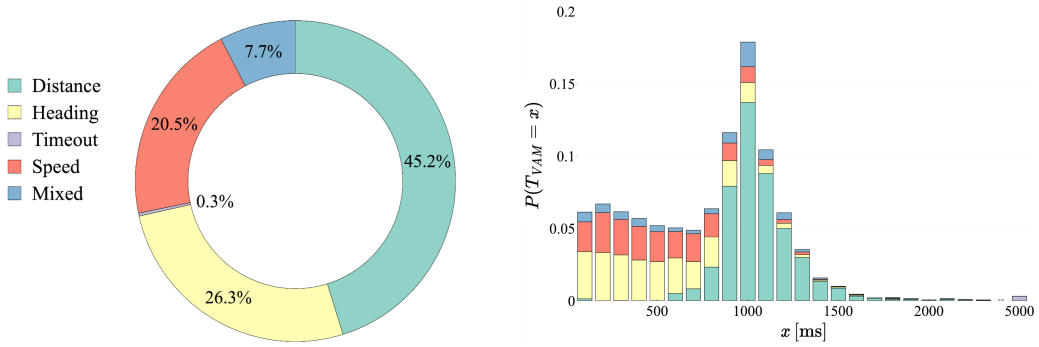

(a) Triggers distribution. (b) PMF.

Figure 4.7: Bicycle: VAM triggers and T_{VAM} PMF for $\Delta_h = 10°$.

portion of the bicycle rides, and plots on the map the GNSS coordinates where VAMs generated by heading variations were broadcasted. Fig. 4.6(a) refers to $\Delta_h = 4°$ and shows that a remarkable fraction of such VAMs is transmitted under non-relevant circumstances, i.e., on a straight road segment. These VAMs are useless to other road traffic participants. Figs. 4.6(b) and (c) correspond to a progressively increased value of Δ_h, $\Delta_h = 7°$ and $10°$, respectively. The figures show that an increase of the Δ_h threshold reduces the number of heading-generated VAMs, without diminishing their effectiveness for Δ_h up to $10°$. In this case, Fig. 4.6(c) indicates that heading variations generate new VAMs only when the VRU performs a relevant steering movement.

Figs. 4.7(a)-(b) analyze the VAMs generated when $\Delta_h = 10°$, the values of Δ_d and Δ_s being unmodified. Fig. 4.7(a) quantifies the impact of the $\Delta_h = 10°$ setting on the percentage of each trigger type. With respect to Fig. 4.4(a), the percentage of VAMs triggered by heading variations decreases from 74% to 26.3%, whereas the percentage of VAMs triggered by speed and distance increases from 6.8% to 20.5%, and from 11% to 45.2%, respectively. Fig. 4.7(b) completes the analysis, reporting the PMF of the T_{VAM}, where a significant reduction in heading-generated VAMs is observed, with respect to Fig. 4.4(a). As a result, the contribution of speed and distance triggers is more evident, and

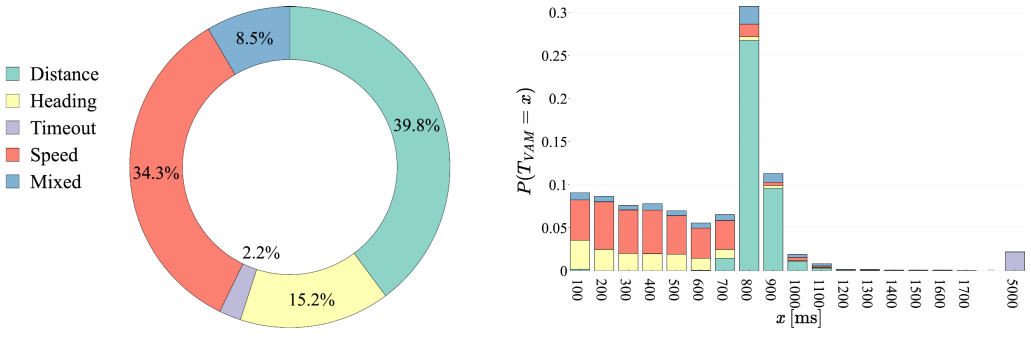

(a) Triggers distribution. (b) PMF.

Figure 4.8: E-scooter: VAM triggers and T_{VAM} PMF for $\Delta_h = 10°$.

the PMF profile is shifted towards larger T_{VAM} values.

The impact of a larger Δ_h threshold on the e-scooter VAMs is also analyzed in Figs. 4.8(a)-(b). Fig. 4.8(a) quantifies the impact of the $\Delta_h = 10°$ setting on the percentages of VAMs due to each trigger type; the reduction of heading-generated VAMs is evident from the comparison with Fig. 4.5(a). Fig. 4.8(b) reports the PMF of T_{VAM}, revealing that the probabilities associated with the most populated T_{VAM} bins (at 800 and 900 ms) are markedly higher.

4.4.3 Vehicle and Motorcycle CAMs

The test results about the generation times of CAMs by cars and motorbikes are reported next, for $\Delta_d = 4$ m, $\Delta_h = 4°$, and $\Delta_s = 0.5$ m/s. Although motorcycles are classified as VRUs [104], they transmit CAMs instead of VAMs [104]. For cars, the measurements were performed in three different settings, namely, an urban environment, a suburban area, and a highway. For motorbikes, the urban scenario only was examined. Cars and motorbikes traveled among other vehicles, experiencing normal daytime traffic conditions. The macroscopic description of the propagation environments that vehicles experienced is the following: (i) the urban route mainly consists of large avenues surrounded by trees and residential buildings; (ii) the suburban route is made of road portions that run close to constructions, while other segments run through open areas; (iii) the highway route crosses a mixed area, with cultivated fields and also industrial plants.

The car repeatedly traveled along the urban path reported in red in Fig. 4.9, where varied traffic conditions were encountered, traffic lights forced vehicles to stop and roundabouts slowed them down. Fig. 4.10(a) reports the percentage of the different triggering conditions and Fig. 4.10(b) the PMF of T_{CAM}, the time interval between two consecutive

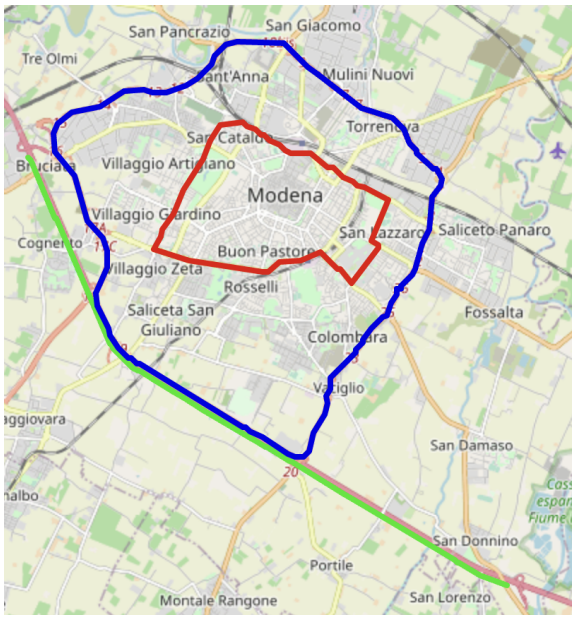

Figure 4.9: The examined scenarios.

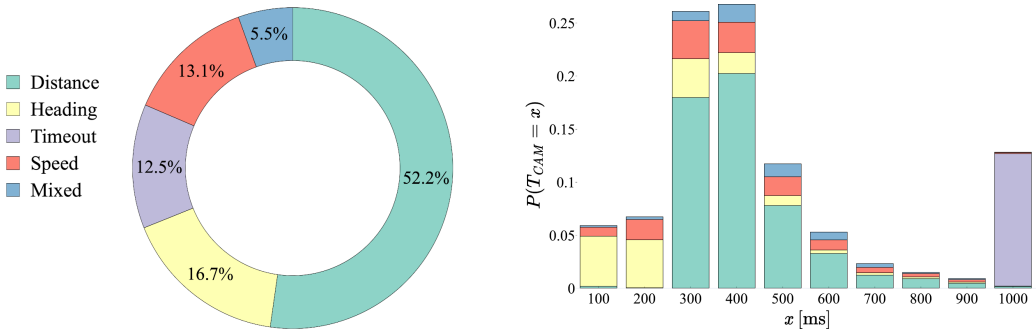

(a) Triggers distribution. (b) PMF.

Figure 4.10: Car: CAM triggers and T_{CAM} PMF, *urban* scenario.

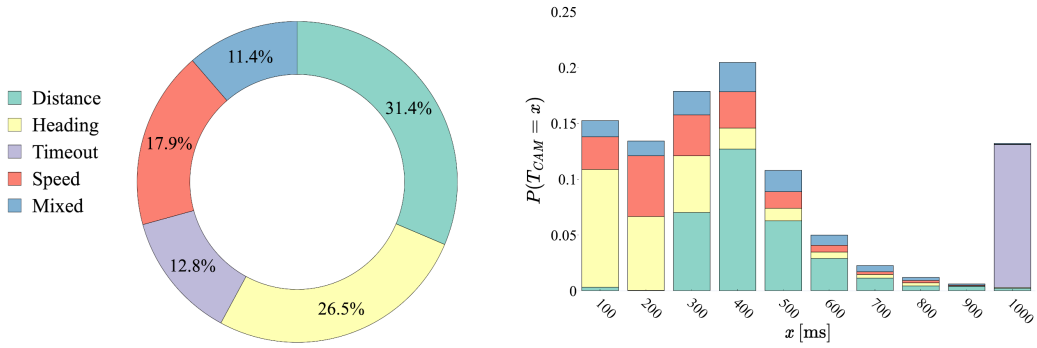

(a) Triggers distribution. (b) PMF.

Figure 4.11: Motorcycle: CAM triggers and T_{CAM} PMF, *urban* scenario.

CAMs. Fig. 4.10(a) indicates that distance variations greater than 4 m are the main cause behind CAM generation in this setting; they trigger more than half of the recorded messages. In order of decreasing weight, distance is followed by heading, speed, timeout, and mixed. Except for the latter, the figure indicates that the other conditions have a comparable weight. With respect to the VAMs generated by bicycles and e-scooters with $\Delta_h = 4°$, the heading variation weighs much less, although, as visible from Fig. 4.10(b), it still accounts for the majority of triggering conditions when $T_{CAM} = 100$ ms and 200 ms. Moreover, CAMs due to distance variations greater than 4 m are generated with a lower T_{CAM}. This is due to the higher speeds of the vehicle: we recorded 29.7 km/h for the average and 71.9 km/h for the maximum. The other relevant difference with respect to VAMs pertains to the timeouts, nearly absent in all results presented in Section 4.4.2. Here, they account for 12.5% of the triggering conditions, the totality of them corresponding to a T_{CAM} of 1000 ms.

With reference to the same urban scenario, Fig. 4.11 reports the percentage of the triggering conditions and the PMF of T_{CAM} for a two-wheel motorbike. Comparing the percentages of the motorbike triggering conditions in Fig. 4.11(a) with those of the car in Fig. 4.10(a), it is concluded that they are more diversified. Distance variations account for approximately 30% of the total, heading variations are significantly more frequent, and the

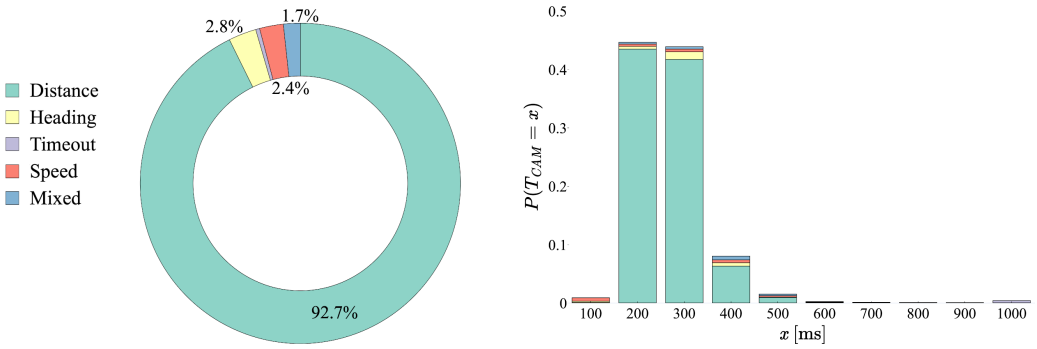

(a) Triggers distribution. (b) PMF.

Figure 4.12: Car: CAM triggers and T_{CAM} PMF, *suburban* scenario.

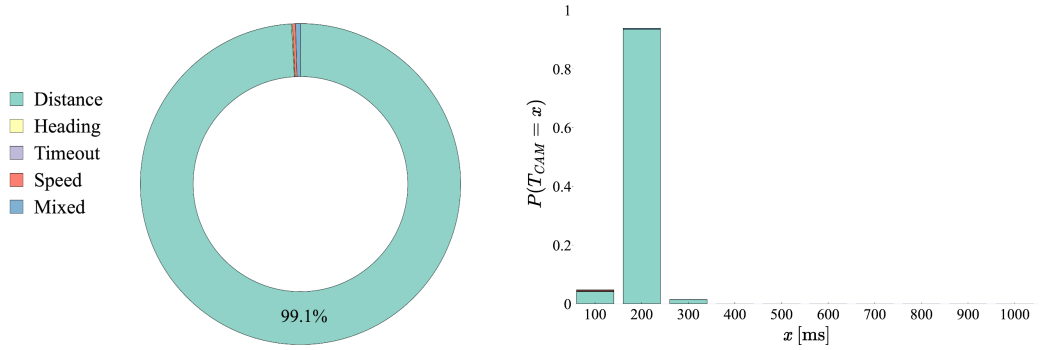

(a) Triggers distribution. (b) PMF.

Figure 4.13: Car: CAM triggers and T_{CAM} PMF, *highway* scenario.

occurrences of the speed and mixed conditions are increased. Only the timeout percentage stays approximately the same. The most significant phenomenon is that CAMs generated by heading variations become more frequent. Indeed, a motorbike follows a more irregular trajectory than a car, due to its smaller size and agility in slipping through traffic. Fig. 4.10(b) also reveals that heading variations are mainly responsible for generating CAMs whose T_{CAM} is lower than 300 ms; above this value, CAMs caused by distance variations, i.e., higher speeds, prevail.

The measurement campaign was completed by tests performed having the cars also traveling on suburban roads and on a highway. In the first case, the route followed by the car is reported in blue in Fig. 4.9, and refers to the freeway belt surrounding Modena. The corresponding percentages of triggering conditions and the PMF of T_{CAM} are provided in Figs. 4.12(a) and 4.12(b), respectively. Here, distance accounts for the vast majority of the generated CAMs. Timeouts are negligible, while heading, velocity, and mixed altogether cause $\sim 7\%$ of the CAMs. The vast majority of the CAMs is triggered in the time interval $[200, 300]$ ms, which corresponds to the $[48, 144]$ km/h range of speed values.

As regards the highway, the drive route is evidenced in green in Fig. 4.9. Fig. 4.13(a) reports the percentages of the triggering conditions, and Fig. 4.13(b) the PMF of T_{CAM}. Distance variations cause more than 99% of the CAMs. Due to the higher speeds, the

(a) Tx bicycle: setup.

(b) Tx and Rx car: rooftop antennas.

(c) Rx car: in-vehicle setup.

Figure 4.14: LTE-V2X SL PRR: employed equipment.

vast majority of CAMs are triggered because of distance with a time interval between two subsequent messages of 200 ms. This is because $T_{CAM} = 300$ ms corresponds to velocities in the $[48, 72)$ km/h range, $T_{CAM} = 200$ ms to velocities in the $[72, 144)$ km/h range, and $T_{CAM} = 100$ ms to speeds higher than 144 km/h [14]. As a matter of fact, the car speed exhibited an average of 116.7 km/h and a maximum of 148.5 km/h. Because of varying traffic conditions, there were cases in which vehicles were traveling below 72 km/h, which in turn account for $T_{CAM} = 300$ ms. Overall, the PMF resulted nearly unimodal. Although some of the results referring to vehicles could be drawn by simulation, we observe that a simulator cannot perfectly replicate the vehicle dynamics. This is particularly true in environments where the driver's behavior heavily depends on the road topology and the dynamics of other road users. To conclude, only the measurement campaign offers a realistic picture of what to expect.

The number of collected messages, the test duration, and the traveled distance are summarized in Tables 4.1 and 4.2 for all the examined road users.

4.4.4 LTE-V2X SL: PRR Analysis

Next, we determined the PRR achieved for two distinct cases: (i) a bicycle broadcasting VAMs and a car receiving them; (ii) a car broadcasting CAMs, while a second car was driving behind and acting as the receiver. Note that message losses were exclusively due to varying propagation conditions, given no additional traffic was injected on the radio channel to cause congestion and harm the VAM and CAM dissemination.

The cars and the bicycle were equipped with the Quectel AG15 LTE-V2X SL evaluation boards. Fig. 4.14 displays the equipment deployed in the two cases. When positioned on the bicycle, the antennas were directly connected to the board and their height was 0.80 m; on the car rooftop, the antennas' height was 1.45 m. Three evaluation boards were utilized for this set of measurements. One board was placed in the front vehicle or in the bicycle basket, generating the awareness messages. Two boards were placed on the receiving vehicle, one acted as the receiver, and another recorded the GNSS outputs every 50 ms. This allowed for the continuous tracking of the receiving car position, which would otherwise occur only at the successful reception of a message. By doing so, we

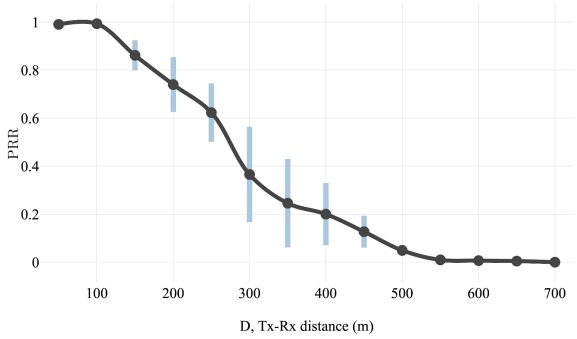

Figure 4.15: Bicycle-to-Vehicle communications: PRR.

were able to reliably compute the PRR as a function of the distance between the two road occupants. All other relevant settings, such as transmit power, MCS, code rate, and receiver sensitivity are listed in Subsection 4.4.1.

When considering VAMs transmitted from the bicycle, Fig. 4.15 portrays the VAM PRR as a function of D, the distance between the transmitting bicycle and the receiving car, in the urban environment. The black curve shows the PRR mean values; the box plots, i.e., the vertical, light blue lines, indicate the 95% confidence interval of the measurements. The figure shows that the reliability range, i.e., the distance at which the PRR falls below 0.9 is 125 m, and that the communication range, i.e., the distance beyond which the PRR goes to zero, is 600 m.

Figs. 4.16(a), 4.16(b) and 4.16(c) portray the CAM PRR as a function of D, the distance between the transmitting and the receiving cars, in the urban, suburban and highway scenarios, respectively. In the urban setting, the PRR curve is consistently higher than in the case of VAM transmission; this is explained by a more favorable placement of the antennas on the car, which benefits from a higher elevation from the ground than in the case of the bicycle. The reliability range is 250 m and the communication range is 700 m. The former value is far lower than the one obtained in NLOS condition in [103], where the authors report 875 m at an intersection and 425 m in the case of a strong obstruction. In the suburban setting, packets were delivered over an extended distance, as very few buildings could disrupt the communications. Yet, many vehicles, including trucks, were on the two-lane road of the drive route. Fig. 4.16(b) accordingly shows that the reliability range is 250 m, whereas the communication range extends to 1000 m. Last, Fig. 4.16(c) refers to the highway scenario. On the highway, cars were able to communicate more reliably and over longer distances because there were seldom any obstacles between them. In this case, the reliability range is 700 m; moreover, the communication range is slightly lower than 2000 m. Comparing these outcomes with the findings in [103] that refer to LOS conditions, we observe that in [103] the reliability range was measured at nearly 1200 m. There are fewer discrepancies when our results are compared with those in [102], where the viewgraphs show PRR values around 0.9 for $D = 500$ m.

In parallel, we monitored the delay incurred by packets and verified that it respected the packet delay budget, set to 50 ms. This is a direct consequence of the nearly deter-

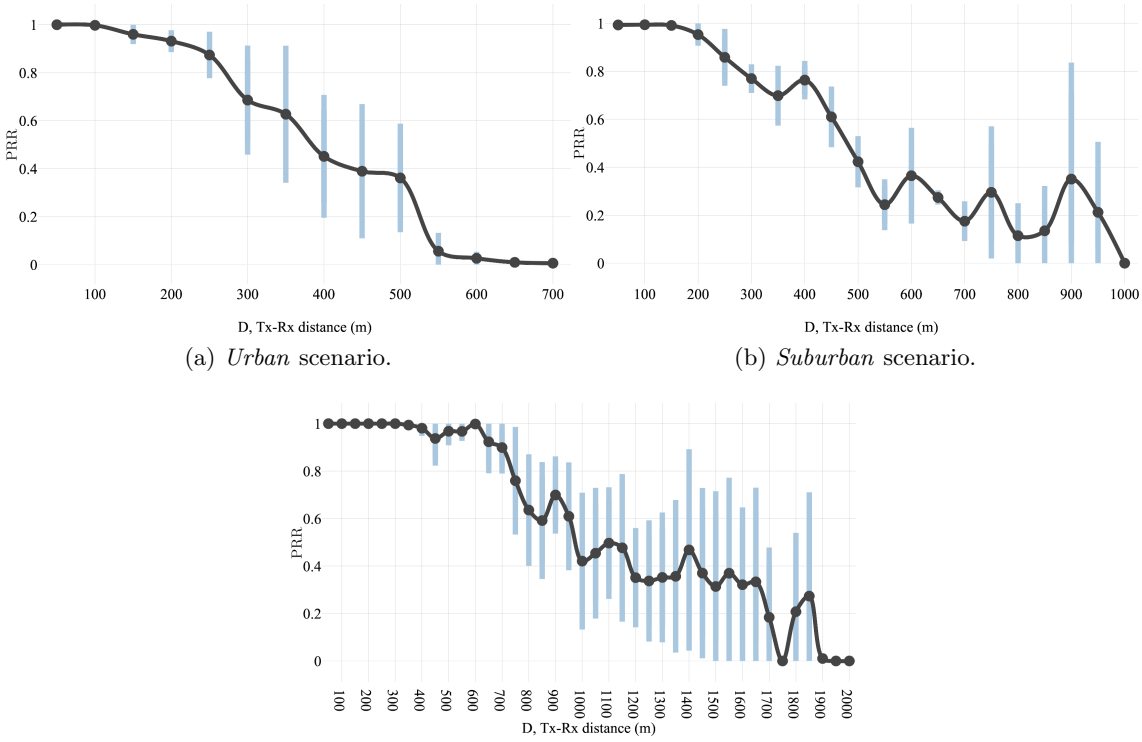

(a) *Urban* scenario.

(b) *Suburban* scenario.

(c) *Highway* scenario.

Figure 4.16: V2V communications: PRR.

ministic channel access mechanism that LTE-V2X Mode 4 adopts.

4.4.5 Discussion

4.4.5.1 VAM and CAM generation times

The tests have revealed that the Δ_h value bikes and e-scooters inherit by the vehicular setting is excessively low. It causes a significant number of VAMs to be generated under non-relevant circumstances, providing neighboring road users no useful clues about the VRU dynamics and wasting precious radio resources. Properly setting Δ_h to the slightly larger value of 10° solves the problem, with no loss of meaningful information about the VRU status. The significant increase in the average T_{VAM} that this modification achieves can be appreciated by comparing the values of Table 4.3. The benefit is relevant, especially in the light of an increase in the number of connected VRUs, and if the operation of each VRU will be subject to spectrum usage constraints [104].

The field tests performed with the motorbike indicate that the number of distance and heading triggers is comparable for CAMs generated by this category of VRUs. It is furthermore observed that motorbikes generate more varied CAM traffic than cars. In the urban setting, the average T_{CAM} is 270 ms, far lower than for cars, whose value is 460 ms, and it reflects the typical hasty nature of motorbike drivers. Table 4.3 provides an overall numerical picture of the T_{VAM} and T_{CAM} values we measured.

Table 4.3: Average T_{VAM} and T_{CAM} values.

Vehicle type	Urban	Suburban	Highway
Bicycle - $\Delta_h = 4°$	435 ms	-	-
Bicycle - $\Delta_h = 10°$	831 ms	-	-
E-scooter - $\Delta_h = 4°$	474 ms	-	-
E-scooter - $\Delta_h = 10°$	702 ms	-	-
Motorbike	270 ms	-	-
Car	460 ms	270 ms	197 ms

The tests performed with cars show that: (i) distance variations are the main CAM triggering cause, followed by speed variations; (ii) the T_{CAM} PMF is strongly multimodal in the urban scenario and reduces to being almost unimodal in the highway scenario. Moreover, the highway PMF largely departs from the one reported in previous studies [14]. We believe that the discrepancy is due to the limited size of the data set employed by the authors of the latter study.

4.4.5.2 PRR of LTE-V2X SL Communications

We assessed the PRR of the LTE-V2X SL technology for direct VRU-to-vehicle and V2V communications. No tests were ever presented in the literature for VRU communications, and our experiments highlighted that in the urban setting, the transmission of VAMs can count on reduced reliability and communication range with respect to CAMs.

For V2V communications, we offered a complete and real picture referring to distinct settings and found more pessimistic range values than the authors of [103]. In this respect, we would like to stress that we considered actual environments, not a test track. Our tests indicate that the CAM reliability and communication range significantly depend on the examined context, and increase when LOS conditions are more frequently encountered. The dataset collected during our field tests referring to VAMs, CAMs, and PRR values is freely accessible on GitHub [111].

4.5 Conclusions and Future Work

This Chapter has presented the results of a vast measurement campaign, conducted to experimentally analyze the temporal patterns of the awareness messages generated by different road occupants. VAMs and CAMs were gathered through field tests performed in urban, suburban, and highway settings. The analysis of the collected data revealed the impact of the various triggering conditions of both message types in the distinct environments where road occupants travel.

Concerning the urban scenario, this study has highlighted the unique characteristics of VAMs issued by bicycles and e-scooters, and has pointed to an amendment of their generation rule that diminishes their frequency without any information loss. This turns out particularly relevant if many VRUs are equipped with a transmitting device and the

risk of creating congestion on the radio channel becomes tangible. In the urban setting, this study has also shown that the VRU category of motorbikes exhibits a distinctive PMF of the time between consecutive CAMs, and that the conditions which trigger its CAM broadcasting are peculiar as well. As regards cars, distance variations are always the main reason behind CAM generation; their CAMs become more frequent and the PMF of T_{CAM} accordingly shifts toward lower values when moving from the busy traffic conditions of the urban environment to the smoother driving scene of the highway.

Besides evaluating the PRR performance of LTE-V2V SL communications in the urban, suburban, and highway settings, this study has also presented a preliminary analysis of the PRR attained by bicycle-to-vehicle communications in the urban scenario. The actual communication range and the reliability range that the VAM and CAM broadcasting accomplishes were identified. Furthermore, the dataset collected during the field tests has been openly released, to ease the work of the community researching on road safety solutions.

To conclude, several research paths are suggested by the measurement campaign, the most immediate being the proposal of alternative algorithms for the generation of VAMs by different categories of VRUs. Namely, bikes and e-scooters may still adhere to ETSI specifications to broadcast VAMs, yet with adaptive thresholds, or may generate VAMs with adaptive frequencies, e.g., dependent on their speed.

Chapter 5

Fundamental Limits on the Performance of PD-NOMA

Scanning the 5G and beyond horizon, wireless connectivity appears as one of the key enabling technologies for future Internet of Things (IoT) and Internet of Vehicles (IoV). According to Cisco [112], the number of connected devices is yet growing at an extraordinary pace and is expected to reach a total of 29.3 billion devices by 2023, with IoT connections accounting for half of the total. At the same time, the advent of V2X has paved the way for the realization of the IoV vision, where an exponentially increasing number of vehicles will be able to share data collected through on-board sensors over the Internet.

Such a massive demand for Internet connectivity, along with the heterogeneous set of performance requirements which characterizes IoT devices [113] and V2X-enabled vehicles [19], transcends the capabilities of 4G and 5G systems. Relying on orthogonal transmission techniques, these systems can only support a limited number of simultaneously active users, and advocate for the design of new connectivity solutions. In this direction, sixth generation (6G) cellular systems will feature novel NGMA schemes able to guarantee the coexistence of traditional UEs and connected vehicles, supporting a wide range of data rate, reliability, and latency requirements. Under the NGMA umbrella, NOMA techniques are expected to play a pivotal role in the support of unprecedented connectivity capabilities [114]. The key idea behind NOMA, i.e., serving multiple users over the same radio spectrum, has been widely investigated over the last years, breeding an abundant body of scientific literature and the proposal of many distinct NOMA approaches. Among them, there appear PD-NOMA, Sparse Code (SC)-NOMA, and Resource Spread (RS)-NOMA, to name a few examples.

5.1 Research Motivation and Main Contribution

In this Chapter, we analytically assess the performance of a Single Input Single Output (SISO) uplink PD-NOMA system when a dynamic-ordered SIC receiver is considered.

PD-NOMA is able to accommodate multiple users (UEs and connected vehicles) on the same frequencies by carefully assigning different power levels to the active devices. A dynamic-ordered SIC receiver adaptively varies the decoding order on the basis of the instantaneously received signals power. The assumption of only two or three superimposed signals usually found in literature is removed to disclose the fundamental limits on the achievable performance of the receiver. Furthermore, the analysis in the presence of Rayleigh fading is extended by the investigation of the combined effects of fading and lognormal shadowing. To characterize system behavior, the outage probabilities $P_{out}^{(j)}$, $j = 1, 2, \ldots, n$, are determined, the generic $P_{out}^{(j)}$ being defined as the probability that the receiver fails to decode the j-th strongest signal and therefore cannot recover the remaining $n - j$ weaker signals. With respect to the state-of-the-art, several novel contributions are offered:

- a general method to analytically evaluate the outage probabilities is provided, based on the unique properties of the joint probability density function of the ordered received powers, as the latter are dependent, non identically distributed random variables. The approach can be profitably employed for any number n of simultaneously received signals;
- when Rayleigh fading is considered, the exact analytical expression of $P_{out}^{(1)}$, the probability that the strongest signal cannot be decoded and that the SIC receiver fails to recover any of the simultaneous signals, is provided for an arbitrary value of n. When the first strongest signal can be decoded, a closed-form approximation of $P_{out}^{(j)}$, $j \geq 2$, is also put forth in order to characterize the performance of the remaining $n - 1$ active users;
- when the signals are affected by Rayleigh-lognormal shadowed fading, an approximation of the outage probability $P_{out}^{(j)}$, $j \geq 1$, is offered, demonstrating that it achieves an excellent accuracy, again for an arbitrary n;
- the proposed approximations show that, j being fixed, $P_{out}^{(j)}$, $j \geq 2$, can be recursively evaluated as a function of the probabilities $P_{out}^{(1)}$ obtained in the presence of $n, n - 1, \ldots, n - j + 1$ simultaneously active users.

Overall, the analysis discloses the limits that the dynamic-ordered SIC receiver faces for an increasing number of superimposed signals, when an uplink NOMA system is considered. Furthermore, it reveals that lognormal shadowing is responsible for a non-negligible performance worsening, with respect to the circumstance where Rayleigh fading only is considered. The deterioration is quantified for different values of σ_L, the standard deviation of the slow, lognormal fading in dB.

The remainder of the Chapter is organized as follows. Section 5.2 explores the available literature on NOMA. Section 5.3 introduces the system model, it illustrates the analysis and it puts forth the approximations to the outage probabilities experienced on the uplink. Section 5.4 specializes the study to the cases of Rayleigh and Rayleigh-lognormal

shadowed fading. Section 5.5 provides several numerical results that validate the approach and Section 5.6 draws the conclusions, discussing the potential of NOMA in V2V communications.

5.2 Related Work

Due to the abundant body of scientific literature about NOMA that has been published over the last years, the risk that every NOMA bibliography incurs is to forget or inadvertently miss some relevant works among the multitude of published papers. For this reason, the surveys reported in [115–118] are mentioned in this Section together with [119] and [120], which provide a comprehensive comparison between NOMA techniques and alternative NGMA schemes such as Multi-User MIMO (MU-MIMO), and Rate-Splitting Multiple Access (RSMA).

Confining the attention to PD-NOMA, simply referred to as NOMA in the rest of this Chapter, its behavior has been assessed in numerous studies. However, the only work that examined the potential of NOMA in the V2X domain is [121], which showed the potential of NOMA in mitigating traffic congestion and reducing latency in NR-V2X downlink communications. The performance of random access uplink NOMA was evaluated in [122] from a system-level perspective, i.e., in terms of throughput and access delay, whereas an analytical framework for the modeling and the analysis of large-scale uplink and downlink NOMA systems has been proposed in [123, 124]. The studies in [125, 126] highlighted the strengths and the limitations which characterize a typical SIC-based decoding process in uplink and downlink NOMA communications, and put forth the design of a new hybrid SIC receiver. Moreover, it is worth recalling the contributions in [127] and [128]. In the former, the performance of uplink NOMA, paired with a dedicated power control scheme, was analyzed in terms of outage probability and achievable data rate. In the latter, the authors proposed a novel uplink NOMA scheme able to achieve higher spectral efficiency and lower receiver complexity with respect to conventional techniques. Stemming from [127], [129] focused on the optimization of the power allocation strategy. Differently from the contributions mentioned so far, where a fixed decoding order SIC receiver was considered, the authors of [130] and [131] analyzed NOMA systems employing dynamic-ordered SIC receivers. These works determined closed-form expressions of the outage probability for the case of three users, without however providing a systematic analysis. In [132], the Signal-to-Interference Ratio (SIR) coverage probability of uplink NOMA was evaluated, comparing the performance of two SIC receivers; namely, the dynamic-ordered SIC receiver was confronted against a SIC receiver that ranks and decodes the users' signals on the basis of their mean received powers.

It is worth pointing out that all the previous investigations were performed in the exclusive presence of Rayleigh fading, an assumption that greatly simplifies the study. In this regard, the work in [133] introduced the hypothesis of generalized fading channels encompassing statistics such as Rayleigh, Rice, and Nakagami, and considered a fixed

decoding order SIC. Due to the complex nature of the analysis, the authors exclusively considered the circumstance of two superimposed users for mathematical tractability. In [134], the channel model was the generalized α-μ fading and, also in this work, the SIC decoding order was fixed; here too, the analysis was limited to the circumstance of two or at the most three simultaneously active users.

5.3 Performance Analysis

5.3.1 System Model and Performance Metric Evaluation

The current work focuses on the uplink communications in a cellular system. Power-domain NOMA is considered and the reference scenario features n UEs (or vehicles) that transmit to a central-located base station on the same radio spectrum. Let $p_{t,i}$ denote the transmission power of the i-th UE and h_i the envelope of the channel between such UE and the base station. Let

$$X_i = p_{t,i}|h_i|^2 \qquad i = 1, 2, \ldots, n \tag{5.1}$$

denote the instantaneously received power at the base station from the i-th UE. Further assume that every UE experiences independent channel conditions while transmitting to the base station; it follows that h_i and h_j are independent random variables, $\forall\, i$ and $\forall\, j$, $i \neq j$, and so are X_i and X_j. Moreover, assume that the signal recovery is performed using a dynamic-ordered SIC receiver. This choice implies that: (i) the instantaneously received signal powers from the UEs are first sorted in descending order at the base station; (ii) the receiver attempts to decode the signals in accordance to the same sequence.

Indicate by $\mathcal{S}_\mathcal{N}$ the set of the $n!$ permutations of $N = \{1, 2, \ldots, n\}$ and by $Z = \{z_1, z_2, \ldots, z_n\}$, $Z \in \mathcal{S}_\mathcal{N}$, the permutation that corresponds to the descending order of the instantaneously received powers. It follows that

$$p_{t,z_1}|h_{z_1}|^2 \geq p_{t,z_2}|h_{z_2}|^2 \geq \ldots \geq p_{t,z_n}|h_{z_n}|^2 \,. \tag{5.2}$$

where p_{t,z_1} is the transmitted power of the UE that exhibits the highest received power, p_{t,z_2} the transmitted power of the UE that exhibits the second highest received power, and so forth.

Next, introduce the random variables

$$X_{(i)} = p_{t,z_i}|h_{z_i}|^2 \qquad i = 1, 2, \ldots, n \tag{5.3}$$

and observe that the $X_{(i)}$s are no longer independent. Rather, owing to (5.2) they constitute an order statistics; for the notation employed, $X_{(1)}$ is the largest order statistic, $X_{(n)}$ is the smallest.

The receiver first attempts to decode the strongest signal. If the decoding process is

successful, the receiver removes the first strongest signal and then proceeds to decode the second strongest. For the base station to decode the message from the j-th strongest user $\mathrm{UE}_{(j)}$, the $j-1$ received signals with the strongest power have to be successfully recovered and removed first.

Recalling Shannon's capacity theorem, the achievable data rate of $\mathrm{UE}_{(j)}$ is

$$R_{(j)} = \log_2 \left(1 + \frac{X_{(j)}}{\sum_{i=j+1}^{n} X_{(i)} + \sigma^2} \right) \quad \text{bits/s/Hz}\,, \tag{5.4}$$

for $j = 1, 2, \ldots, n-1$, where σ^2 is the noise power, and

$$R_{(n)} = \log_2 \left(1 + \frac{X_{(n)}}{\sigma^2} \right) \quad \text{bits/s/Hz} \tag{5.5}$$

for the last user $\mathrm{UE}_{(n)}$, whose received power is the weakest.

Denote by $\hat{R}_{(j)}$ the target data rate of $\mathrm{UE}_{(j)}$ and define the outage probability $P_{out}^{(j)}$, $j = 1, 2, \ldots, n$, as the probability that the SIC receiver can successfully recover the first strongest signal, the second strongest, up to the $j-1$, but it fails to decode the j-th strongest and all the subsequent signals. Analytically,

$$P_{out}^{(j)} = 1 - P\{R_{(j)} \geq \hat{R}_{(j)}\}\,. \tag{5.6}$$

If we indicate by $\mathcal{E}_k$, $k = 1, 2, \ldots, n-1$, the random event identified by the condition

$$\frac{X_{(k)}}{\sum_{i=k+1}^{n} X_{(i)} + \sigma^2} \geq \hat{\gamma}_k \tag{5.7}$$

where

$$\hat{\gamma}_k = 2^{\hat{R}_{(k)}} - 1, \quad k = 1, 2, \ldots, n\,, \tag{5.8}$$

and by $\mathcal{E}_n$ the event in which the condition

$$\frac{X_{(n)}}{\sigma^2} \geq \hat{\gamma}_n \tag{5.9}$$

holds, then (5.6) is equivalently re-written as

$$P_{out}^{(j)} = 1 - P\left\{ \cap_{k=1}^{j} \mathcal{E}_k \right\}\,, \quad j = 1, 2, \ldots, n\,, \tag{5.10}$$

where it is observed that the random events $\mathcal{E}_1, \mathcal{E}_2, \ldots, \mathcal{E}_n$ are statistically dependent.

Indicate by $f_{joint_n}(x_{(1)}, x_{(2)}, \ldots, x_{(n)})$ the joint probability density function (pdf) of the ordered set of random variables $X_{(i)}$, $i = 1, 2, \ldots, n$, and by $\mathcal{D}_j$ the region of the

$X_{(1)}, X_{(2)}, \ldots, X_{(n)}$ space identified by the conditions:

$$\mathcal{D}_j = \begin{cases} X_{(1)} \geq \hat{\gamma_1} \cdot (\sum_{i=2}^{n} X_{(i)} + \sigma^2) \\ X_{(2)} \geq \hat{\gamma_2} \cdot (\sum_{i=3}^{n} X_{(i)} + \sigma^2) \\ \vdots \\ X_{(j)} \geq \hat{\gamma_j} \cdot (\sum_{i=j+1}^{n} X_{(i)} + \sigma^2) \\ 0 \leq X_{(n)} \leq X_{(n-1)} \leq \cdots \leq X_{(2)} \leq X_{(1)} \end{cases} . \tag{5.11}$$

It follows that $P_{out}^{(j)}$, $j = 1, 2, \ldots, n$, is determined solving the integral

$$P_{out}^{(j)} = 1 - \int \cdots \int_{\mathcal{D}_j} f_{joint_n}(x_{(1)}, x_{(2)}, \ldots, x_{(n-1)}, x_{(n)}) \times dx_{(n)} dx_{(n-1)} \ldots dx_{(2)} dx_{(1)} \,. \tag{5.12}$$

Among the different outage probabilities $P_{out}^{(j)}$, observe that $P_{out}^{(1)}$ deserves a special place, as it coincides with the probability that the strongest signal cannot be correctly recovered; in this circumstance, not even one, out of the n simultaneous transmissions, can be successfully decoded and power-domain NOMA fails. Indeed, the inequality $X_{(1)} \geq \hat{\gamma_1} \cdot (\sum_{i=2}^{n} X_{(i)} + \sigma^2)$ identifying the outage domain $\mathcal{D}_1$ represents the necessary condition for the SIC decoding process to begin. Equivalently stated, $P_{out}^{(1)}$ gives the probability that the adoption of power-based NOMA turns out detrimental, as not even the best signal is correctly decoded.

When evaluating $P_{out}^{(j)}$, the first non-trivial problem at hand is to determine $f_{joint_n}(x_{(1)}, x_{(2)}, \ldots, x_{(n)})$. In this respect, let $f_i(x_i)$ be the pdf of the *unordered* random variable X_i, $i = 1, 2, \ldots, n$, defined in (5.1), whose pdf is available once the pdf of h_i is known, as $p_{t,i}$ is a constant, and define F_n as the following $n \times n$ matrix

$$F_n = \begin{bmatrix} f_1(x_{(1)}) & f_2(x_{(1)}) & \ldots & f_n(x_{(1)}) \\ f_1(x_{(2)}) & f_2(x_{(2)}) & \ldots & f_n(x_{(2)}) \\ \vdots & \vdots & \ddots & \vdots \\ f_1(x_{(n)}) & f_2(x_{(n)}) & \ldots & f_n(x_{(n)}) \end{bmatrix} \tag{5.13}$$

where $f_j(x_{(i)})$ denotes the pdf of the unordered random variable X_j, $j = 1, 2, \ldots, n$, when the function argument is the random sample $x_{(i)}$ of the ordered random variable $X_{(i)}$. For the purpose of what follows, recall that the permanent of a square matrix A, written as $\overset{+}{|}\, A \,\overset{+}{|}$, is defined like the determinant, except that all signs are positive. For an arbitrary n, it can be demonstrated that the joint pdf $f_{joint_n}(x_{(1)}, x_{(2)}, \ldots, x_{(n)})$ of the ordered statistics $X_{(1)}, X_{(2)}, \ldots, X_{(n)}$ is

$$f_{joint_n}(x_{(1)}, x_{(2)}, \ldots, x_{(n)}) = \overset{+}{|}\, F_n \,\overset{+}{|}, \tag{5.14}$$

where F_n is given by (5.13). Last result is substantiated by the reasoning in [135], where the arguments of [136] are extended to prove the formulation in (5.14) with the use of permanents.

At first sight, (5.14) gives the impression that evaluating the integral in (5.12) might be quite cumbersome for an arbitrary value of n. However, the joint pdf obeys a highly peculiar structure, that allows a more convenient rewriting of it in the following terms: let $S_i = \{i_1, i_2, \ldots, i_n\}$ denote the generic permutation of $N = \{1, 2, \ldots, n\}$ in $\mathcal{S}_\mathcal{N}$, where we recall that the latter symbol indicates the set of all possible permutations. It follows that $f_{joint_n}(x_{(1)}, x_{(2)}, \ldots, x_{(n)})$ can be equivalently written as

$$f_{joint_n}(x_{(1)}, x_{(2)}, \ldots, x_{(n)}) = \sum_{S_i \in \mathcal{S}_\mathcal{N}} f_1(x_{(i_1)}) f_2(x_{(i_2)}) \ldots f_n(x_{(i_n)}). \tag{5.15}$$

Last expression highlights that the joint pdf exhibits the presence of $n!$ terms, wherein the permutations of the arguments of the $f_1(\cdot), f_2(\cdot), \ldots, f_n(\cdot)$ pdfs appear. Replacing (5.15) in (5.12) leads to

$$P_{out}^{(j)} = 1 - \int \cdots \int_{\mathcal{D}_j} \sum_{S_i \in \mathcal{S}_\mathcal{N}} f_1(x_{(i_1)}) f_2(x_{(i_2)}) \ldots f_n(x_{(i_n)}) \times \, dx_{(n)} dx_{(n-1)} \ldots dx_{(2)} dx_{(1)} \tag{5.16}$$

and denoting by I_{S_i} the result of the integral

$$I_{S_i} = \int \cdots \int_{\mathcal{D}_j} g_{i_1 i_2 \ldots i_n}(x_{(1)}, x_{(2)}, \ldots, x_{(n)}) \, dx_{(n)} \ldots dx_{(1)} \,, \tag{5.17}$$

where

$$g_{i_1 i_2 \ldots i_n}(x_{(1)}, x_{(2)}, \ldots, x_{(n)}) = f_1(x_{(i_1)}) f_2(x_{(i_2)}) \cdot \ldots \cdot f_n(x_{(i_n)}) \,, \tag{5.18}$$

then $P_{out}^{(j)}$ can be rewritten as:

$$P_{out}^{(j)} = 1 - \sum_{S_i \in \mathcal{S}_\mathcal{N}} I_{S_i} \,, \quad \forall j \,, j = 1, 2, \ldots, n \,. \tag{5.19}$$

Luckily, the random variables $X_1, X_2, \ldots, X_n$ obey the same statistical description, although with different mean values. It follows that it is not necessary to compute every single I_{S_i} term in (5.19). Rather, the *n-th* fold integral in (5.17) has to be solved only once, for a specific S_i. For instance, I_{S_1} can be determined, $S_1 = \{1, 2, \ldots, n\}$. Then, all the remaining I_{S_i} terms are obtained through the proper permutation of the $f_i(\cdot)$'s arguments $x_{(i_k)}$, $k = 1, 2, \ldots, n$ in (5.18). This significantly reduces $P_{out}^{(j)}$ computational complexity in n regardless of the channel envelope statistics, i.e., no matter what pdf the random variables h_i, $i = 1, 2, \ldots, n$, obey to.

Once $P_{out}^{(j)}$ has been obtained, the sum data rate that power-domain NOMA achieves is evaluated as:

$$R_{NOMA} = \sum_{j=1}^{n} \hat{R}_{(j)} \cdot \left(1 - P_{out}^{(j)}\right) \,. \tag{5.20}$$

5.3.2 $P_{out}^{(j)}$ Approximation for $j \geq 2$

The previous development highlighted how to reduce the complexity that hinders behind the exact analytical evaluation of the outage probability $P_{out}^{(j)}$, $j = 1, 2, \ldots, n$. The approach turns out particularly effective when evaluating $P_{out}^{(1)}$. When $j \geq 2$, the difficulty in evaluating $P_{out}^{(j)}$ has also to be ascribed to an increasing complexity of the integration domain $\mathcal{D}_j$ in (5.11), as well as to the dependency among the events $\mathcal{E}_1, \mathcal{E}_2, \ldots, \mathcal{E}_j$. To alleviate the computational burden, this subsection explores the following approximation to $P_{out}^{(j)}$, $j \geq 2$:

$$P_{out}^{(j)} \approx 1 - \prod_{k=1}^{j} P\{\mathcal{E}_k\}\,, \quad j \geq 2\,, \tag{5.21}$$

that holds under the assumption that the random events $\mathcal{E}_k$, $k = 1, 2, \ldots, n$, be weakly dependent. To the authors' knowledge, there is no general result in the vast literature on ordered statistics that come to help in corroborating the above approximation. It has however been employed before, e.g., in [127, 130, 131]. In this work, we follow the same approach, and *a posteriori* demonstrate that it holds.

Let us begin considering the case of $n = 2$ UEs. In addition to $P_{out}^{(1)}$, only $P_{out}^{(2)}$, the probability that the receiver fails to decode the second strongest signal, has to be determined. If we recall (5.7) and (5.9), $P_{out}^{(2)}$ specializes to

$$P_{out}^{(2)} \approx 1 - P\left\{\frac{X_{(1)}}{X_{(2)} + \sigma^2} \geq \hat{\gamma}_1\right\} \cdot P\left\{\frac{X_{(2)}}{\sigma^2} \geq \hat{\gamma}_2\right\}\,; \tag{5.22}$$

it is easy to recognize that the first term in the product on the right-hand side of (5.22) coincides with $1 - P_{out}^{(1)}$. Moreover, indicating by $G_{(2)}(\cdot)$ the CDF of the random variable $X_{(2)}$, (5.22) is equivalently re-written as

$$P_{out}^{(2)} \approx 1 - \left(1 - P_{out}^{(1)}\right) \cdot \left(1 - G_{(2)}(\hat{\gamma}_2 \sigma^2)\right)\,. \tag{5.23}$$

With no loss in generality, let the unordered random variables X_i, $i = 1, 2, \ldots, n$, be numbered in accordance to the descending order of their mean received powers, so that $\overline{X}_1 > \overline{X}_2 > \ldots > \overline{X}_n$. Moreover, let us assume that the random variable δ_i measuring the distance between X_i and X_{i-1}, $\delta_i = |X_i - X_{i-1}|$, $i = 2, \ldots, n$, takes on large values with probability close to 1. In the scenarios where uplink NOMA is profitably employed, such approximation is verified, i.e., δ_i is sufficiently wide; as a matter of fact, this is the condition that allows to better discriminate among simultaneously received signals. Given this assumption holds, observe that it is possible to leverage upon a further approximation, namely, $G_{(2)}(\cdot)$ that appears in (5.23) is replaced by $G_2(\cdot)$, the CDF of the unordered random variable X_2. This leads to

$$P_{out}^{(2)} \approx 1 - \left(1 - P_{out}^{(1)}\right) \cdot \left(1 - G_2(\hat{\gamma}_2 \sigma^2)\right) \tag{5.24}$$

that represents the final, approximated $P_{out}^{(2)}$ expression when $n = 2$.

When $n = 3$ UEs are present, $P_{out}^{(2)}$ and also $P_{out}^{(3)}$ have to be determined. The probability $P_{out}^{(2)}$ modifies in

$$P_{out}^{(2)} \approx 1 - P\left\{\frac{X_{(1)}}{X_{(2)} + X_{(3)} + \sigma^2} \geq \hat{\gamma}_1\right\} \cdot P\left\{\frac{X_{(2)}}{X_{(3)} + \sigma^2} \geq \hat{\gamma}_2\right\}, \tag{5.25}$$

and denoting by $P_{out_{2,3}}^{(1)}$ the probability $P\left\{\frac{X_{(2)}}{X_{(3)}+\sigma^2} \geq \hat{\gamma}_2\right\}$, then $P_{out}^{(2)}$ is expressed as

$$P_{out}^{(2)} \approx 1 - \left(1 - P_{out}^{(1)}\right) \cdot \left(1 - P_{out_{2,3}}^{(1)}\right) \tag{5.26}$$

As regards $P_{out}^{(3)}$, the same approximation leveraged in (5.24) leads to

$$P_{out}^{(3)} \approx 1 - \left(1 - P_{out}^{(1)}\right) \cdot \left(1 - P_{out_{2,3}}^{(1)}\right) \cdot \left(1 - G_3(\hat{\gamma}_3\sigma^2)\right), \tag{5.27}$$

$G_3(\cdot)$ being the CDF of the unordered random variable X_3.

When an arbitrary number n of UEs is considered, the expression of $P_{out}^{(j)}$, $j \leq n$ can be approximated as

$$P_{out}^{(j)} \approx \begin{cases} 1 - \left(1 - P_{out}^{(1)}\right) \cdot \prod_{h=1}^{j-1} \left(1 - P_{out_{h+1,h+2,\ldots,n}}^{(1)}\right) & j < n \\ 1 - \left(1 - P_{out}^{(1)}\right) \cdot \prod_{h=1}^{n-1} \left(1 - P_{out_{h+1,h+2,\ldots,n}}^{(1)}\right) \cdot \left(1 - G_n(\hat{\gamma}_n\sigma^2)\right) & j = n \end{cases} \tag{5.28}$$

where $P_{out_{h+1,h+2,\ldots,n}}^{(1)}$ is defined as $P\left\{\frac{X_{(h+1)}}{X_{(h+2)}+\ldots+X_{(n)}} \leq \hat{\gamma}_{h+1}\right\}$, and $G_n(\cdot)$ is the CDF of the n-th random variable X_n, i.e., the CDF of the power received from the most distant user from the base station.

Last expression is illuminating, as it reveals that: (i) the outage probability $P_{out}^{(j)}$, $j \geq 2$, depends on $P_{out}^{(1)}$, the probability that NOMA fails in the presence of the same number of users; (ii) moreover, $P_{out}^{(j)}$ can be readily computed, given the expression of $P_{out}^{(1)}$ in the presence of $n, n-1, \ldots, n-j+1$ users is known. Finally, observe that the expressions we have obtained can be employed when different fading conditions are examined.

The numerical results highlight that it is possible to rely upon the proposed approximation of $P_{out}^{(j)}$, $j \geq 2$, in several meaningful settings. To the authors' knowledge, there is however no means to conclude whether (5.28) provides an upper or lower bound to the outage. When Rayleigh fading is considered, next Section reports the exact analytical expression of $P_{out}^{(1)}$ derived in the Appendix, and the closed-form approximation of $P_{out}^{(j)}$, $j \geq 2$, for an arbitrary number of simultaneously transmitting UEs. When shadowing is also introduced, the closed-form approximation of $P_{out}^{(j)}$ is provided for $j \geq 1$.

5.4 Fading Models

5.4.1 Rayleigh Fading

When the presence of Rayleigh fading is assumed, the probability density function of the received signal power X_i is exponential:

$$f_i(x_i) = \frac{1}{\overline{X}_i} \exp\left(\frac{-x_i}{\overline{X}_i}\right) \tag{5.29}$$

with mean $\overline{X}_i$,

$$\overline{X}_i = p_{t,i} \cdot k_p D_i^{-\alpha}\,, \tag{5.30}$$

k_p being a constant that depends on the operating frequency and D_i the distance between UE_i and the base station. In Appendix A, it is proved that when $\hat{\gamma}_1 \geq 1$, for an arbitrary number of users n, $P_{out}^{(1)}$ obeys the expression:

$$P_{out}^{(1)} = 1 - \sum_{k=1}^{n} \frac{\exp\left(\frac{-\hat{\gamma}_1\sigma^2}{\overline{X}_k}\right)}{\prod\limits_{\substack{i=1\\i\neq k}}^{n}\left(1 + \frac{\overline{X}_i}{\overline{X}_k}\hat{\gamma}_1\right)} \tag{5.31}$$

that reveals what limits NOMA faces if the number n of simultaneously active users is increased from 2 to higher values. In turn, taking advantage of (5.31), the fading-independent approximation of $P_{out}^{(j)}$ provided by (5.28), $j \geq 2$, specializes to

$$P_{out}^{(j)} \approx \begin{cases} 1 - \prod_{h=1}^{j}\left(\sum_{k=h}^{n} \frac{\exp\left(\frac{-\hat{\gamma}_h\sigma^2}{\overline{X}_k}\right)}{\prod\limits_{\substack{i=h\\i\neq k}}^{n}\left(1 + \frac{\overline{X}_i}{\overline{X}_k}\hat{\gamma}_h\right)}\right) & j < n \\ 1 - \prod_{h=1}^{n}\left(\sum_{k=h}^{n} \frac{\exp\left(\frac{-\hat{\gamma}_h\sigma^2}{\overline{X}_k}\right)}{\prod\limits_{\substack{i=h\\i\neq k}}^{n}\left(1 + \frac{\overline{X}_i}{\overline{X}_k}\hat{\gamma}_h\right)}\right) \cdot \exp\left(-\frac{\hat{\gamma}_n\sigma^2}{\overline{X}_n}\right) & j = n \end{cases} \tag{5.32}$$

The constraint on $\hat{\gamma}_1$ deserves a careful remark: it is the authors' belief that it does not represent a limiting factor, rather, a fairly widespread requirement in upcoming settings. As an example, enhanced V2X and high-end industrial IoT use cases are expected to require conspicuous data rates, exceeding the system available bandwidth, to support a wide range of services [19, 113].

5.4.2 Rayleigh-lognormal Shadowed Fading

When the envelope of the received signal is subject to both Rayleigh fading and lognormal shadowing, the X_i pdf is

$$f_i(x_i) = \int_0^{+\infty} \frac{1}{\overline{x}} \exp\left(\frac{-x_i}{\overline{x}}\right) \frac{1}{\sqrt{2\pi}\frac{\sigma_L}{h}\overline{x}} \times \exp\left(-\frac{\left(\ln(\overline{x}) - \mu_i\right)^2}{\frac{2\sigma_L^2}{h^2}}\right) d\overline{x} \tag{5.33}$$

where σ_L is the standard deviation of the Gaussian random variable modeling lognormal shadowing in dB, μ_i depends on the distance attenuation law, $\mu_i = \ln(p_{t,i} \cdot k_p D_i^{-\alpha})$, and $h = \frac{10}{\ln(10)}$. In our analysis, σ_L is assumed to be the same for all signals.

An additional hurdle is present here, because of the integral in (5.33). We therefore propose an approximation to (5.33), exploiting the approach put forth by Holtzman in [137]. According to [137], given a function $\psi(\theta)$ of a Gaussian random variable θ with mean μ_θ and variance σ_θ^2, the expectation $\mathbb{E}\left[\psi(\theta)\right]$ can be approximated by

$$\mathbb{E}\left[\psi(\theta)\right] \approx \frac{2}{3}\psi(\mu_\theta) + \frac{1}{6}\psi\left(\mu_\theta + \sqrt{3}\sigma_\theta\right) + \frac{1}{6}\psi\left(\mu_\theta - \sqrt{3}\sigma_\theta\right). \tag{5.34}$$

For the examined case, it is observed that the change of variable

$$y = h \cdot \ln(\overline{x}) \tag{5.35}$$

that is, $\overline{x} = \exp\left(\frac{y}{h}\right)$, leads to rewrite (5.33) in the form

$$f_i(x_i) = \int_{-\infty}^{+\infty} \frac{1}{\exp\left(\frac{y}{h}\right)} \exp\left(\frac{-x_i}{\exp\left(\frac{y}{h}\right)}\right) \frac{1}{\sqrt{2\pi}\sigma_L} \times \exp\left(-\frac{\left(y - \mu_i\right)^2}{2\sigma_L^2}\right) dy \tag{5.36}$$

that can therefore be approximated as:

$$f_i(x_i) \approx \sum_{k=1}^{3} \frac{a_k}{b_{i,k}} \exp\left(\frac{-x_i}{b_{i,k}}\right) \tag{5.37}$$

where $a_1 = \frac{2}{3}$, $a_2 = a_3 = \frac{1}{6}$, $b_{i,1} = \exp\left(\frac{\mu_i}{h}\right)$, $b_{i,2} = \exp\left(\frac{(\mu_i + \sqrt{3}\sigma_L)}{h}\right)$ and $b_{i,3} = \exp\left(\frac{(\mu_i - \sqrt{3}\sigma_L)}{h}\right)$. This linear combination of exponential functions allows to leverage the results obtained in the previous case of Rayleigh fading. Hence, when lognormal shadowing is added, for n superimposed signals $P_{out}^{(1)}$ is approximated by

$$P_{out}^{(1)} \approx 1 - \sum_{k_1=1}^{3} a_{k_1} \cdot \ldots \cdot \sum_{k_{n-1}=1}^{3} a_{k_{n-1}} \sum_{k_n=1}^{3} a_{k_n} \times \left[\sum_{i=1}^{n} \frac{\exp\left(\frac{-\hat{\gamma}_1 \sigma^2}{b_{i,k_i}}\right)}{\prod_{\substack{j=1 \\ j \neq i}}^{n} \left(1 + \frac{b_{j,k_j}}{b_{i,k_i}}\hat{\gamma}_1\right)}\right]. \tag{5.38}$$

In the $j \geq 2$ case, $P_{out}^{(j)}$ follows from (5.28) and it is approximated by

$$P_{out}^{(j)} \approx 1 - \prod_{h=1}^{j} \left(\sum_{k_h=1}^{3} a_{k_h} \cdot \ldots \cdot \sum_{k_{n-1}=1}^{3} a_{k_{n-1}} \sum_{k_n=1}^{3} a_{k_n} \cdot \left(\sum_{i=h}^{n} \frac{\exp\left(\frac{-\hat{\gamma}_h \sigma^2}{b_{i,k_i}}\right)}{\prod_{\substack{j=h \\ j \neq i}}^{n} \left(1 + \frac{b_{j,k_j}}{b_{i,k_i}} \hat{\gamma}_{h+1}\right)} \right) \right) \tag{5.39}$$

if $j < n$. When $j = n$, the value of $P_{out}^{(j)}$, $j \geq 2$, is approximated by

$$P_{out}^{(j)} \approx 1 - \prod_{h=1}^{n} \left(\sum_{k_h=1}^{3} a_{k_h} \cdot \ldots \cdot \sum_{k_{n-1}=1}^{3} a_{k_{n-1}} \sum_{k_n=1}^{3} a_{k_n} \cdot \left(\sum_{i=h}^{n} \frac{\exp\left(\frac{-\hat{\gamma}_h \sigma^2}{b_{i,k_i}}\right)}{\prod_{\substack{j=h \\ j \neq i}}^{n} \left(1 + \frac{b_{j,k_j}}{b_{i,k_i}} \hat{\gamma}_h\right)} \right) \right) \times \sum_{k=1}^{3} a_k \exp\left(-\frac{\hat{\gamma}_n \sigma^2}{b_{i,k}}\right) . \tag{5.40}$$

In the next Subsection, the excellent accuracy of the approximations in (5.32), (5.38), (5.39), and (5.40) will be demonstrated for several choices of system parameters.

5.5 Numerical Results

An exemplary set of numerical results is reported next in order to highlight the accuracy of the proposed analytical approaches, as well as to provide useful insights on the uplink performance of power-domain NOMA, when the dynamic-ordered SIC receiver is employed.

The results have been obtained for the following configuration: in (5.6), the target data rate $\hat{R}_{(j)}$ is set to 1.2 bits/s/Hz, $\forall\ j,\ j = 1, 2, \ldots, n$; in (5.30), the k_p constant is $\left(\frac{c}{4\pi f_c}\right)$, where c is the speed of light and omnidirectional antennas are assumed. The operating frequency is $f_c = 2$ GHz, the pathloss exponent is $\alpha = 2$ and the cell radius is $R = 1000$ m. As regards lognormal shadowing, unless otherwise stated, in (5.33) $\sigma_L = 4$ dB. To improve the base station capability to recover the signals coming from distinct UEs, the transmitted powers are set so as to attribute higher power levels to UEs closer to the base station. Unless otherwise stated, the UEs location along the cell radius is the one illustrated in Fig. 5.1. Namely, the ratio between the transmitted powers of UE_i and UE_j, with distances D_i and D_j from the base, $D_i < D_j$, is set to

$$\frac{p_{t,i}}{p_{t,j}} = 10^{\frac{(j-i)\Delta}{10}} , \tag{5.41}$$

and the power back-off step is $\Delta = 6$ dB.

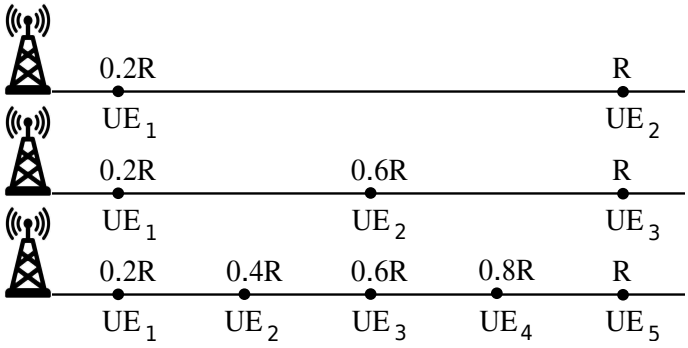

Figure 5.1: Users location along the cell radius when $n = 2, 3, 5$.

In the first set of figures, the outage probabilities are reported as a function of the largest average received SNR. Recalling that the UEs are indexed so that $\overline{X}_1 > \overline{X}_2 > \cdots > \overline{X}_n$, it follows that SNR $= \overline{X}_1/\sigma^2$. In other words, as the power law assignment privileges users that are closer to the base, the SNR is the average received signal-to-noise ratio of the UE that is the nearest to the base station.

When Rayleigh fading is considered, Fig. 5.2 shows $P_{out}^{(1)}$ as a function of the SNR, if $n = 2, 3, 5$ users are simultaneously transmitting. As indicated in Fig. 5.2, when $n = 2$, the distance D_1 of UE_1 from the base is $0.2R$ and the distance D_2 of UE_2 from it is R; when $n = 3$, $D_1 = 0.2R$, $D_2 = 0.6R$ and $D_3 = R$; when $n = 5$, $D_1 = 0.2R$, $D_2 = 0.4R$, $D_3 = 0.6R$, $D_4 = 0.8R$ and $D_5 = R$. Solid lines refer to the exact analytical evaluation, markers to $P_{out}^{(1)}$ values determined through Monte Carlo simulation, considering 10^5 samples for each plotted value. The perfect match between the analytical and the simulation results confirms the correctness of the exact closed-form of $P_{out}^{(1)}$ provided by (5.31). The figure also indicates that $P_{out}^{(1)}$ is close to 1 when the SNR is smaller than 5 dB, regardless of n. As a matter of fact, in this low SNR region, the achievable data rate of the strongest user is limited by the weak level of the received signal, rather than by the presence of simultaneously transmitting users. As the SNR increases above 10 dB, the impact of a larger number of interfering users gradually becomes more evident. Yet, observe that for

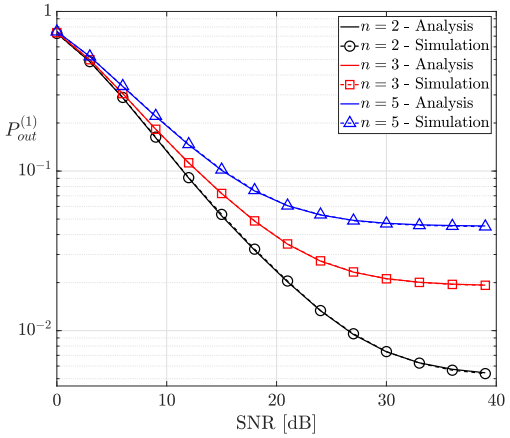

Figure 5.2: Rayleigh fading: $P_{out}^{(1)}$ as a function of the SNR, $n = 2, 3, 5$.

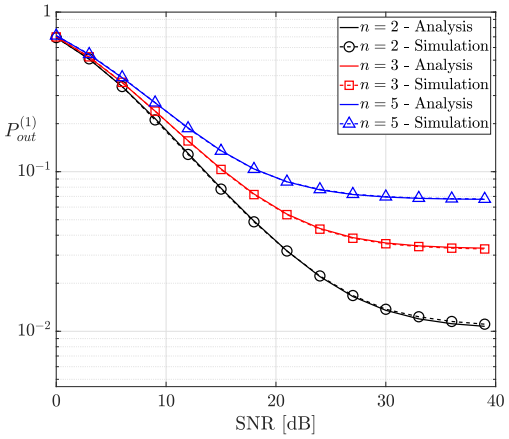

Figure 5.3: Rayleigh fading and log-normal shadowing: $P_{out}^{(1)}$ as a function of the SNR, $n = 2, 3, 5$.

$n = 5$, $P_{out}^{(1)}$ is always below 0.1 for all SNR values in the $[15, +\infty]$ range.

Fig. 5.3 shows $P_{out}^{(1)}$ as a function of the SNR, when the received signal is subject to both Rayleigh fading and log-normal shadowing, and it reveals that the closed-form in (5.38) is an excellent approximation to the exact $P_{out}^{(1)}$ computed via simulation. In this scenario too, the impact on $P_{out}^{(1)}$ of a larger number of users can be appreciated only for SNR values greater than 10 dB. The relative position of the curves is the same as observed in Fig. 5.2. However, $P_{out}^{(1)}$ takes on higher values than in the presence of Rayleigh fading only. For instance, when SNR = 30 dB, for $n = 2$ $P_{out}^{(1)}$ increases from 7×10^{-3} determined in the presence of Rayleigh fading to 1.4×10^{-2} when lognormal shadowing is also taken into account, and raises from 4.7×10^{-2} to 6.9×10^{-2} for $n = 5$. This indicates that the shadowing plays a non-negligible role in the outage probability evaluation.

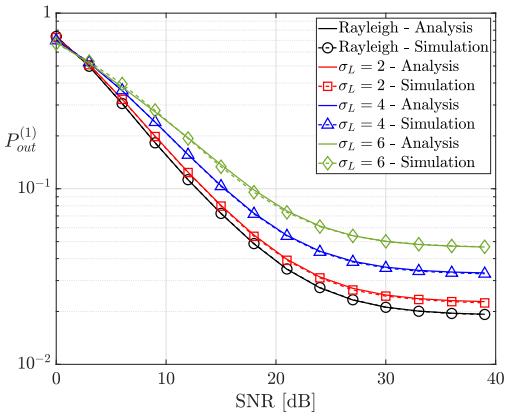

Figure 5.4: $P_{out}^{(1)}$, Rayleigh fading and Rayleigh plus lognormal, $n = 3$.

To better understand the influence of lognormal shadowing, Fig. 5.4 compares $P_{out}^{(1)}$ in the presence of Rayleigh fading against $P_{out}^{(1)}$ in the additional presence of lognormal shadowing, when $n = 3$ and three different values of σ_L, namely, 2, 4, and 6 dB are

considered. The black lowest curve and the circle markers refer to the benchmark case of Rayleigh fading; the red, blue and green curves, paired with the square, triangle and diamond markers, respectively, refer to the Rayleigh plus lognormal case. The results corroborate what was previously anticipated, quantifying the remarkable impact of the shadowing on $P_{out}^{(1)}$ for increasing values of σ_L. For instance, when $\sigma_L = 6$ dB and SNR > 30 dB, $P_{out}^{(1)}$ is 2.4 times larger than for the case of Rayleigh fading only. Also note the tightness of the approximation provided by (5.38): here too, the results obtained by simulation are nearly undistinguishable from the analytical outcomes, no matter what σ_L value is examined.

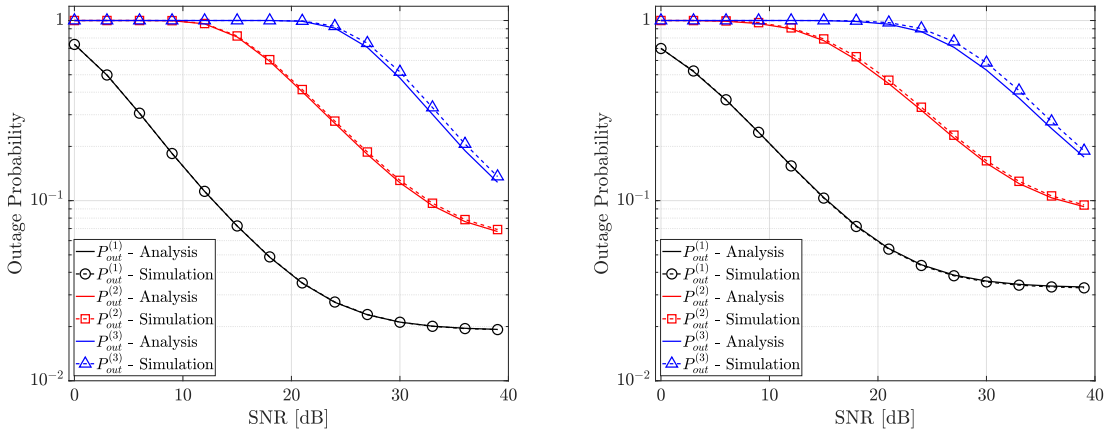

(a) Rayleigh Fading.

(b) Rayleigh fading plus log-normal shadowing.

Figure 5.5: Outage probability vs SNR, $n = 3$.

Next, Fig. 5.5(a) and 5.5(b) report $P_{out}^{(1)}$, $P_{out}^{(2)}$ and $P_{out}^{(3)}$ for the case of Rayleigh fading and Rayleigh plus lognormal shadowing, respectively, when $n = 3$. As regards $P_{out}^{(2)}$ and $P_{out}^{(3)}$, these figure show the impressive accuracy of the approximation proposed in Section 5.4 and detailed in eqs. (5.32), (5.39), and (5.40). Furthermore, they reveal that $P_{out}^{(2)}$ and $P_{out}^{(3)}$ take on significantly high values, with and without lognormal shadowing. As expected, $P_{out}^{(3)}$ takes on the worst values, as the signal coming from the third strongest user can be decoded only if both the second and the first strongest signals have already been decoded.

When considering both Rayleigh fading and lognormal shadowing, an alternative view is provided by Fig. 5.6, that shows $P_{out}^{(1)}$, $P_{out}^{(2)}$ and $P_{out}^{(3)}$ as a function of D_2/R, the normalized distance of UE_2 from the base, when UE_1 and UE_3 distances are $D_1 = 0.2R$ and $D_3 = R$, respectively. This figure indicates that the values of $P_{out}^{(2)}$ and $P_{out}^{(3)}$ lie in the range of a few percentage points; equivalently stated, approximately in 90% of the cases it is possible to support 2 simultaneous communications (and in 80% of the cases even 3). If the involved UEs all require maximum reliability, this is unacceptable. Yet, resorting to power-domain NOMA becomes truly interesting in alternative settings: for instance, whenever communication redundancy can be introduced without an excessive

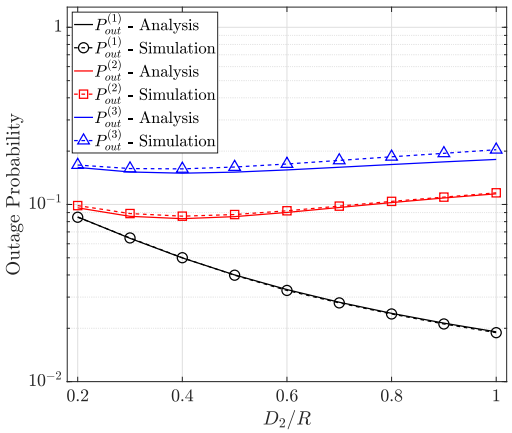

Figure 5.6: Outage probability as a function of D_2/R, $n = 3$, SNR = 30 dB.

overhead, as it happens when a modest number of packet retransmissions are introduced and packets exhibit a modest size. The figure also shows that $P_{out}^{(2)}$ and $P_{out}^{(3)}$ minima lie at $D_2 = 0.4R$ and that they are not so critical, revealing that the location of the UEs does not have to be identified with extreme accuracy. Also observe that the tightness of the proposed approximation in evaluating $P_{out}^{(3)}$ slightly worsens as D_2 approaches D_3; this happens since the spacing δ_3 no longer verifies the assumption of taking on large values with probability close to 1.

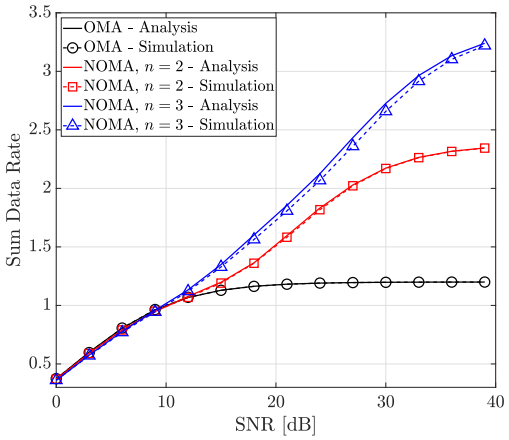

Figure 5.7: Sum data rate of OMA and NOMA as a function of SNR, Rayleigh and lognormal shadowing.

Last, Fig. 5.7 displays R_{NOMA}, the sum data rate of power-domain NOMA of (5.20), as a function of the SNR in the simultaneous presence of Rayleigh fading and lognormal shadowing, when $n = 2$ and $n = 3$, and compares it against the OMA data rate. The latter scheme is examined under the hypothesis that the signal-to-noise ratio of the OMA user coincides with SNR, the signal-to-noise ratio of the NOMA UE that is the nearest to the base station. Here too, the accuracy of the proposed approximation is striking. The gain of NOMA over OMA becomes more and more evident for increasing SNR values.

Moreover, at high SNR regimes the NOMA system with $n = 3$ users achieves a sum data rate significantly greater than 2.4 bits/s/Hz, the maximum data rate NOMA attains when $n = 2$.

5.6 Conclusions and Future Work

This Chapter has proposed a novel analytical approach to evaluate the outage probability of uplink power-domain NOMA, when a dynamic-ordered SIC receiver is employed. The method has been employed in the presence of Rayleigh fading, and Rayleigh plus lognormal shadowing. In the former setting, it has allowed to determine the probability that NOMA fails through an exact analytical expression, for a generic number of superimposed signals; in the second examined scenario, such probability has been obtained in closed-form via an excellent approximation. Moreover, the current study has put forth an approximated expression of the probability that the SIC receiver does not succeed in decoding the second, third, n-th strongest user. Monte Carlo simulations have demonstrated the accuracy of the results obtained through the proposed approximations, that clearly quantify the effects of an increasing number of simultaneously active devices (UEs and connected vehicles) on system performance. The analysis has also disclosed to what extent lognormal shadowing affects PD-NOMA behavior, revealing that its presence significantly deteriorates performance with respect to the case of Rayleigh fading only.

As of today, the research on PD-NOMA (or any other NOMA technique) has concentrated on infrastructured scenarios, as PD-NOMA requires the assistance of a central orchestration unit to carefully assign different power levels to the active users. Nevertheless, the introduction of more sophisticated V2X use cases, like vehicles platooning, opens new research paths towards the application of PD-NOMA in distributed settings, where there is no infrastructure support. In platooning, the leader vehicle is followed by platoon members which employ Cooperative-Adaptive Cruise Control (C-ACC) to minimize air drag and fuel consumption. C-ACC is an extension to the traditional ACC that relies on the V2V data exchange. Besides C-ACC, V2X communications also foresee the support of cooperative perception and maneuvering applications that require the frequent and significant exchange of information. To minimize the amount of occupied radio resources while guaranteeing the required performance levels, PD-NOMA can be profitably employed in intra-platoon communications, multiplexing simultaneously active vehicles on the same frequencies. The leader vehicle can behave as the central orchestration unit and be in charge of assigning different power levels to the platoon members. The potential of PD-NOMA in intra-platoon communications is still unexplored, and we think that it deserves careful consideration in 5G and beyond studies.

Appendix A: Derivation of $P_{out}^{(1)}$ for an arbitrary number of users

In the presence of Rayleigh fading, the outage probability $P_{out}^{(1)}$ is evaluated beginning with the special case $n = 2$. From (5.19), $P_{out}^{(1)}$ specializes to

$$P_{out}^{(1)} = 1 - (I_{S_1} + I_{S_2}) \tag{5.42}$$

where $S_1 = \{1, 2\}$, $S_2 = \{2, 1\}$,

$$I_{S_1} = \iint_{\mathcal{D}_1} f_1(x_{(1)}) f_2(x_{(2)}) dx_{(2)} dx_{(1)} = \iint_{\mathcal{D}_1} \frac{1}{\overline{X}_1} \exp\left(-\frac{x_{(1)}}{\overline{X}_1}\right) \cdot \frac{1}{\overline{X}_2} \exp\left(-\frac{x_{(2)}}{\overline{X}_2}\right) dx_{(2)} dx_{(1)} \tag{5.43}$$

and $\mathcal{D}_1$ is

$$\mathcal{D}_1 = \begin{cases} X_{(1)} \geq \hat{\gamma}_1 \cdot \left(X_{(2)} + \sigma^2\right) \\ X_{(1)} \geq X_{(2)} \geq 0 \end{cases} . \tag{5.44}$$

When the target data rate of the strongest user $\hat{\gamma}_1$ is at least equal to 1 bit/s/Hz, solving the integral in (5.43) gives

$$I_{S_1} = \frac{\exp\left(\frac{-\hat{\gamma}_1 \sigma^2}{\overline{X}_1}\right)}{1 + \frac{\overline{X}_2}{\overline{X}_1} \hat{\gamma}_1} . \tag{5.45}$$

From (5.45), I_{S_2} is readily determined as

$$I_{S_2} = \frac{\exp\left(\frac{-\hat{\gamma}_1 \sigma^2}{\overline{X}_2}\right)}{1 + \frac{\overline{X}_1}{\overline{X}_2} \hat{\gamma}_1} \tag{5.46}$$

and $P_{out}^{(1)}$ follows:

$$P_{out}^{(1)} = 1 - \left(\frac{\exp\left(\frac{-\hat{\gamma}_1 \sigma^2}{\overline{X}_1}\right)}{1 + \frac{\overline{X}_2}{\overline{X}_1} \hat{\gamma}_1} + \frac{\exp\left(\frac{-\hat{\gamma}_1 \sigma^2}{\overline{X}_2}\right)}{1 + \frac{\overline{X}_1}{\overline{X}_2} \hat{\gamma}_1} \right) . \tag{5.47}$$

When $n = 3$, $P_{out}^{(1)}$ exhibits 3! distinct contributions. The first of them, I_{S_1}, $S_1 = \{1, 2, 3\}$, is

$$I_{S_1} = \iiint_{\mathcal{D}_1} f_1(x_{(1)}) f_2(x_{(2)}) f_3(x_{(3)}) \, dx_{(3)} dx_{(2)} dx_{(1)} \tag{5.48}$$

where $\mathcal{D}_1$ is identified by the conditions

$$\mathcal{D}_1 = \begin{cases} X_{(1)} \geq \hat{\gamma}_1 \cdot \left(X_{(2)} + X_{(3)} + \sigma^2\right) \\ X_{(1)} \geq X_{(2)} \geq X_{(3)} \geq 0 \end{cases} . \tag{5.49}$$

So,

$$I_{S_1} = \iiint_{\mathcal{D}_1} p_{123}(x_{(1)}, x_{(2)}, x_{(3)})\, dx_{(3)}dx_{(2)}dx_{(1)} = \\ = \iiint_{\mathcal{D}_1} \frac{1}{\overline{X}_1}\exp\left(-\frac{x_{(1)}}{\overline{X}_1}\right)\cdot\frac{1}{\overline{X}_2}\exp\left(-\frac{x_{(2)}}{\overline{X}_2}\right)\cdot\frac{1}{\overline{X}_3}\exp\left(-\frac{x_{(3)}}{\overline{X}_3}\right)\times dx_{(3)}dx_{(2)}dx_{(1)}\,. \tag{5.50}$$

After a few lengthy steps, last integral is solved and leads to the following result

$$I_{S_1} = \frac{\exp\left(\frac{-\hat{\gamma}_1\sigma^2}{\overline{X}_1}\right)\overline{X_1}^2\overline{X_2}}{(\overline{X_1}+\overline{X_2}\hat{\gamma_1})(\overline{X_1X_2}+\overline{X_1X_3}\hat{\gamma_1}+\overline{X_2X_3}\hat{\gamma_1}+\overline{X_2X_3}\hat{\gamma_1}^2)} \tag{5.51}$$

that more aptly is written as

$$I_{S_1} = \frac{\exp\left(\frac{-\hat{\gamma}_1\sigma^2}{\overline{X}_1}\right)}{(1+\frac{\overline{X}_2}{\overline{X}_1}\hat{\gamma_1})(1+\frac{\overline{X}_3}{\overline{X}_2}\hat{\gamma_1}+\frac{\overline{X}_3}{\overline{X}_1}\hat{\gamma_1}(\hat{\gamma_1}+1))}\,. \tag{5.52}$$

Once I_{S_1} is determined, all the remaining contributions can be obtained permuting over $\mathcal{S}_\mathcal{N}$. On purpose, the set $S_2 = \{1, 3.2\}$ is considered next, which provides the result

$$I_{S_2} = \frac{\exp\left(\frac{-\hat{\gamma}_1\sigma^2}{\overline{X}_1}\right)\overline{X}_1^2\overline{X}_3}{(\overline{X}_1+\overline{X}_3\cdot\hat{\gamma_1})(\overline{X}_1\overline{X}_2+\overline{X}_1(\overline{X}_3+2\overline{X}_2\overline{X}_3\hat{\gamma_1}))}\,. \tag{5.53}$$

Observe that, when the sum $I_{S_1}+I_{S_2}$ is computed, it gives

$$C_1 = I_{S_1}+I_{S_2} = \frac{\exp\left(\frac{-\hat{\gamma}_1\sigma^2}{\overline{X}_1}\right)\overline{X}_1^2}{(\overline{X}_1+\overline{X}_2\hat{\gamma_1})(\overline{X}_1+\overline{X}_3\hat{\gamma_1})} = \frac{\exp\left(\frac{-\hat{\gamma}_1\sigma^2}{\overline{X}_1}\right)}{(1+\frac{\overline{X}_2}{\overline{X}_1}\hat{\gamma_1})(1+\frac{\overline{X}_3}{\overline{X}_1}\hat{\gamma_1})}\,. \tag{5.54}$$

If we now introduce the sets $S_3 = \{2, 1, 3\}$ and $S_4 = \{3, 1, 2\}$, permuting $\overline{X}_1$,$\overline{X}_2$ and $\overline{X}_3$ in (5.52), I_{S_3} and I_{S_4} are also determined. Their sum gives

$$C_2 = I_{S_3}+I_{S_4} = \frac{\exp\left(\frac{-\hat{\gamma}_1\sigma^2}{\overline{X}_2}\right)\overline{X}_2^2}{(\overline{X}_2+\overline{X}_1\hat{\gamma_1})(\overline{X}_2+\overline{X}_3\hat{\gamma_1})} = \frac{\exp\left(\frac{-\hat{\gamma}_1\sigma^2}{\overline{X}_2}\right)}{(1+\frac{\overline{X}_1}{\overline{X}_2}\hat{\gamma_1})(1+\frac{\overline{X}_3}{\overline{X}_2}\hat{\gamma_1})}\,. \tag{5.55}$$

The same applies to $S_5 = \{3, 2, 1\}$ and $S_6 = \{2, 3, 1\}$, whose sum C_3 is

$$C_3 = I_{S_5}+I_{S_6} = \frac{\exp\left(\frac{-\hat{\gamma}_1\sigma^2}{\overline{X}_3}\right)\overline{X}_3^2}{(\overline{X}_3+\overline{X}_1\hat{\gamma_1})(\overline{X}_3+\overline{X}_2\hat{\gamma_1})} = \frac{\exp\left(\frac{-\hat{\gamma}_1\sigma^2}{\overline{X}_3}\right)}{(1+\frac{\overline{X}_1}{\overline{X}_3}\hat{\gamma_1})(1+\frac{\overline{X}_2}{\overline{X}_3}\hat{\gamma_1})}\,. \tag{5.56}$$

Therefore, in the presence of $n = 3$ users and under the previous hypothesis $\hat{\gamma}_1 \geq 1$, $P_{out}^{(1)}$

is amenable to the writing

$$P_{out}^{(1)} = 1 - \sum_{S_i \in S_N} I_{S_i} = 1 - \sum_{k=1}^{3} C_k = \\ = 1 - \left(\frac{\exp\left(\frac{-\hat{\gamma_1}\sigma^2}{\overline{X}_1}\right)}{(1+\frac{\overline{X}_2}{\overline{X}_1}\hat{\gamma_1})(1+\frac{\overline{X}_3}{\overline{X}_1}\hat{\gamma_1})} + \frac{\exp\left(\frac{-\hat{\gamma_1}\sigma^2}{\overline{X}_2}\right)}{(1+\frac{\overline{X}_1}{\overline{X}_2}\hat{\gamma_1})(1+\frac{\overline{X}_3}{\overline{X}_2}\hat{\gamma_1})} + \frac{\exp\left(\frac{-\hat{\gamma_1}\sigma^2}{\overline{X}_3}\right)}{(1+\frac{\overline{X}_1}{\overline{X}_3}\hat{\gamma_1})(1+\frac{\overline{X}_2}{\overline{X}_3}\hat{\Gamma})} \right). \tag{5.57}$$

When the case $n = 4$ is examined, $4! = 24$ terms contribute to $P_{out}^{(1)}$; yet, it is sufficient to compute the term that corresponds to the $S_1 = \{1, 2, 3, 4\}$ set, that is expressed by

$$I_{S_1} = \iiiint_{\mathcal{D}_1} f_1(x_{(1)}) f_2(x_{(2)}) f_3(x_{(3)}) f_4(x_{(4)}) \times dx_{(4)} dx_{(3)} dx_{(2)} dx_{(1)} \tag{5.58}$$

with $\mathcal{D}_1$ now being given by

$$\mathcal{D}_1 = \begin{cases} X_{(1)} \geq \hat{\gamma_1} \cdot \left(X_{(2)} + X_{(3)} + X_{(4)} + \sigma^2\right) \\ X_{(1)} \geq X_{(2)} \geq X_{(3)} \geq X_{(4)} \geq 0 \end{cases} . \tag{5.59}$$

Here too, the assumption of exponential pdf allows to solve (5.58) in closed-form, resulting in

$$I_{S_1} = \exp\left(\frac{-\hat{\gamma_1}\sigma^2}{\overline{X}_1}\right) \\ \times \frac{\overline{X}_1^3 \overline{X}_2^2 \overline{X}_3}{(\overline{X}_1 + \overline{X}_2\hat{\gamma_1})(\overline{X}_1\overline{X}_2 + \overline{X}_1\overline{X}_3 + 2\overline{X}_2\overline{X}_3\hat{\gamma_1})(\overline{X}_1\overline{X}_2\overline{X}_3 + \overline{X}_1\overline{X}_2\overline{X}_4 + 3\overline{X}_2\overline{X}_3\overline{X}_4\hat{\gamma_1})} . \tag{5.60}$$

At first sight, last expression might look unmanageable and hinder $P_{out}^{(1)}$ determination. Yet, in analogy with the previous $n = 3$ case, the contribution in (5.60) has to be grouped with other conveniently identified terms, namely, those that correspond to the sets $S_2 = \{1, 2, 4, 3\}$, $S_3 = \{1, 3, 2, 4\}$, $S_4 = \{1, 3, 4, 2\}$, $S_5 = \{1, 4, 2, 3\}$, and $S_6 = \{1, 4, 3, 2\}$, leading to the partial sum

$$C_1 = \sum_{i=1}^{6} S_i = \frac{\exp\left(-\frac{\hat{\gamma_1}\sigma^2}{\overline{X}_1}\right)}{(1+\frac{\overline{X}_2}{\overline{X}_1}\hat{\gamma_1})(1+\frac{\overline{X}_3}{\overline{X}_1}\hat{\gamma_1})(1+\frac{\overline{X}_4}{\overline{X}_1}\hat{\gamma_1})} . \tag{5.61}$$

In an analogous manner, 3 more partial sums are computed, so that altogether 4 terms

are identified, C_k, $k = 1, 2, 3, 4$, the generic C_k being

$$C_k = \frac{\exp\left(-\frac{\hat{\gamma}_1 \sigma^2}{\overline{X}_k}\right)}{\prod\limits_{\substack{i=1 \\ i \neq k}}^{4} \left(1 + \frac{\overline{X}_i}{\overline{X}_k} \hat{\Gamma}\right)} \tag{5.62}$$

and $P_{out}^{(1)}$ is then computed as

$$P_{out}^{(1)} = 1 - \sum_{k=1}^{4} C_k \,. \tag{5.63}$$

For an arbitrary number n of superimposed signals, n partial sums, each with $(n-1)!$ elements, have to be determined. By induction, the generic sum C_k turns out to be

$$C_k = \frac{\exp\left(-\frac{\hat{\gamma}_1 \sigma^2}{\overline{X}_k}\right)}{\prod\limits_{\substack{i=1 \\ i \neq k}}^{n} (1 + \frac{\overline{X}_i}{\overline{X}_k} \cdot \hat{\gamma}_1)} \,, \tag{5.64}$$

so that $P_{out}^{(1)}$, the probability that power-domain NOMA cannot guarantee the target data rate to the strongest user, is finally written as

$$P_{out}^{(1)} = 1 - \sum_{k=1}^{n} C_k = 1 - \sum_{k=1}^{n} \frac{\exp\left(\frac{-\hat{\gamma}_1 \sigma^2}{\overline{X}_k}\right)}{\prod\limits_{\substack{i=1 \\ i \neq k}}^{n} \left(1 + \frac{\overline{X}_i}{\overline{X}_k} \hat{\gamma}_1\right)} \,. \tag{5.65}$$

under the condition $\hat{\gamma}_1 \geq 1$ bits/s/Hz.

Chapter 6

Final Remarks

Vehicular communications are instrumental for the accomplishment of more secure, efficient, and environment-friendly transports. As of today, the IEEE 802.11p and C-V2X SL technologies compete to gain the technological superiority and be selected for supporting the ITS evolution towards cooperative and automated driving. In this thesis, the C-V2X SL solution has been thoroughly analyzed from the simulative, experimental, and analytical perspective.

To do so, this thesis leveraged the MoReV2X simulator presented in Chapter 1. The implementation of MoReV2X adheres to 3GPP specifications and evaluation guidelines introduced in Release 14 through Release 16. As a result, MoReV2X allows the simulation of LTE-V2X SL and NR-V2X SL communications, including a specific set of channel models, traffic models, and performance metrics. In Chapter 1, the value of MoReV2X has been illustrated by reporting an exemplary set of results. The presented results offered a preliminary insight on the achievable performance of NR-V2X SL for two alternative SCS choices and two different traffic types, namely, periodic and aperiodic. Moreover, Chapter 1 has reported the BLER curves obtained during the link-level analysis of the C-V2X SL technology, highlighting the impact of PHY layer aspects on the performance of C-V2X SL communications. In this Chapter, the relevance of the BLER curves to the development of the simplified PHY layer abstraction models employed by MoReV2X has also been discussed.

Next, Chapter 2 leveraged the MoReV2X simulator to exhaustively analyze C-V2X SL communications, concentrating on the LTE-V2X Mode 4 and NR-V2X Mode 2 distributed resource allocation modes. Chapter 2 studied the coexistence of periodic and aperiodic traffic in LTE-V2X Mode 4, performing an exhaustive set of simulations that considered different percentages and arrival rates of aperiodic flows, size of aperiodic packets, and vehicular densities. As the obtained results showed that LTE-V2X Mode 4 is not able to effectively serve aperiodic traffic, a new reservation-less scheduling scheme has been proposed, and its superiority with respect to legacy solutions has been numerically demonstrated.

In the NR-V2X Mode 2 case, Chapter 2 has examined the impact of the latest MAC

sublayer features introduced by 3GPP in Release 16 on system performance. Specifically, Chapter 2 concentrated on the analysis of the re-evaluation mechanism and on the comparison between the SPS and DS schemes. First, numerical simulations revealed that the re-evaluation mechanism is not able to improve NR-V2X SL performance despite a significant implementation complexity. In the case of periodic traffic, the re-evaluation mechanism is able to detect only a negligible number of collisions. The amount of detected collisions increases with aperiodic traffic, but the selection of collision-free resources after a re-evaluation is not guaranteed. The obtained results demonstrated that the MAC sublayer of NR-V2X Mode 2 and LTE-V2X Mode 4 faces the same challenges when the dissemination of aperiodic traffic is considered.

Then, further valuable insights on the NR-V2X Mode 2 operation have been reported by the comparative analysis of the SPS and DS schemes. Analytical models and simulation results have demonstrated that varying PDB requirements do not affect the DS strategy operation, whereas they can significantly deteriorate the SPS strategy. Moreover, simulation results have also shown that the SPS strategy represents the best solution for serving fixed size periodic traffic, while the DS strategy attains the best performance in the (fixed or variable size) aperiodic traffic case. In Chapter 2, a novel adaptive strategy that allows vehicles to adapt their scheduling scheme on the basis of the generated traffic has also been proposed.

The MoReV2X simulator has been employed also in Chapter 3 to assess the performance of a new AI-based resource allocation strategy for the dissemination of CAMs in LTE-V2X Mode 4. The proposed solution leverages ML to forecast future CAM interarrival times and accordingly optimize the MAC layer configuration. The obtained results show that a simple AI technique like the KNN algorithm can achieve an excellent accuracy and that the AI-enhanced resource allocation outperforms the legacy LTE-V2X Mode 4 approach. Although Chapter 3 concentrates on LTE-V2X Mode 4, it discloses the potential of ML for disseminating aperiodic traffic also in NR-V2X Mode 2.

In Chapter 4, the simulative analysis of C-V2X SL communications has been complemented with experimental results, reporting the findings of a vast measurement campaign conducted to analyze the temporal patterns that characterize the dissemination of CAMs and VAMs via LTE-V2X SL. VAMs and CAMs generated by bicycles, e-scooters, cars and motorbikes were collected through numerous field tests performed in urban, suburban, and highway settings. The analysis of the experimental results revealed the impact of the various triggering conditions on the temporal patterns of VAMs and CAMs. In the case of VAMs, an amendment to the current generation rules to reduce the amount of generated messages, without losing relevant information about the VRU movements, has been proposed and proved successful. Furthermore, Chapter 4 has experimentally assessed the PRR of LTE-V2V SL communications, revealing worse communication range levels with respect to simulations and other experimental studies performed in controlled environments. Chapter 4 has also presented the first experimental analysis of the PRR attained by bicycle-to-vehicle communications based on LTE-V2X SL, providing prelimi-

nary indications on the achievable communication range of safety-related VRU protection use cases.

Lastly, Chapter 5 has presented a novel analytical approach to evaluate the outage probability of uplink power-domain NOMA when a dynamic-ordered SIC receiver and an arbitrary number of simultaneously active devices (UEs or V2X-enabled vehicles) is considered. The method has been employed both in the presence of Rayleigh fading and Rayleigh plus lognormal shadowing, providing closed-form expressions whose correctness and excellent accuracy have been demonstrated through Monte Carlo simulations. Chapter 5 has also put forth approximated closed-form expressions to determine the probability that the SIC receiver does not succeed in decoding the second, third, n-th strongest user. In this Chapter, the application of power-domain NOMA to V2V communications in distributed scenarios, where there is no infrastructure support, has also been discussed.

This thesis has provided an accurate and exhaustive analysis of the C-V2X SL technology, clearly identifying its strengths and limitations. At the same time, this thesis has highlighted the need for a V2X technology able to serve different traffic types and cope with a wide range of latency, reliability, and communication range requirements. To this end, the PHY layer of the C-V2X SL technology features an extremely flexible design and its configuration can be tailored to the needs of each specific V2X application. However, the C-V2X SL MAC sublayer is not characterized by the same flexibility, and the key question of how to accommodate different traffic types still remains unanswered. In order to demonstrate its superiority with respect to its IEEE 802.11p (and IEEE 802.11bd) competitor, the C-V2X SL technology will require a partial re-design of the MAC sublayer during the "beyond 5G" standardization phase. In this regard, this thesis has demonstrated the potential of AI-based resource allocation strategies for the support of realistic V2X applications and laid the analytical foundations for the deployment of novel NGMA solutions also in the V2X ecosystem. The design of new and effective MAC sublayer solutions will have to be complemented by an exhaustive characterization of V2X traffic sources, since future CAVs will disseminate more sophisticated messages in addition to CAMs, e.g., CPMs and MCMs, which exhibit peculiar features in terms of message size and inter-arrival time that cannot be neglected.

Bibliography

[1] H. Hartenstein and L. P. Laberteaux, "A tutorial survey on vehicular ad hoc networks," *IEEE Communications Magazine*, vol. 46, no. 6, pp. 164-171, June 2008.

[2] The PATH project. https://path.berkeley.edu/ (accessed, Dec. 21, 2022).

[3] The PROMETHEUS project, EUREKA website. https://web.archive.org/web/20180814201633/http://www.eurekanetwork.org/project/id/45 (accessed, Dec. 21, 2022).

[4] 3GPP, "TR 21.914 Release 14 Description; Summary of Rel-14 Work Items (v14.0.0, Release 14)," 3GPP, Tech. Report, May 2018.

[5] 3GPP, "TR 37.885; Study on evaluation methodology of new Vehicle-to-Everything (V2X) use cases for LTE and NR (v15.3.0, Release 15)," 3GPP, Tech. Rep., June 2019.

[6] 3GPP, "TR 21.916 Release 16 Description; Summary of Rel-16 Work Items (v16.2.0, Release 16)," 3GPP, Tech. Report, June 2022.

[7] L. Lusvarghi and M. L. Merani, "MoReV2X - A New Radio Vehicular Communication Module for ns-3," *2021 IEEE 94th Veh. Technol. Conf. (VTC2021-Fall)*, 2021, pp. 1-7.

[8] L. Lusvarghi, B. Coll-Perales, J. Gozalvez and M. L. Merani, "Link Level Evaluation of 5G NR V2X Sidelink Communication," submitted to *IEEE Transactions on Vehicular Technology*.

[9] L. Lusvarghi and M. L. Merani, "On the Coexistence of Aperiodic and Periodic Traffic in Cellular Vehicle-to-Everything," *IEEE Access*, vol. 8, pp. 207076-207088, 2020.

[10] A. Molina-Galan, B. Coll-Perales, L. Lusvarghi, J. Gozalvez and M. L. Merani, "How does 5G NR V2X Mode 2 Handle Aperiodic Packets and Variable Packet Sizes?," *2022 IEEE 23rd International Conference on High Performance Switching and Routing (HPSR)*, 2022, pp. 183-188.

[11] A. Molina-Galan, L. Lusvarghi, B. Coll-Perales, J. Gozalvez and M. L. Merani, "On the Impact of Re-evaluation in 5G NR V2X Mode 2," submitted to *IEEE Transactions on Vehicular Technology*.

[12] L. Lusvarghi, A. Molina-Galan, B. Coll-Perales, J. Gozalvez and M. L. Merani, "A Comparative Analysis of the Semi-Persistent and Dynamic Scheduling Schemes in NR-V2X Mode 2," submitted to *Elsevier Vehicular Communications*.

[13] L. Lusvarghi, C. A. Grazia, M. Klapez, M. Casoni and M. L. Merani, "Awareness Messages by Vulnerable Road Users and Vehicles: Field Tests via LTE-V2X," submitted to *IEEE Transactions on Intelligent Vehicles*.

[14] CAR 2 CAR Communication Consortium, "Survey on ITS-G5 CAM statistics," TR2052, V1.0.1, Dec. 2018.

[15] L. Lusvarghi and M. L. Merani, "Machine Learning for Disseminating Cooperative Awareness Messages in Cellular V2V Communications," *IEEE Transactions on Vehicular Technology*, vol. 71, no. 7, pp. 7890-7903, July 2022.

[16] European Telecommunications Standards Institute (ETSI), "Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Specification of Cooperative Awareness Basic Service," EN 302 637-2 V1.4.1, April 2019.

[17] European Telecommunications Standards Institute (ETSI), "Intelligent Transport Systems (ITS); Vulnerable Road Users (VRU) awareness; Part 3: Specification of VRU awareness basic service; Release 2," TS 103 300-3 V2.1.1, November 2020.

[18] L. Lusvarghi and M. L. Merani, "Fundamental Limits on the Uplink Performance of the Dynamic-Ordered SIC Receiver," *IEEE Access*, vol. 10, pp. 73178-73189, 2022.

[19] 3GPP, "Study on enhancement of 3GPP Support for 5G V2X Services," TR 22.886, Rel-16 V16.2.0, Dec. 2018.

[20] P. A. Lopez et al., "Microscopic Traffic Simulation using SUMO," *21st Int. Conf. on Intell. Transp. Syst. (ITSC)*, 2018, pp. 2575-2582.

[21] 3GPP, "Sidelink physical layer structure for NR V2X," R1-1911882, 3GPP TSG RAN WG1 Meeting #99, Nov. 2019.

[22] 3GPP, "Link level evaluations of NR PSCCH," R1-1903180, 3GPP TSG RAN WG1 Meeting #96, March 2019.

[23] 3GPP, "Link level evaluations of NR PSSCH," R1-1903181, 3GPP TSG RAN WG1 Meeting #96, March 2019.

[24] ETSI, "Intelligent Transport Systems (ITS); LTE-V2X and NR-V2X Access layer specification for Intelligent Transport Systems operating in the 5 GHz frequency band; Release 2," ETSI EN 303 798 V 1.1.7, June 2021.

[25] F. Eckermann, M. Kahlert and C. Wietfeld, "Performance Analysis of C-V2X Mode 4 Communication Introducing an Open-Source C-V2X Simulator," *2019 IEEE 90th Veh. Technol. Conf. (VTC2019-Fall)*, 2019, pp. 1-5.

[26] B. McCarthy and A. O'Driscoll, "OpenCV2X Mode 4: A Simulation Extension for Cellular Vehicular Communication Networks," *2019 IEEE 24th International Workshop on Computer Aided Modeling and Design of Communication Links and Networks (CAMAD)*, 2019, pp. 1-6.

[27] G. Cecchini, A. Bazzi, B. M. Masini and A. Zanella, "LTEV2Vsim: An LTE-V2V simulator for the investigation of resource allocation for cooperative awareness," *2017 5th IEEE International Conference on Models and Technologies for Intelligent Transportation Systems (MT-ITS)*, 2017, pp. 80-85.

[28] V. Todisco, S. Bartoletti, C. Campolo, A. Molinaro, A. O. Berthet and A. Bazzi, "Performance Analysis of Sidelink 5G-V2X Mode 2 Through an Open-Source Simulator," *IEEE Access*, vol. 9, pp. 145648-145661, 2021.

[29] N. Patriciello, S. Lagen, B. Bojovic and L. Giupponi, "An E2E Simulator for 5G NR Networks," *Elsevier Simulation Modelling Practice and Theory (SIMPAT)*, vol. 96, 101933, Nov. 2019.

[30] Z. Ali, S. Lagén, L. Giupponi and R. Rouil, "3GPP NR V2X Mode 2: Overview, Models and System-Level Evaluation," *IEEE Access*, vol. 9, pp. 89554-89579, 2021.

[31] S. Lagen, K. Wanuga, H. Elkotby, S. Goyal, N. Patriciello and L. Giupponi, "New Radio Physical Layer Abstraction for System-Level Simulations of 5G Networks," *IEEE International Conference on Communications (ICC)*, 2020, pp. 1-7.

[32] M. H. C. Garcia *et al.*, "A Tutorial on 5G NR V2X Communications," *IEEE Commun. Surveys & Tutorials*, vol. 23, no. 3, pp. 1972-2026, 3rd quarter 2021.

[33] 3GPP, "Link level evaluations on sidelink for NR V2X," R1-1900852, 3GPP TSG RAN WG1 Ad-Hoc Meeting 1901, Jan. 2019.

[34] T. Dominguez-Bolano, J. Rodriguez-Pineiro, J. A. Garcia-Naya and L. Castedo, "The GTEC 5G link-level simulator," *2016 1st International Workshop on Link- and System Level Simulations (IWSLS)*, 2016, pp. 1-6.

[35] S. Pratschner et al., "Versatile mobile communications simulation: the Vienna 5G Link Level Simulator," *EURASIP Journal on Wireless Communications and Networking*, pp. 1-17, 2018.

[36] W. Anwar, N. Franchi and G. Fettweis, "Physical Layer Evaluation of V2X Communications Technologies: 5G NR-V2X, LTE-V2X, IEEE 802.11bd, and IEEE 802.11p," *2019 IEEE 90th Vehicular Technology Conference (VTC2019-Fall)*, 2019, pp. 1-7.

[37] D. Wang, R. R. Sattiraju, A. Qiu, S. Partani and H. D. Schotten, "Methodologies of Link-Level Simulator and System-Level Simulator for C-V2X Communication," *2019 IEEE 2nd International Conference on Electronics and Communication Engineering (ICECE)*, 2019, pp. 178-184.

[38] S. -Y. Lien et al., "3GPP NR Sidelink Transmissions Toward 5G V2X," *IEEE Access*, vol. 8, pp. 35368-35382, 2020.

[39] 3GPP, "TS 38.331 NR; Radio Resource Control (RRC) protocol specification (v16.3.1, Release 16)," 3GPP, Tech. Spec., Jan. 2021.

[40] 3GPP, "TS 38.214 NR; Physical layer procedures for data (v16.5.0, Release 16)," 3GPP, Tech. Spec., March 2021.

[41] 3GPP, "TS 38.212 NR; Multiplexing and channel coding (v16.5.0, Release 16)," 3GPP, Tech. Spec., March 2021.

[42] 3GPP, "TS 38.211 NR; Physical channels and modulation (v16.5.0, Release 16)," 3GPP, Tech. Spec., March 2021.

[43] 3GPP, "TR 36.885; Study on LTE-based V2X Services (v14.0.0, Release 14)," 3GPP, Tech. Rep., June 2016.

[44] 3GPP, "TS 38.215 NR; Physical layer measurements (v16.4.0, Release 16)," 3GPP, Tech. Spec., Dec. 2020.

[45] R. Rouil, F. J. Cintrón, A. Ben Mosbah, and S. Gamboa, "Implementation and Validation of an LTE D2D Model for Ns-3," *Proceedings of the Workshop on Ns-3 (WNS3 '17)*, Porto, Portugal, 2017, pp. 55-62.

[46] L. Gibellini and M. L. Merani, "Out-of-Coverage Multi-Hop Road Safety Message Distribution via LTE-A Cellular V2V (C-V2V)," *2018 IEEE 88th Vehicular Technology Conference (VTC-Fall)*, 2018, pp. 1-6.

[47] R. Molina-Masegosa, M. Sepulcre, J. Gozalvez, F. Berens, and V. Martinez, "Empirical models for the realistic generation of cooperative awareness messages in vehicular networks," *IEEE Transactions on Vehicular Technology*, vol. 69, no. 5, pp. 5713-5717, May 2020.

[48] 3GPP, "TS 38.101-1 NR; User Equipment (UE) radio transmission and reception; Part 1: Range 1 Standalone (Release 17)," 3GPP, Tech. Spec., March 2021.

[49] J. Gozalvez, M. Sepulcre and R. Bauza, "Impact of the radio channel modelling on the performance of VANET communication protocols," *Telecommunication Systems 50*, pp. 149–167, 2012.

[50] M. Mezzavilla, M. Miozzo, M. Rossi, N. Baldo and M. Zorzi, "A lightweight and accurate link abstraction model for the simulation of LTE networks in ns-3," *Proc. of the 15th ACM international conference on Modeling, analysis and simulation of wireless and mobile systems (MSWiM '12)*, New York, NY, USA, 2012, pp. 55-60.

[51] I. Tal and A. Vardy, "List Decoding of Polar Codes," *IEEE Transactions on Information Theory*, vol. 61, no. 5, pp. 2213-2226, May 2015.

[52] Z. Ali, S. Lagén and L. Giupponi, "On the impact of numerology in NR V2X Mode 2 with sensing and random resource selection," *2021 IEEE Vehicular Networking Conference (VNC)*, 2021, pp. 151-157.

[53] S. -Y. Lien et al., "3GPP NR Sidelink Transmissions Toward 5G V2X," *IEEE Access*, vol. 8, pp. 35368-35382, 2020.

[54] 3GPP, "TS 38.321 NR; Medium Access Control (MAC) protocol specification (v16.4.0, Release 16)," 3GPP, Tech. Spec., March 2021.

[55] 3GPP, "5G; Service requirements for enhanced V2X scenarios," TS 22.186, Rel-17 V17.0.0, April 2022.

[56] R. Molina-Masegosa, J. Gozalvez and M. Sepulcre, "Comparison of IEEE 802.11p and LTE-V2X: An Evaluation With Periodic and Aperiodic Messages of Constant and Variable Size," *IEEE Access*, vol. 8, pp. 121526-121548, 2020.

[57] R. Molina-Masegosa and J. Gozalvez, "System Level Evaluation of LTE-V2V Mode 4 Communications and Its Distributed Scheduling," *2017 IEEE 85th Vehicular Technology Conference (VTC Spring)*, Sydney, NSW, 2017, pp. 1-5.

[58] A. Bazzi, G. Cecchini, A. Zanella and B. M. Masini, "Study of the Impact of PHY and MAC Parameters in 3GPP C-V2V Mode 4," *IEEE Access*, vol. 6, pp. 71685-71698, 2018.

[59] M. Gonzalez-Martín, M. Sepulcre, R. Molina-Masegosa and J. Gozalvez, "Analytical Models of the Performance of C-V2X Mode 4 Vehicular Communications," *IEEE Transactions on Vehicular Technology*, vol. 68, no. 2, pp. 1155-1166, Feb. 2019.

[60] C. Campolo, A. Molinaro, F. Romeo, A. Bazzi and A. O. Berthet, "Full duplex-aided Sensing and Scheduling in Cellular-V2X Mode 4," *Proc. 1st ACM MobiHoc Workshop on Technologies, mOdels, and Protocols for Cooperative Connected Cars (TOP-Cars '19)*, Catania, Italy, July 2019, pp. 19-24.

[61] A. Bazzi, C. Campolo, A. Molinaro, A. O. Berthet, B. M. Masini and A. Zanella, "On Wireless Blind Spots in the C-V2X Sidelink," *IEEE Transactions on Vehicular Technology*, vol. 69, no. 8, pp. 9239-9243, Aug. 2020.

[62] R. Molina-Masegosa, M. Sepulcre and J. Gozalvez, "Geo-Based Scheduling for C-V2X Networks," *IEEE Transactions on Vehicular Technology*, vol. 68, no. 9, pp. 8397-8407, Sept. 2019.

[63] F. Romeo, C. Campolo, A. Molinaro, A. O. Berthet, "DENM Repetitions to Enhance Reliability of the Autonomous Mode in NR V2X Sidelink," *2020 IEEE 91st Vehicular Technology Conference (VTC2020-Spring)*, Antwerp, Belgium, May 2020, pp. 1-5.

[64] M. Muhammad Saad, M. Ashar Tariq, M. Mahmudul Islam, M. Toaha Raza Khan, J. Seo and D. Kim, "Enhanced Semi-persistent scheduling (e-SPS) for Aperiodic Traffic in NR-V2X," *Proceedings of the 2022 International Conference on Artificial Intelligence in Information and Communication (ICAIIC)*, Feb. 2022, pp. 171-175.

[65] M. Bezmenov, Z. Utkovski, K. Sambale and S. Stanczak, "Semi-Persistent Scheduling with Single Shot Transmissions for Aperiodic Traffic," *2021 IEEE 93rd Vehicular Technology Conference (VTC2021-Spring)*, 2021, pp. 1-7.

[66] Y. Yoon and H. Kim, "A Stochastic Reservation Scheme for Aperiodic Traffic in NR V2X Communication," *2021 IEEE Wireless Communications and Networking Conference (WCNC)*, 2021, pp. 1-6.

[67] C. Campolo, V. Todisco, A. Molinaro, A. Berthet, S. Bartoletti and A. Bazzi, "Improving Resource Allocation for beyond 5G V2X Sidelink Connectivity," *Proceedings of the 2021 55th Asilomar Conference on Signals, Systems, and Computers*, Nov. 2021, pp. 55-60.

[68] 5GAA, "V2X Functional and Performance Test Report; Test Procedures and Results," 5GAA Report, Apr. 2019.

[69] CAR 2 CAR Communication Consortium, "Guidance for day 2 and beyond roadmap," White Paper WP2072, V1.1.0, Sep. 2019.

[70] F. Hussain, R. Hussain, A. Anpalagan and A. Benslimane, "A New Block-Based Reinforcement Learning Approach for Distributed Resource Allocation in Clustered IoT Networks," *IEEE Transactions on Vehicular Technology*, vol. 69, no. 3, pp. 2891-2904, March 2020.

[71] H. Xiang, M. Peng, Y. Sun and S. Yan, "Mode Selection and Resource Allocation in Sliced Fog Radio Access Networks: A Reinforcement Learning Approach," *IEEE Transactions on Vehicular Technology*, vol. 69, no. 4, pp. 4271-4284, April 2020.

[72] H. Ye, Y. Li and B.-H. F. Juang, "Deep Reinforcement Learning based Resource Allocation for V2V Communications," *IEEE Transactions on Vehicular Technology*, vol. 68, no. 4, pp. 3163-3173, April 2019.

[73] H. Ye, L. Liang, G. Ye Li, J. Kim, L. Lu and M. Wu, "Machine Learning for Vehicular Networks: Recent Advances and Application Examples," *IEEE Veh. Technol. Mag.*, vol. 13, no. 2, pp. 94-101, June 2018.

[74] F.A. Ghaleb *et al.*, "Deep Kalman Neuro Fuzzy-Based Adaptive Broadcasting Scheme for Vehicular Ad Hoc Network: a Context-Aware Approach," *IEEE Access*, vol. 8, pp. 217744-217761, Dec. 2020.

[75] J. Aznar-Poveda, E. Egea-López and A. J. García-Sanchez, "Cooperative Awareness Message Dissemination in EN 302 637-2: an Adaptation for Winding Roads," *2020 IEEE 91st Veh. Technol. Conf. (VTC2020-Spring)*, 2020, pp. 1-5.

[76] O. Amador, I. Soto, M. Uruena and M. Calderon, "GoT: Decreasing DCC Queueing for CAM Messages," *IEEE Communications Letters*, vol. 24, no. 12, pp. 2974-2978, Dec. 2020.

[77] M. Sepulcre, J. Gozalvez, G. Thandavarayan, B. Coll-Perales, J. Schindler and M. Rondinone, "On the Potential of V2X Message Compression for Vehicular Networks," *IEEE Access*, vol. 8, pp. 214254-214268, December 2020.

[78] A. Bazzi, A. Zanella and B. M. Masini, "Optimizing the Resource Allocation of Periodic Messages with Different Sizes in LTE-V2V," *IEEE Access*, vol. 1, pp. 43820-43830, March 2019.

[79] A. Bazzi, C. Campolo, B. M. Masini, A. Molinaro and A. Zanella, "Enhancing Cooperative Driving in IEEE 802.11 Vehicular Networks Through Full-Duplex Radio," *IEEE Transactions on Wireless Communications*, vol. 17, No. 4, pp. 2402-2416, April 2018.

[80] N. Lyamin, A. Vinel, M. Jonsson and B. Bellalta, "Cooperative Awareness in VANETs: on ETSI EN 302 637-2 Performance," *IEEE Transactions on Vehicular Technology*, vol. 67, no. 1, pp. 17-28, Jan. 2018.

[81] G. Giambene, M. S. Rahman and A. Vinel, "Analysis of V2V Sidelink Communications for Platoon Applications," *2020 IEEE Int. Conf. on Commun. (ICC)*, 2020, pp. 1-6.

[82] S. Roger *et al.*, "Low-Latency Layer-2-Based Multicast Scheme For Localized V2X Communications," *IEEE Transactions on Intelligent Transportation Systems*, vol. 20, no. 8, pp. 2962-2975, Aug. 2019.

[83] B. Coll-Perales, J. Gozalvez, M. Gruteser, "Sub-6GHz Assisted MAC for Millimeter Wave Vehicular Communications," *IEEE Communications Magazine*, vol. 57, no. 3, pp. 125-131.

[84] Z. Li, L. Xiang, X. Ge, G. Mao and H. -C. Chao, "Latency and Reliability of mmWave Multi-Hop V2V Communications Under Relay Selections," *IEEE Transactions on Vehicular Technology*, vol. 69, no. 9, pp. 9807-9821, Sept. 2020.

[85] P. Harrington, *Machine Learning in Action*, Manning Publications, April 2012.

[86] OpenStreetMap contributors, "Planet dump retrieved from https://planet.osm.org," OpenStreetMap.org, https://www.openstreetmap.org (accessed Feb. 1, 2021).

[87] A. Wegener, M. Piorkowski, M. Raya, H. Hellbruck, S. Fischer and J. Hubaux, "TraCI: An Interface for Coupling Road Traffic and Network Simulators," *Proc. 11th Commun. and Netw. Siml. Symp. (CNS '08)*, April 2008, pp. 155-163.

[88] N. V. Chawla, K. W. Bowyer, L. O. Hall and W. P. Kegelmeyer, "SMOTE: synthetic minority over-sampling technique," *Journal of Artif. Intell. Res.*, Jan. 2002, pp. 321-357.

[89] A. Houenou, P. Bonnifait, V. Cherfaoui and W. Yao, "Vehicle trajectory prediction based on motion model and maneuver recognition," in *2013 IEEE/RSJ Int. Conf. on Intell. Robots and Systems*, Tokyo, 2013, pp. 4363-4369.

[90] N. Puller, H. -J. Günther, G. Lucas, A. Leschke and V. Rocco, "Towards Increasing VRU Safety: A Map-Based and Data-Driven Analysis of Accident Black Spots," *2021 IEEE Veh. Netw. Conf. (VNC)*, 2021, pp. 60-67.

[91] A. Willecke, K. Garlichs, F. Schulze and L. C. Wolf, "Vulnerable Road Users Are Important As Well: Persons in the Collective Perception Service," *2021 IEEE Veh. Netw. Conf. (VNC)*, 2021, pp. 24-31.

[92] P. Ghorai, A. Eskandarian, Y. Kim and G. Mehr, "State Estimation and Motion Prediction of Vehicles and Vulnerable Road Users for Cooperative Autonomous Driving: A Survey," *IEEE Transactions on Intelligent Transportation Systems*, 2022, Early Access article.

[93] K. Saleh, M. Hossny and S. Nahavandi, "Contextual Recurrent Predictive Model for Long-Term Intent Prediction of Vulnerable Road Users," *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 8, pp. 3398-3408, Aug. 2020.

[94] M. Goldhammer, S. Kohler, S. Zernetsch, K. Doll, B. Sick and K. Dietmayer, "Intentions of Vulnerable Road Users-Detection and Forecasting by Means of Machine Learning," *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 7, pp. 3035-3045, July 2020.

[95] A. Yáñez and S. Céspedes, "Pedestrians also Have Something to Say: Integration of Connected VRU in Bidirectional Simulations," *2020 IEEE Veh. Netw. Conf. (VNC)*, 2020, pp. 1-4.

[96] Y. Jia and D. Cebon, "Measuring the Motion of Vulnerable Road Users Relative to Moving HGVs," *IEEE Transactions on Intelligent Transportation Systems*, vol. 20, no. 4, pp. 1404-1415, April 2019.

[97] R.B. Zadeh, M. Ghatee and H.R. Eftekhari, "Three-Phases Smartphone-Based Warning System to Protect Vulnerable Road Users under Fuzzy Conditions," *IEEE Transactions on Intelligent Transportation Systems*, vol. 19, no. 7, pp. 2086-2098, July 2018.

[98] V. R. S. Banjade, S. C. Jha, K. Sivanesan, L. Gomes Baltar, S. A. Sehra and S. J. Tan, "Vulnerable Road Users Safety in Infrastructure Assisted Intelligent Transportation System," *2021 IEEE Int. Smart Cities Conf. (ISC2)*, 2021, pp. 1-7.

[99] M. Shan, K. Narula, Y. F. Wong, S. Worrall, M. Khan, P. Alexander and E. Nebot, "Demonstration of Cooperative Perception: Safety and Robustness in Connected and Automated Vehicle Operations," *MDPI Sensors*, vol. 21, no. 1, 2021.

[100] P. Merdrignac, O. Shagdar and F. Nashashibi, "Fusion of Perception and V2P Communication Systems for the Safety of Vulnerable Road Users," *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 7, pp. 1740-1751, July 2017.

[101] M. Boban and P. M. d'Orey, "Exploring the Practical Limits of Cooperative Awareness in Vehicular Communications," *IEEE Transactions on Vehicular Technology*, vol. 65, no. 6, pp. 3904-3916, June 2016.

[102] V. Maglogiannis, D. Naudts, S. Hadiwardoyo, D. van den Akker, J. Marquez-Barja and I. Moerman, "Experimental V2X Evaluation for C-V2X and ITS-G5 Technologies in a Real-Life Highway Environment," *IEEE Trans. on Netw. and Service Manage.*, November 2021.

[103] 5G Automotive Association (5GAA), "V2X Functional and Performance Test Report; Test Procedures and Results," 5GAA P-190033, April 2019.

[104] European Telecommunications Standards Institute (ETSI), "Intelligent Transport Systems (ITS); Vulnerable Road Users (VRU) awareness; Part 1: Use Cases definition; Release 2," ETSI TR 103 300-1 V2.1.1, September 2019.

[105] European Telecommunications Standards Institute (ETSI), "Intelligent Transport System (ITS); Vulnerable Road Users (VRU) awareness; Part 2: Functional Architecture and Requirements definition; Release 2," ETSI TS 103 300-2 V2.1.1, May 2020.

[106] Society of Automotive Engineers (SAE), "Vulnerable Road User Safety Message Minimum Performance Requirements," SAE J 2945/9, March 2017.

[107] I. Rashdan, M. Schmidhammer, F. de Ponte Mueller and S. Sand, "Performance Evaluation of Vehicle-to-Vehicle Communication for Cooperative Collision Avoidance at Urban Intersections," *2017 IEEE 86th Veh. Tech. Conf. (VTC-Fall)*, 2017, pp. 1-5.

[108] J. Aznar-Poveda, E. Egea-López and A. García-Sánchez, "Cooperative Awareness Message Dissemination in EN 302 637-2: an Adaptation for Winding Roads," *2020 IEEE 91st Veh. Technol. Conf. (VTC2020-Spring)*, 2020, pp. 1-5.

[109] Quectel AG15 Module. [Online]. Available: https://www.quectel.com/product/c-v2x-ag15.

[110] European Telecommunications Standards Institute (ETSI), "Intelligent Transport Systems (ITS); Radiocommunications equipment operating in the 5855 MHz to 5925 MHz frequency band; Harmonised Standard covering the essential requirements of article 3.2 of Directive 2014/53/EU," ETSI EN 302 571 V2.1.1, February 2017.

[111] Open source repository with the presented CAMs and VAMs dataset: https://github.com/LLusvarghi/CAM-VAM-dataset.

[112] Cisco, "Cisco Annual Internet Report (2018–2023) White Paper". [Online]. Available: https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-c11-741490.html

[113] N. Varsier, L.-A. Dufrène, M. Dumay, Q. Lampin and J. Schwoerer, "A 5G New Radio for Balanced and Mixed IoT Use Cases: Challenges and Key Enablers in FR1 Band," *IEEE Communications Magazine*, vol. 59, no. 4, pp. 82-87, April 2021.

[114] Y. Liu et al., "Evolution of NOMA Toward Next Generation Multiple Access (NGMA) for 6G," *IEEE Journal on Selected Areas in Communications*, vol. 40, no. 4, pp. 1037-1071, April 2022.

[115] L. Dai, B. Wang, Y. Yuan, S. Han, C. I and Z. Wang, "Non-orthogonal Multiple Access for 5G: Solutions, Challenges, Opportunities, and Future Research Trends," *IEEE Communications Magazine*, vol. 53, no. 9, pp. 74-81, September 2015.

[116] S. M. Riazul Islam, N. Avazov, O. A. Dobre and K.-S. Kwak, "Power-domain Non-Orthogonal Multiple Access (NOMA) in 5G Systems: Potentials and Challenges," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 2, pp. 721-742, October 2016.

[117] Z. Ding, X. Lei, G. K. Karagiannidis, R. Schober, J. Yuan, and V. J. Bhargava, "A Survey on Non-Orthogonal Multiple Access for 5G networks: Research Challenges and Future Trends," *IEEE Journal on Selected Areas in Communications*, vol. 35, no. 10, pp. 2181-2195, October 2017.

[118] B. Makki, K. Chitti, A. Behravan and M. Alouini, "A Survey of NOMA: Current Status and Open Research Challenges," *IEEE Open Journal of the Communications Society*, vol. 1, pp. 179-189, January 2020.

[119] K. Senel, H. V. Cheng, E. Björnson and E. G. Larsson, "What Role can NOMA Play in Massive MIMO?," *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 3, pp. 597-611, June 2019.

[120] B. Clerckx et al., "Is NOMA Efficient in Multi-Antenna Networks? A Critical Look at Next Generation Multiple Access Techniques," *IEEE Open Journal of the Communications Society*, vol. 2, pp. 1310-1343, 2021.

[121] G. Liu, Z. Wang, J. Hu, Z. Ding and P. Fan, "Cooperative NOMA Broadcasting/-Multicasting for Low-Latency and High-Reliability 5G Cellular V2X Communications," *IEEE Internet of Things Journal*, vol. 6, no. 5, pp. 7828-7838, Oct. 2019.

[122] J. Seo, B. C. Jung and H. Jin, "Performance Analysis of NOMA Random Access," *IEEE Communications Letters*, vol. 22, no. 11, pp. 2242-2245, Nov. 2018.

[123] M. Salehi, H. Tabassum and E. Hossain, "Meta Distribution of SIR in Large-Scale Uplink and Downlink NOMA Networks," *IEEE Transactions on Communications*, vol. 67, no. 4, pp. 3009-3025, April 2019.

[124] H. Tabassum, E. Hossain and J. Hossain, "Modeling and Analysis of Uplink Non-Orthogonal Multiple Access in Large-Scale Cellular Networks Using Poisson Cluster Processes," *IEEE Transactions on Communications*, vol. 65, no. 8, pp. 3555-3570, Aug. 2017.

[125] Z. Ding, R. Schober and H. V. Poor, "Unveiling the Importance of SIC in NOMA Systems—Part 1: State of the Art and Recent Findings," *IEEE Communications Letters*, vol. 24, no. 11, pp. 2373-2377, Nov. 2020.

[126] Z. Ding, R. Schober and H. V. Poor, "Unveiling the Importance of SIC in NOMA Systems—Part II: New Results and Future Directions," *IEEE Communications Letters*, vol. 24, no. 11, pp. 2378-2382, Nov. 2020.

[127] N. Zhang, J. Wang, G. Kang and Y. Liu, "Uplink Nonorthogonal Multiple Access in 5G Systems," *IEEE Communications Letters*, vol. 20, no. 3, pp. 458-461, March 2016.

[128] M. Al-Imari, P. Xiao, M. A. Imran, R. Tafazolli, "Uplink non-orthogonal multiple access for 5G wireless networks," *2014 11th International Symposium on Wireless Communications Systems (ISWCS)*, Barcelona, Spain, 2014.

[129] Y. Liu, M. Derakhshani and S. Lambotharan, "Outage Analysis and Power Allocation in Uplink Non-Orthogonal Multiple Access Systems," *IEEE Communications Letters*, vol. 22, no. 2, pp. 336-339, February 2018.

[130] Y. Gao, B. Xia, K. Xiao, Z. Chen, X. Li and S. Zhang, "Theoretical Analysis of the Dynamic Decode Ordering SIC Receiver for Uplink NOMA Systems," *IEEE Communications Letters*, vol. 21, no. 10, pp. 2246-2249, October 2017.

[131] Y. Gao, B. Xia, K. Xiao, Z. Chen, X. Li and S. Zhang, "Analysis of the Dynamic Ordered Decoding for Uplink NOMA systems with imperfect CSI," *IEEE Transactions on Vehicular Technology*, vol. 67, no. 7, pp. 6647-6651, July 2018.

[132] M. Salehi and E. Hossain, "On Coverage Probability in Uplink NOMA With Instantaneous Signal Power-Based User Ranking," *IEEE Wireless Communications Letters*, vol. 8, no. 6, pp. 1683-1687, Dec. 2019.

[133] A. Agarwal, R. Chaurasiya, S. Rai and A. K. Jagannatham, "Outage Probability Analysis for NOMA Downlink and Uplink Communication Systems With Generalized Fading Channels," *IEEE Access*, vol.8, pp. 220461-220481, December 2020.

[134] A. Alqahtani, E. Alsusa, A. Al-Dweik and M. Al-Jarrah, "Performance Analysis for Downlink NOMA Over α-μ Generalized Fading Channels," *IEEE Transactions on Vehicular Technology*, vol. 70, no. 7, pp. 6814-6825, July 2021.

[135] R.J. Vaughan and W.N. Venables, "Permanent Expressions for Order Statistics Densities," *Journal of the Royal Statistical Society*, vol. 34, no. 2, pp. 308-310, 1972.

[136] M.G. Kendall and A. Stuart, "The Advanced Theory of Statistics," 2nd edition, vol. 1, 1958.

[137] J. M. Holtzman, "A simple, accurate method to calculate spread-spectrum multiple-access error probabilities," *IEEE Transactions on Communications*, vol. 40, no. 3, pp. 461-464, March 1992.