Modelling and Enhanced Scheduling in the Cellular Vehicular Sidelink Standards

Brian McCarthy
BSc
114302146

Thesis submitted for the degree of Doctor of Philosophy

NATIONAL UNIVERSITY OF IRELAND, CORK

FACULTY OF SCIENCE

DEPARTMENT OF COMPUTER SCIENCE

January 2024

Head of Department: Prof Utz Roedig

Supervisor: Dr Aisling O'Driscoll

Advisor: Prof Cormac Sreenan

Contents

List of Figures

List of Tables

List of Publications

The following list contains peer reviewed publications that emanated from the thesis, as well as manuscripts currently under review.

Peer Reviewed:

1. **[Pub 1]** **B. McCarthy** and A. O'Driscoll, "*Adapting the Resource Reservation Interval for Improved Congestion Control in NR-V2X*", IEEE 23^{rd} International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM), Belfast, United Kingdom, 2022, pp. 261-270, doi: 10.1109/WoWMoM54355.2022.00045.
2. **[Pub 2]** **B. McCarthy** and A. O'Driscoll, "*Exploring NR-V2X Dynamic Grant Limitations for Aperiodic Traffic*", 8^{th} International Conference on Vehicle Technology and Intelligent Transport Systems (VEHITS), 2022, Vol. 1, pp. 17-25, doi: 10.5220/0010937500003191.
3. **[Pub 3]**[1] A. Burbano-Abril, **B. McCarthy**, M. Lopez-Guerrero, V. Rangel, and A. O'Driscoll, "*MCS Adaptation within the Cellular V2X Sidelink*", IEEE Conference on Standards for Communications and Networking (CSCN), Thessaloniki, Greece, 2021, pp. 111-117, doi: 10.1109/CSCN53733.2021.9686168.
4. **[Pub 4]** **B. McCarthy** and A. O'Driscoll, "*OpenCV2X Mode 4: A Simulation Extension for Cellular Vehicular Communication Networks*", IEEE 24^{th} International Workshop on Computer Aided Modeling and Design of Communication Links and Networks (CAMAD), Limassol, Cyprus, 2019, pp. 1-6, doi: 10.1109/CAMAD.2019.8858436.
5. **[Pub 5]** **B. McCarthy**, A. Burbano-Abril and A. O'Driscoll, "*MCS Adaptation and its role in improving MAC Scheduling for Variable Packet Sizes in NR-V2X*", 2023 6^{th} International Conference on Advanced Communication Technologies and Networking (CommNet), Rabat, Morocco, 2023, pp. 1-7, doi: 10.1109/CommNet60167.2023.10365302.
6. **[Pub 6]** **B. McCarthy** and A. O'Driscoll, "*RRI Adaptive: A Standards Compliant Approach for Equitable and Stable Congestion Control in C-V2X Networks*", 2023 6^{th} International Conference on Advanced Communication Technologies and Networking (CommNet), Rabat, Morocco, 2023, pp. 1-8, doi: 10.1109/CommNet60167.2023.10365283.

Under Review:

7. **[Pub 7]** **B. McCarthy** and A. O'Driscoll, "*Prediction of Inter-Packet Times for Enhanced NR-V2X Sidelink Scheduling*", submitted to the IEEE 35^{th} Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC), 2024.

Arxiv (not peer-reviewed):

8. **[Pub 8]** **B. McCarthy**, A. Burbano-Abril, V.R., Licea, and A. O'Driscoll, "*OpenCV2X: Modelling of the V2X cellular sidelink and performance evaluation for aperiodic traffic*", 2021 arXiv preprint arXiv:2103.13212.

[1]The main author was involved in the formulation of key paper ideas, the collection of the data set, interpretation of results and the writing of the paper. **B. McCarthy** provided the models and scenarios on which the results were generated, contributed to discussion of the key paper ideas and results analysis.

9. **[Pub 9]** B. McCarthy, and A. O'Driscoll, "*Prediction of inter packet arrival times for enhanced NR-V2X sidelink scheduling*", 2023 arXiv preprint arXiv:2311.08227.

I, Brian McCarthy, certify that this thesis is my own work and has not been submitted for another degree at University College Cork or elsewhere.

Brian McCarthy

In loving memory of Michael, Abby and Tim,

thoughts of them guided me through the challenging times.

Abstract

Vehicular networking can address significant challenges in the areas of vehicular safety, traffic efficiency, and enables higher levels of autonomous driving for sustainable mobility. The Third Generation Partnership Project (3GPP) Cellular Vehicle-To-Everything (C-V2X) standard and subsequent New Radio Vehicle-To-Everything (NR-V2X) standard are relatively new yet disruptive technology enablers in the regulatory landscape. Unfortunately, the relative infancy of the standards has resulted in shortcomings in their ability to effectively handle real-life vehicular service characteristics and network management techniques. Specifically, the scheduler that forms the basis of both C-V2X and NR-V2X cannot effectively schedule variable packet inter arrival rates, variable packet sizes or facilitate congestion control techniques.

This thesis will address these challenges by providing a rigorous analysis of the scheduling problems caused by these contemporary vehicular applications, emphasising the complexity posed by European Telecommunications Standards Institute (ETSI) and 3GPP standard application models. It further addresses the pressing issue of congestion management within diverse and dense vehicular scenarios, all while preserving scheduling performance. Specifically, this thesis proposes, evaluates, and analyses solutions based on machine learning enabled prediction of inter packet arrival times, scheduler compliant congestion control mechanisms, and a deep understanding of the role of MCS adaptation in V2X scheduling.

Acknowledgements

Firstly, I want to thank my family, Mom, Dad, Kate, Colm, and Peggy, as well as my Grandmother Ann. It has been a long road, and they have always been so enthusiastic even when the details were incomprehensible to them. I am very lucky to have them, and their support, patience, and great pride has always been a source of strength when the work was most difficult.

I must thank my supervisor, Dr. Aisling O'Driscoll. I do not believe there is another person who works harder than she does. I also believe that she has incredible patience, as I am most certain that I have caused her innumerable headaches. Her support throughout the process especially considering her own circumstances has been immense and for that I am truly appreciative.

I would like to thank some of the faculty in the Department of Computer Science specifically. First, Prof. Dirk Pesch and Dr. Md. Noor-A-Rahim, both provided guidance in the period when Aisling was on leave and their input in the project was of great help. Second, I would like to thank Prof. Cormac Sreenan, he provided guidance coming up to conferences and the viva itself and made the MISL lab a nice environment with regular coffee mornings. In addition, I would like to thank the wider computer science faculty and support staff. Everyone was always helpful and kind in my interactions with them.

I want to give a special thanks to Mary Noonan, again she provided great support, and I will miss starting the day with our chats.

I want to thank my MISL lab mates, in my time there has been a lot of change. The early days with Darijo Raca and our coffee mornings will be fondly remembered. My second batch of lab-mates have also been great friends. So thanks to Afshan, Evi, Japneet, Mohamed, Nasrin, Rahmin and Zohaib and I wish you all the best of luck with finishing your own work.

I also would like to acknowledge the support provided by the Science Foundation Ireland (SFI).

Finally, I want to thank Megan Ross for her love, support, and understanding throughout. She has been there through it all and always put the pieces back together if they ever fell apart, I will always love her for it and I consider myself lucky to have her in my life.

Chapter 1

Introduction

1.1 Background

Vehicular networking is the application of networking technologies to a transportation scenario, enabling vehicles to communicate with each other and other entities around them. Cooperative, Connected and Automated Mobility (CCAM) is expected to transform the transport sector, with enhanced safety, traffic efficiency, and increased automation. As the time to market for widespread deployment of vehicles exhibiting higher levels of autonomy can be expected to be lengthy, CCAM systems have initially focused on connected human-driven vehicles. In such scenarios, warnings and updates are presented to the driver via an integrated On-Board Unit (OBU), which have been learnt cooperatively from other vehicles (Vehicle-To-Vehicle (V2V)) or infrastructure (Vehicle-To-Infrastructure (V2I)). When faced with higher levels of on-board automation, it can still be expected that leveraging connectivity will be a vital enabler for removing the driver from the equation. On-board sensors are limited to Line-Of-Sight (LOS) scenarios and exhibit drawbacks under a variety of weather and environmental conditions. Connectivity, in-conjunction with sensing, is expected to be the key enabler to realise the widespread and reliable deployment of autonomous vehicles.

The EU has set a goal to halve all fatalities in road traffic accidents by 2030 and completely eliminate them by 2050 [4]. In 2022, the Irish national level of road fatalities increased by 18 deaths compared to the previous year, with a total of 155 people killed in road traffic accidents [5]. A particular challenge exists to ensure the safety of Vulnerable Road User, e.g. pedestrians, cyclists,

etc., with 36 fatalities recorded in 2022. Existing Advanced Driver Assistance Systems (ADAS), designed to address such safety concerns, are wholly sensor-based, which reduces their effectiveness when operating in scenarios with Non-Line Of Sight (NLOS) or to adequately respond in time given limited sensing range. The application of connectivity can enable both NLOS scenarios and increase the "sensing" range significantly.

Additionally, there are significant challenges in achieving enhanced traffic efficiency, including a target reduction in carbon emissions. The European Commission has significant climate change goals with the ambition to reduce greenhouse gas emissions to at least 55% below 1990 levels by 2030 [6]. For Ireland to meet its climate obligations by 2030 road traffic emissions must be reduced by 50% [7]. This is a policy goal for the national strategy of Project Ireland 2040, with a focus on sustainable transport (NSO 4) and also relates to (NS0 8) focused on a low carbon society [8]. The current generation of ADAS technologies will not be able to provide significant improvements in traffic efficiency on a scale that will affect emissions, due to the fact that they cannot act collaboratively or predict traffic scenarios such as traffic light change times. V2X technologies can improve efficiency by enabling vehicles to operate in a more cooperative manner; a key application in this area is vehicle platooning. They can also improve the efficiency of traffic lights through applications such as Green Light Optimised Speed Advisory (GLOSA). Furthermore, connectivity can also allow vehicles to access the core network, edge, and cloud platforms. Access to greater levels of compute, as well as the wider Internet, will enable many more advanced use-cases. Significant scope exists to develop digital twins as a result of real-time recording of information about traffic flows which will in turn enable better city planning.

The basis for many of these cooperative services and future advanced use-cases is Cooperative Awareness (CA), Decentralised Environmental Notification (DEN) and the more recently defined Cooperative Perception (CP) services. These services enable vehicles to understand the dynamics of their neighbours (CA), be alerted to any significant events, such as a vehicle breaking in an emergency (DEN), and finally, cooperative communication of sensor data to more comprehensively understand their surroundings such as the location of Vulnerable Road Users (VRUs) and other environmental objects (CP). These services form the fundamental basis on which increased autonomy is built and must provide highly reliable performance in diverse vehicular environments. Overcoming issues with constrained channel resources, diverse

application requirements and dense user environments is vital.

1.2 Research Motivation

As described in Section 1.1, the effectiveness and reliability of vehicular technologies such as Cellular Vehicle-To-Everything (C-V2X) and New Radio Vehicle-To-Everything (NR-V2X) is vital when deployed in dense and diverse vehicular scenarios. This is difficult to assess in real-world deployments as it is cost prohibitive and has safety implications. As such, there needs to be a low cost yet rigorous environment to test these technologies without the need for large-scale physical testing. Given the relative infancy of C-V2X and NR-V2X these environments were not commonly available.

The first challenge of this work was to implement a simulation model of the C-V2X and NR-V2X standards so that their communication characteristics, particularly related to MAC layer scheduling, could be thoroughly analysed and these models made available to the wider vehicular networking community for comprehensive and reproducible performance analysis of the standards.

The second challenge is to provide the effective scheduling of the current and future generation of V2X applications. For C-V2X and NR-V2X to be reliable as they become more widespread, it is necessary to operate effectively with the European Telecommunications Standards Institute (ETSI) standards of applications. Thus this work provides a rigorous exploration of their performance and ultimately proposes effective scheduling mechanisms for managing these applications, that exhibit event-driven characteristics, with a specific focus specifically on CA and DEN applications. An extensive analysis of MCS adaptation to manage the variable packet sizes exhibited by the CA services as well as many other vehicular services is provided, and recommendations made. This is an aspect often overlooked in the literature that focuses on fixed MCS which is not how the standard is specified.

The third and final research challenge addressed in this thesis relates to the limited availability of spectrum and the potentially dense characteristics of vehicular networks that can lead to congested radio conditions. Given limited dedicated Cooperative Intelligent Transport Systems (C-ITS) frequencies, along with a proliferation of existing services and a plethora of future envisaged vehicular services, congestion within these bands is envisaged. To

overcome this, congestion control mechanisms are proposed which can effectively manage these limited spectrum scenarios. These solutions ensure the correct maintenance of the underlying scheduling mechanism in C-V2X and NR-V2X, while reducing congestion in the channel, enabling better network performance. A key aspect that has been overlooked in other studies on congestion control in C-V2X/NR-V2X is the fairness (of access among vehicles) and the stability (of channel conditions) of any proposed solution. As such, this thesis provides a complete performance analysis of the shortcomings of the current ETSI standard mechanisms when applied to C-V2X and proposes solutions, that ultimately show improved network performance while maintaining the same or higher levels of fairness and stability, and remaining standards compliant.

1.3 Thesis contributions

This thesis has 4 major contributions [**C1-C4**] and 1 minor contribution [**C5**]. This is summarised in Fig. 1.1 and is now described.

- **Contribution 1** [**C1**]: (*Pub 4, Pub 7*)
  Implementation of an open source full-stack standards compliant C-V2X and NR-V2X model, that has been fully validated and is in use by the wider V2X academic community. This is captured in **Chapter 3**.
- **Contribution 2** [**C2**]: (*Pub 2 and poses the underlying research challenge addressed in Pub 1 & Pubs 5, 6, 8*)
  Thorough investigation of the performance issues faced by Sensing-Based Semi-Persistent Scheduling (SB-SPS) when handling commonly occuring vehicular service characteristics (aperiodic packet inter-arrival rates and variable packet sizes) and network management techniques (congestion control). This includes:
  - An in-depth theorised analysis of the challenges posed to the C-V2X and NR-V2X SB-SPS by Missed Transmissions (MTs) and Grant Breaks (GBs). This is captured in **Sections 4.2.1** and **4.2.2**.

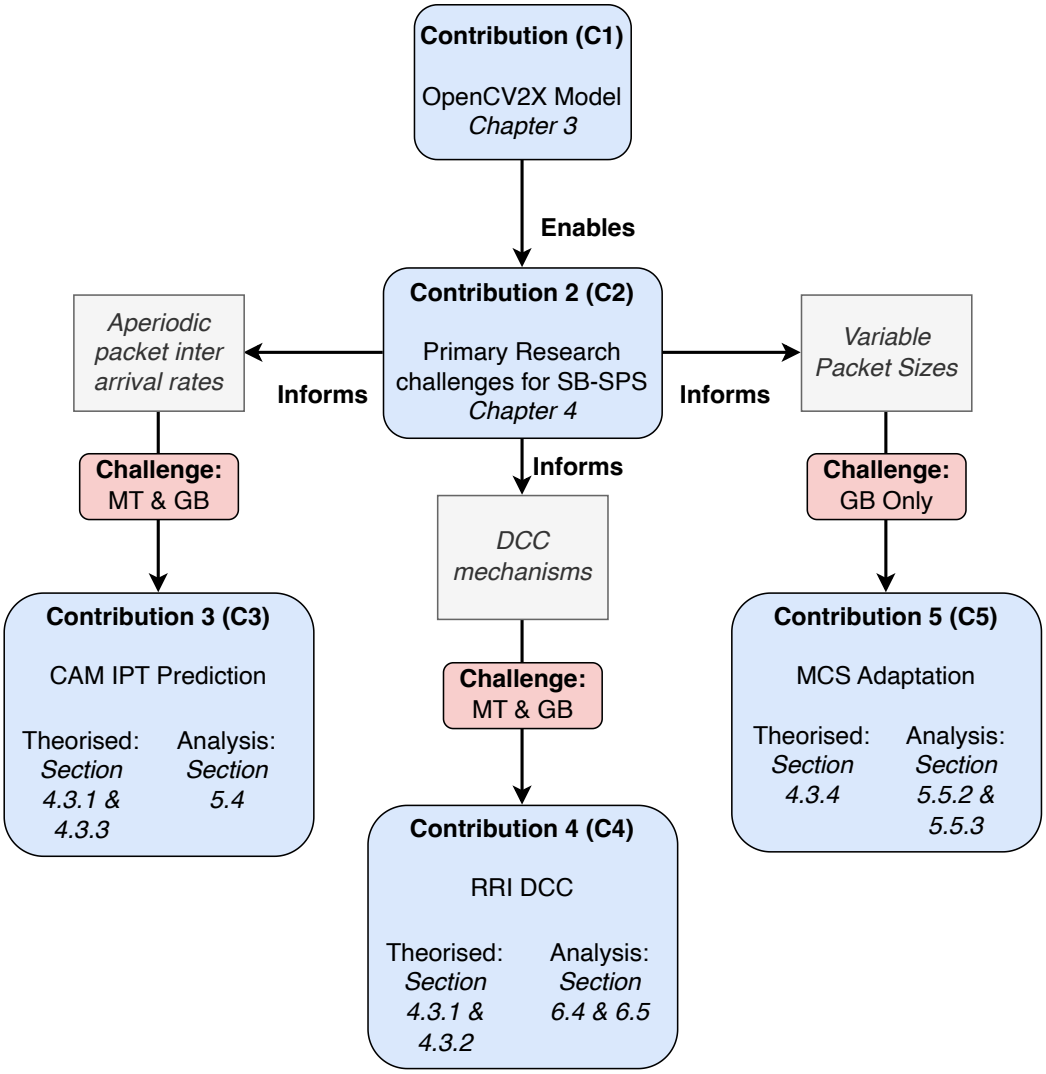

Figure 1.1: Contributions of this thesis

  - A detailed quantitative analysis of the impact of both MT and GB on SB-SPS scheduling performance with a brief quantitative overview provided in **Section 4.2.3** and a detailed analysis of the consequences of a) aperiodic traffic patterns in **Section 5.3**, b) variable packet sizes in **Section 5.5** and c) congestion control mechanisms in **Section 6.3**.

- **Contribution 3** [**C3**]: *(Manuscript under review)*
  A machine learning based approach to improved scheduling of aperiodic packet inter-arrival rates (CAMs) within SB-SPS to overcome MT and GB. This is proposed in **Section 4.3.1 and 4.3.3** and quantitatively evaluated in **Section 5.4**.

- **Contribution 4** [**C4**]: *(Pub 1, Pub 3, Pub 6)*
  Proposal of congestion control mechanisms that are compliant with SB-SPS including a quantitative evaluation of their effectiveness. This comprises:
  - A Resource Reservation Interval Decentralised Congestion Control (RRI DCC) mechanism that reduces congestion while maintaining effective scheduling. This further investigates fairness amongst vehicles and stability of channel conditions, which has yet to be explored in the C-V2X/NR-V2X context. This is proposed in **Sections 4.3.1** and **4.3.2** and quantitatively evaluated in **Section 6.5**.
  - An analysis of the effectiveness of MCS adaptation as a means of congestion control, which was found to be ineffective. This is captured in **Section 6.4**.
- **Contribution 5** [**C5**]: *(Pub 5)*
  This minor contribution analyses the role of MCS adaptation in assisting SB-SPS to effectively schedule variable sized packets. This is described in **Section** 4.3.4. This includes:
  - A quantitative analysis of the performance implications of MCS adaptation versus fixed MCS selection (widely assumed in academic literature). This is captured in **Section 5.5.2**.
  - A quantitative evaluation of the highlighting the importance to consider MCS ranges beyond those defined in the existing standards. This is also captured in **Section 5.5.2**.
  - An alternative MCS selection strategy to minimise SB-SPS grant breaks. This is captured in **Section 5.5.3**.

1.4 Thesis structure

The remainder of the thesis is organised as follows:

- **Chapter 2** has two distinct sections. It firstly presents the relevant background related to C-V2X and NR-V2X, their operation, and standardised techniques for addressing their shortcomings. Secondly, a review of related academic literature is provided with a specific focus on the research challenges addressed in this thesis.

- **Chapter 3** describes an open source full-stack model (OpenCV2X), utilised for the quantitative analysis in this thesis and made available to the wider vehicular communications community. We describe in detail its requirements, core components and implementation details. A full validation is provided, bench-marked against an analytical study.
- **Chapter 4** explores the challenges faced by the challenges of SB-SPS within C-V2X and NR-V2X, when faced with certain vehicular service characteristics (aperiodic inter packet arrival rates and variable packet sizes) as well as network management techniques (congestion control). A brief quantitative study of their impact is also provided. This analysis will form the basis for understanding the research challenges that we aim to address in our contributions. Finally, a theorised discussion of the proposed solutions is provided, specifically addressing how these solutions are applied to aperiodic inter packet arrival rates, variable packet sizes and congestion control techniques. Their implementation is also discussed.
- **Chapter 5** provides a quantitative evaluation of the negative impact of aperiodic inter packet arrival rates and variable packet sizes on SB-SPS and various techniques to address this. A novel machine learning-based prediction approach that intelligently maintains the SB-SPS grant mechanism is evaluated. Finally, we provide an exhaustive analysis of the adaptation of MCS for handling variable packet sizes in both technologies.
- **Chapter 6** provides an in-depth analysis of congestion control within C-V2X and NR-V2X. Specifically, it proposes and analyses the performance of a range of RRI DCC mechanisms that are standards compliant yet effectively maintain the underlying scheduling mechanism. This includes a study on fairness and stability. These mechanism overcome the incompatibility issues with current standardised congestion control mechanisms investigated in Chapter 4. Finally, an in-depth analysis of MCS adaptation as a means of congestion control is provided. The quantitative analysis finds such an approach to be ineffective.
- **Chapter 7** presents conclusions and future outlook.

Chapter 2

Background and Related Work

2.1 Introduction

This chapter has two distinct sections. Firstly the technical background of the vehicular technologies in this thesis is described. Specifically, an overview of the following is given:

- Regulatory spectrum allocations for C-ITS applications.
- Key vehicular services that are fundamental to C-ITS and that are evaluated in this thesis.
- C-V2X and NR-V2X operation including channelisation, sensing and scheduling and transmission management as well as their key differences.
- Dynamic Grant overview including dedicated aperiodic scheduling mechanisms.
- MCS adaptation with a focus on variable packet scheduling.
- Standardised congestion control mechanisms for vehicular environments.

Secondly, a review of the state of the art literature related to the research challenges is provided. Specifically, this includes:

- Performance analysis of C-V2X and NR-V2X.
- Research that addresses aperiodic inter-packet arrival times.
- Research considering MCS adaptation and variable packet sizes.

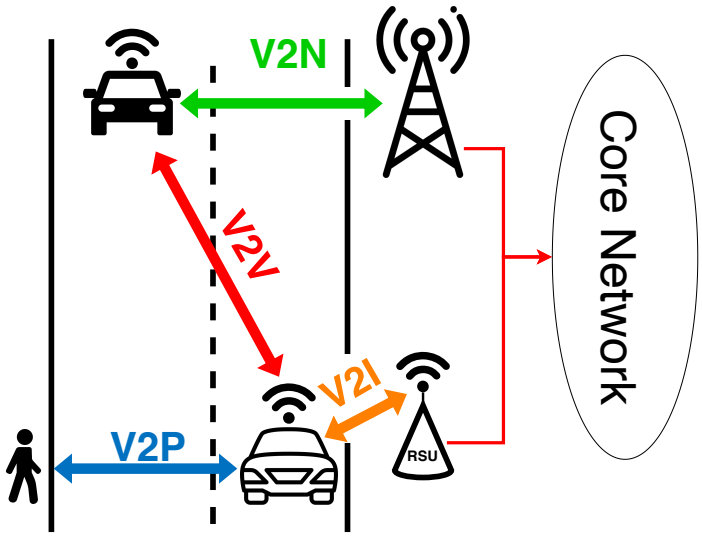

Figure 2.1: V2X communication patterns

- Congestion control within C-V2X and NR-V2X.
- Trajectory prediction and its application to CAM Inter-Packet Time (IPT) prediction.
- Vehicular simulation environments.

2.2 Overview of vehicular networking

Vehicular networks are the application of network technologies in a vehicular scenario. The goal is to allow the vehicle to communicate with neighbouring vehicles, roadside infrastructure, and even vulnerable road users, such as pedestrians. This can be broadly categorised as Vehicle-To-Everything (V2X). The goal of vehicular networks is primarily safety, i.e. reducing the possibility of collisions, and secondarily road traffic efficiency and finally infotainment services. In general, safety applications will involve Vehicle-To-Vehicle (V2V) or Vehicle-To-Pedestrian (V2P) communication. Traffic efficiency may include V2V, but also Vehicle-To-Infrastructure (V2I) communication. Finally, infotainment services will require Vehicle-To-Network (V2N) communication for services such as video streaming. All these patterns are captured in Fig. 2.1.

The initial deployment of these services is called "Day 1" services; these are currently in early stage deployment in projects such as C-Roads in the EU [9]. These day 1 services focus primarily on the initial two areas of concern, that is, safety and traffic efficiency. The core element of Day 1 is the CA service, which

sends Cooperative Awareness Messages (CAMs). This application is the main safety application and describes the characteristics and dynamics of the vehicles with neighbouring vehicles to avoid collisions. It will become an integral part of ADAS systems as deployments expand with increasing penetration rates. This will be the main application that is the focus of this thesis, although additionally the DEN service is a concern and key to safety being used for emergency events such as emergency braking or emergency vehicle approaching.

The first vehicular networking technology to be proposed was based on the 802.11p wireless protocol. This evolved into the ITS-G5 standard proposed by ETSI in Europe [10] and WAVE in North America [11]. This is a thoroughly understood technology deep into the standardisation process, though deployments have lagged behind. This led to 3GPP developing a competing standard in LTE-V2X or C-V2X, a cellular vehicular specific solution in 3GPP release 14 [12]. This was the initial definition before its more recent evolution into the NR-V2X standard which is based on C-V2X [13]. This 3GPP standard, with support from the 5G Automotive Association (5GAA), has gained significant momentum and claims to overcome many of the drawbacks associated with channel congestion that have prevented 802.11p variants from being mass-market. Release 14 utilises a new sidelink radio interface, PC5, for direct V2V communication based on the LTE ProSe Device-to-Device (D2D) Release 12 standard [14]. It further specifies two new communication modes, Mode 3 and Mode 4 as well as two new architectural entities. In Mode 3 (scheduled), the cellular infrastructure allocates and manages the resources necessary for V2V PC5 communications, via the Uu interface. In contrast, in Mode 4 (autonomous), each vehicle selects its radio resources for V2V communications using a distributed scheduling algorithm, SB-SPS. Therefore, in Mode 4 vehicles can operate without infrastructure, analogous to wireless ad hoc networks. This is considered the baseline mode, since safety applications cannot depend on the availability of cellular coverage and is a direct alternative to DSRC/WAVE or ITS-G5. As such, this study will focus exclusively on the Mode 4 autonomous resource selection and the later NR-V2X mode 2 standard, which is built on it.

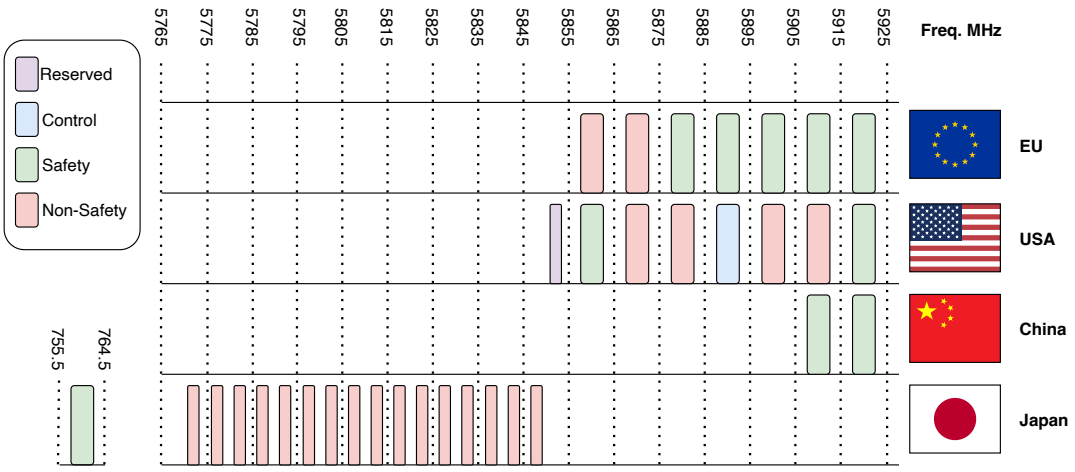

Figure 2.2: Spectrum Allocations around the globe

2.2.1 Regulatory aspects including spectrum allocation

A key concern and major area of regulatory activity with any wireless technology is the allocation of spectrum to the specific task. Intelligent Transport System (ITS) spectrum allocations are diverse throughout the world, the key aspect is that both ITS-G5, C-V2X and NR-V2X are all sub-6 GHz technologies designed for this spectrum range. Initial spectrum allocations would have considered ITS-G5 as the technology of choice, which is why the ranges of the allocated spectrum are allocated to different tasks or functional areas in certain jurisdictions. ITS-G5 allows for multiple channels generally of 10 MHz which can be dedicated to specific tasks or application types, and vehicles can decide to listen in these sub-sections of the spectrum if an application is of interest to them. C-V2X and NR-V2X on the other hand do not operate in quite the same way, the channel can range from 10 to 20 MHz for a specific sidelink configuration.

Ultimately, these two technologies have multiple outcomes which are up to regulators to determine, meaning will they co-exist, or will one prevail and have the spectrum allocated directly to that technology. These questions are beyond the scope of this work, but what is worthwhile considering here is that, in all cases, the spectrum is limited and that poses challenges to any technology adopted. The allocations of spectrum in the EU, USA, China and Japan can be seen in Fig. 2.2. Focusing on the EU, there are 70 MHz of spectrum dedicated to ITS applications, split across 4 categories of use. Considering for a moment only the allocation for safety, the EU allows 20 MHz of spectrum for this. This is actually the maximum allowed for C-V2X to

manage on the sidelink in one contiguous segment, and this ultimately ties into aspects with respect to congestion. 20 MHz, while significant, is not limitless, as this study will show, in 10 MHz even a typical application with a 10 Hz rate and a congested scenario can cause significant levels of congestion in the channel. Considering the number of applications dedicated to safety, CA, CP and DEN, with any additional new services applied on top of this. Additionally, it is important to consider the diversity of vehicular scenarios that may exhibit very high densities in heavy traffic.

2.2.2 Key vehicular services

C-V2X and the more recent NR-V2X are both technologies which have been envisaged to deliver on "Day 1" services; these services focus on the first generation of vehicular communication standards and applications. As such, their focus is primarily on safety and traffic efficiency applications with a best effort model of performance. In the EU, the main agency that standardises these applications is ETSI. There is a very diverse range of standardised ETSI applications, but the primary applications of concern in this study are CA and DEN services, whose messages are CAMs and DENMs.

Taking CA as the first primary service, it is the one that will form the core traffic in any V2X network for all "Day 1" deployments and is standardised in [15]. CA is concerned with vehicles communicating their kinematics to neighbouring vehicles for collision avoidance, as such CAMs contain information about the position, speed, heading, and dimensions of the vehicle. The key when considering the CA service and SB-SPS is the fact that its messages are triggered based on vehicle dynamics. The specific trigger conditions are as follows:

- Change in speed of 0.5 m/s.
- Change in heading of 4°.
- Change in position of 4 m.
- The elapsed time since the last CAM transmission exceeds 1s.

This will be shown later in this study to be a key concern with the SB-SPS standard, as the core "Day 1" application exhibits an aperiodic pattern because of it being designed to trigger based on dynamics, which ultimately makes the pattern irregular.

The other service considered in this study is DEN, in this case DEN services are much more diverse with differing requirements [16]. In general, DEN services are designed for irregular or emergency information, the primary example being vehicle emergency breakdown or an approaching emergency vehicle. Additionally, other more diverse applications have been dedicated to DEN such as weather advisory or roadwork warnings, which is an application for which they are designed on the C-Roads platform. The result of these diverse applications means that they also have diverse requirements from a transmission perspective. Although generally it can be considered that again DEN services are event-triggered such as emergency breaking and may not have a strict pattern, which is why this study discusses them extensively given the strict periodisation of SB-SPS.

2.3 Background: Scheduling in C-V2X & NR-V2X

C-V2X Mode 4, as specified by 3GPP, employs a sensing-based semi-persistent scheduling algorithm for resource allocation at its MAC and PHY layers. This process can be broken into distinct phases of Selection & Sensing, and Transmission. However, it is necessary to understand first how the channel is structured and the information used in scheduling.

2.3.1 Channelisation & Scheduling information

The channel of C-V2X is based on Single-Carrier Frequency Division Multiple Access (SC-FDMA). In the time domain, resources are organised into subframes of 1 ms, which are further grouped into frames of 10 ms. Each subframe contains 14 SC-FDMA symbols, 4 symbols are used for demodulation reference signals, 1 symbol for Tx-Rx switching, and the remaining 9 symbols are left for data transmissions.

In the frequency domain, the channel is divided into subcarriers of 15 kHz. These subcarriers are grouped into Resource Blocks (RBs), each RB containing 12 subcarriers and spanning over 1 subframe. Unlike the conventional LTE resource structure, C-V2X groups RBs into subchannels. The number of RBs per subchannel and the number of subchannels are configurable but limited by the allocated bandwidth, which can be 10 or 20 MHz.

This is shown in fig. 2.3, the subframe labelled in blue, the subchannels

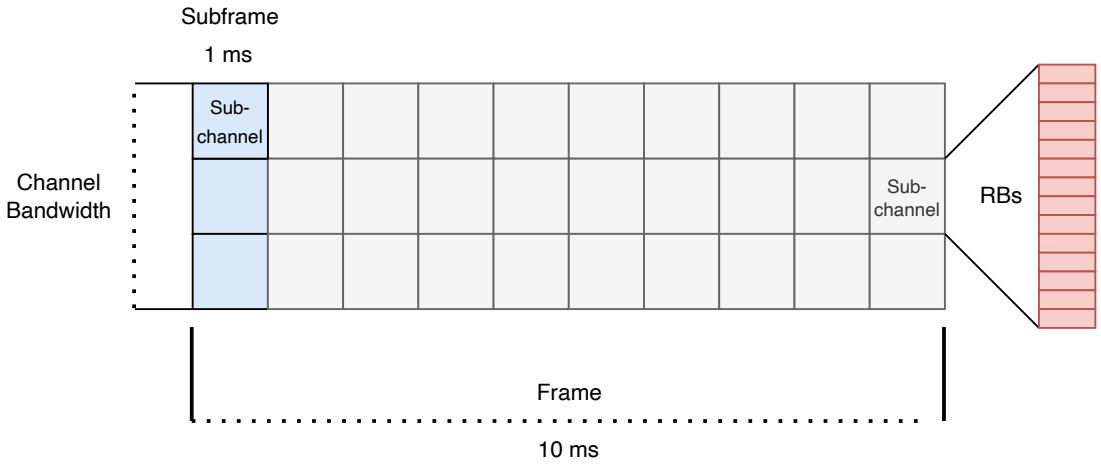

Figure 2.3: C-V2X Channelisation

represented in the grey squares, and the RBs that make up the subchannels in red.

The channel is also broken into the Physical Sidelink Control Channel (PSCCH) and the Physical Sidelink Shared Channel (PSSCH), which are made up of RBs. The PSSCH transmits application layer/data packets, known as Transport Blocks (TBs). The PSCCH is used to transmit scheduling control information known as Sidelink Control Information (SCIs). The PSCCH and PSSCH can both have individual sections of bandwidth dedicated to them (non-adjacent mode), or they can share the whole bandwidth interchangeably (adjacent mode); this is represented in Fig. 2.4. As shown in Fig. 2.4a when in adjacent mode the PSCCH is incorporated within each subchannel alongside the PSSCH. This mode has the advantage that in the case where multiple subchannels are assigned, only two RBs are required for SCI transmission, and as such the additional PSCCH RBs in the second subchannel can be allocated to PSSCH in this case. The non-adjacent configuration is shown in Fig. 2.4b, in this mode a specific segment of the channel is allocated to PSCCH and the remainder is PSSCH; this dedicated area cannot then be allocated to PSSCH when multiple subchannels are reserved. Throughout this study, adjacency is the configuration adopted. The structure of SCI Format 1 is illustrated in Table 2.1. SCI Format 1 indicates the priority of the packet (denoted by Priority field), the time period of the intended next use of the resource (Resource Reservation field), the number and location of the resources for the initial transmission and retransmission, the retransmission time (Time gap between initial transmission and the retransmission), the MCS and, finally, whether or

not a retransmission will occur (Retransmission index field). An SCI is sent in the same time slot as a TB and occupies 2 RBs. A TB cannot be decoded if the associated SCI is not received because the MCS is specified within the SCI.

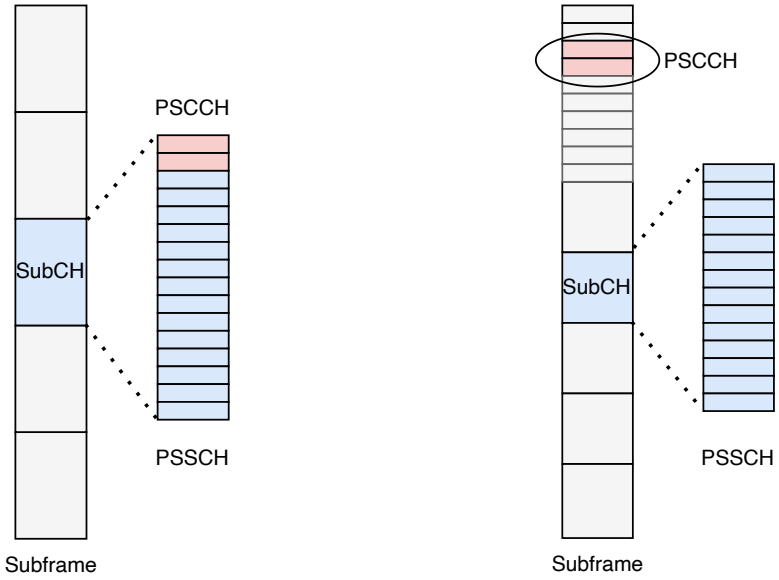

(a) Adjacent. (b) Non-adjacent.

Figure 2.4: PSCCH & PSSCH Adjacency.

Table 2.1: 3GPP SCI Fields

Priority	Resource Reservation	Frequency Resource location of initial and retransmission	Time gap between initial and Frequency Resource location of initial and retransmission	Modulation and coding scheme	Retransmission index	Reserved information bits
3 bits	4 bits	Calculation based on the number of subchannels	4 bits	5 bits	1 bit	add to 32 bit total

2.3.2 Selection & Sensing - MAC & PHY Layer

At the MAC layer, C-V2X implements SB-SPS, dubbed Mode 4, to allow vehicles to select resources autonomously. The process starts with the reception of a packet from the upper layers. Upon reception, the MAC layer creates a scheduling grant containing the number of subchannels and the periodicity between transmissions. If a grant has already been created at the time a packet is received from the upper layers, the transmission is scheduled using the existing grant. The number of subchannels depends on the channel configuration and the application requirements, specifically the current packet size and the configured MCS available. The periodicity between transmissions is defined by the RRI, whose value is set by the upper layers and will be a key component discussed in this study. The RRI in general ranges from 100 to

1000 ms in 100 ms intervals, but is allowed to be 25 ms and 50 ms for some V2P applications [17]; these are not considered in the current study but are implemented in the OpenCV2X model.

The grant is then passed to the PHY layer, which generates a list of all subchannels that meet the grant specifications. These subchannels are known as Candidate Single-Subframe Resources (CSRs) and consist of one or multiple subchannels in the same subframe. The list contains all the CSRs within a selection window that comprises the period between the time the packet is received from the upper layers and the maximum allowed latency defined the upper layer.

The first stage in filtering all possible CSRs is based on the Half-Duplex (HD) challenge for C-V2X. The antennas employed by C-V2X and NR-V2X cannot receive and transmit simultaneously, which results in the fact that on transmission a vehicle has no ability to determine if a collision occurs or if the other subchannels in the subframe are occupied. As such, the SB-SPS algorithm will discard any subframe in the selection window which corresponds to a subframe in the sensing window in which the vehicle transmitted, this is dependent on the configured RRIs.

The list is then filtered using the information in the SCIs received during a sensing window comprised of the last 1000 subframes. On the basis of this information, CSRs are excluded if the SCI indicates the CSR will be reserved during the upcoming selection window and if the average PSSCH Reference Signal Received Power (RSRP) of the CSR exceeds a predefined threshold. RSRP is a measure of the receiver power of a transmission. After excluding all CSRs that meet these two conditions, at least 20% of all CSRs should remain available. If this is not the case, the process is repeated by increasing the RSRP threshold by 3dB. Finally, the PHY selects the 20% of CSRs with the lowest Sidelink Reference Signal Strength Indicator (RSSI) averaged over the sensing window. RSSI is a measure which indicates the total noise on a subchannel, being the sum of RSRP, interference, and noise on the resource. This ensures that the CSRs with the lowest levels of interference are considered for selection.

The grant process is explored in Fig. 2.5, to explore a concrete example of this in operation, at time T_0 the MAC layer requests a set of CSRs from the PHY layer, say for a single subchannel. In this case, the PHY layer generates a selection window with CSRs labelled A through L, assuming that this is the

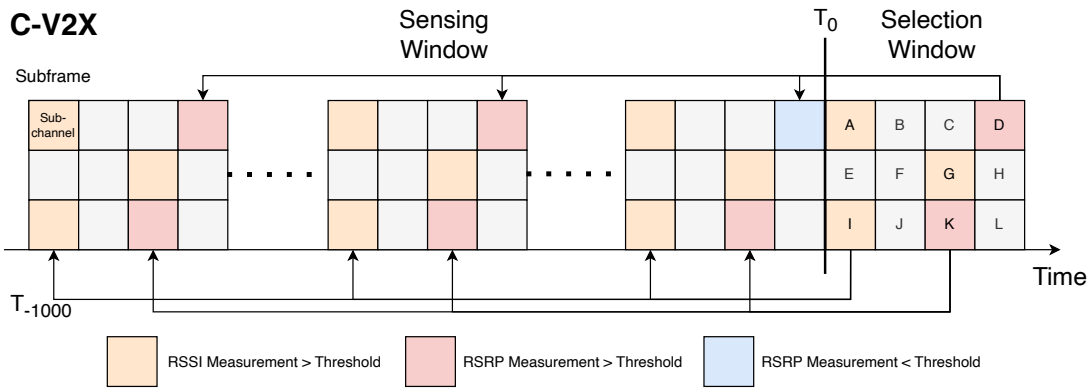

Figure 2.5: C-V2X SB-SPS mechanism

first transmission, then no resources are discarded due to HD. Then there is the SCI and RSRP filtering stage in which resources D and K (red) are filtered out as a result of their average PSSCH RSRP exceeding the threshold amount. Then the SB-SPS mechanism moves to the RSSI filtering stage, in this case, resource A, G, and I (orange) are filtered out due to high recorded RSSI on these subchannels. As a result, the MAC layer will have resources B, C, E, F, H, J, and L to select from.

The remaining CSRs are passed to the MAC layer, where a single CSR is selected at random to reduce the probability that multiple vehicles choose the same CSR. The CSR is selected for a number of recurrent transmissions defined as the Resource Reselection Counter (RC), which is set by randomly selecting an integer value between 5 and 15. The RC, whose value is decremented by one after each transmission. When the RC value reaches zero, each vehicle can maintain the same reservation with probability P or generate a new grant and restart the selection process. The value of P has been configured previously and can take any value between [0,0.8]. This process is represented in Fig. 2.6 where the F subchannel (green) has been selected and the key parameter here is the RRI, which controls the time period between transmissions, allowing the SB-SPS mechanism to maintain an accurate history of which resources are selected by neighbouring vehicles.

It is of note that the RC can additionally be for 25 - 75 and 15 - 30 in the case of RRIs below the 100 ms interval, though this is not considered in this study.

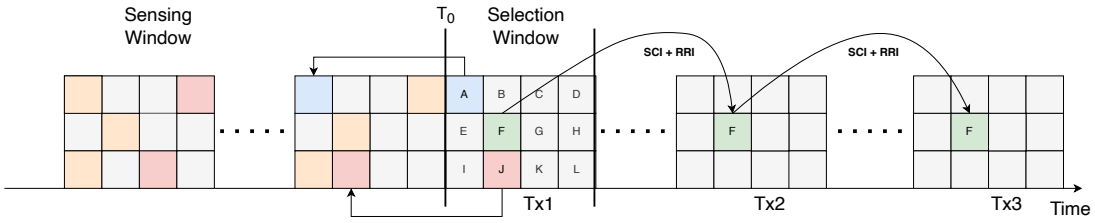

Figure 2.6: SB-SPS Grant mechanism

2.3.3 Packet Transmission

When a TB is ready to be transmitted, the MAC layer must determine the MCS to use. An MCS is selected among the configured range, e.g. should the highest MCS not be able to accommodate the resource, then the resource is relinquished and re-selection of a new resource must occur. Otherwise, on selection of an adequate MCS, the TB is then sent to the PHY layer with the necessary information to create an SCI message. The PHY layer subsequently sending both the TB and SCI message in the same subframe. At the MAC layer, the counter to maintain the reservation is decremented. If the counter is reduced to 1, the MAC layer attempts to maintain the reservation. This is done with probability
$P(X > probResourceKeep)$ with $probResourceKeep \in [0, 0.8]$, otherwise the resource is given up. If the resource is relinquished, then upon transmission of the final TB the SCI informs other users that the resource is available by setting the RRI to 0.

The SCI is the key element of the SB-SPS mechanism as outlined in Table 2.1, it is the message that declares what resources or subchannels the TB will actually occupy, additionally, since it provides the MCS, it is impossible to decode the TB without decoding the SCI first. One of the fields this study will focus on most is the RRI field; this controls the period between transmissions and is the means by which the PHY layer is able to determine which resources in the selection window are, in fact, reserved by neighbouring vehicles.

2.3.4 Key changes from C-V2X to NR-V2X

Generally, 3GPP Rel. 16 NR-V2X Mode 2 maintains much the same approach as described for Mode 4. However, some changes have been introduced at the PHY layer. This includes flexible numerology to allow for reduced subcarrier spacing, enabling low latency transmission. Additional channels have also been specified to enable groupcast and unicast communications which were

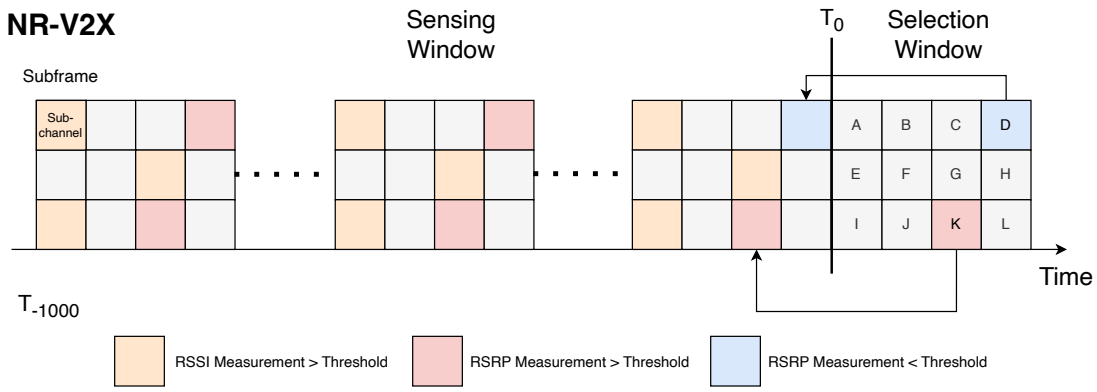

Figure 2.7: NR-V2X SB-SPS

not available for Mode 4. This introduces a multistage SCI format to enable these communication patterns. More information related to the NR-V2X PHY changes can be found in [18].

At the MAC layer, we highlight only the changes that are significant with respect to impacting SB-SPS scheduling performance. The first relates to the RSRP filtering stage; rather than using an average RSRP value for the CSR in the sensing history as is the case in Mode 4, Mode 2 instead uses only the RSRP of the most recent transmission when determining reserved resources. The second key change is the outright removal of the RSSI filtering stage. As such, all CSRs not removed at the RSRP filtering stage will be reported to the MAC layer for selection. These modifications are important, as shown in [19], as they significantly impact the performance of SB-SPS when handling aperiodic traffic.

Fig. 2.7 shows how this variant of SB-SPS operates with a concrete example that mimics what has previously been shown to provide a direct comparison between the two variants. Again, at T_0 the MAC layer requests a set of CSRs for a single subchannel and a selection window containing subchannels A through L is generated, based on application latency requirements. In this case again assuming it is the first transmission, then no HD subframes need be discarded, though in both technologies this approach is applied the same way, as such the outcome would be the same. The second filtering stage now takes place, and the key difference here is that rather than an average RSRP measurement, resources are selected based only on the most recent reservation. This is represented in resources D and K, in the case of resource D which considering Fig. 2.5 was removed as a result of RSRP averaging, in this case as represented by the blue subchannels the RSRP is below the threshold. This change results

in that only resource K is filtered from the CSRs at this stage. Now importantly the RSSI filtering stage is removed entirely and as such all the remaining CSRs are actually available for selection and thus returned to the MAC layer. This change results in both a larger pool of CSRs to select from, but also resources A, G, and I now being available regardless of high recorded RSSI.

There are additional mechanisms introduced in NR-V2X which are not directly explored in this study, those specifically are the re-evaluation mechanism and the resource pre-emption mechanisms.

In the first case re-evaluation is a process under which after selecting a specific CSRs but before its first use a vehicle will still be able to defer its use if it senses another vehicle which has reserved the resource in its place and its use would result in a collision. This is an important introduction for NR-V2X as a result of the fact that the RRI mechanism has been adapted in the variant, allowing for sub-100ms RRIs that are configurable in the range 1-50 ms, not the fixed 100, 200, 300 ms which C-V2X allows.

The second element pre-emption allows for vehicles to defer the use of a selected CSR in the case where the vehicle senses that a neighbouring vehicle which is operating a higher priority service will require that resource for its use. It should be noted that this study has a limited focus on packet prioritisation.

2.3.5 Aperiodic scheduling techniques

NR-V2X made significant efforts to improve aperiodic scheduling performance in the latest standard, and an important element of that was the introduction of a dynamic grant scheduling mechanism. When referring to aperiodic patterns, specifically, these are application patterns which exhibit an unpredictable packet inter-arrival time. A specific example explored in Section 2.2.2 is the ETSI CA service, where packets are triggered based on vehicle dynamics. It operates by using the same underlying SB-SPS algorithm as described for NR-V2X, but assuming a single transmission. The SCI message for this transmission will highlight that the resource will not be maintained in the future. This does not apply when retransmissions are enabled, in which case a grant is maintained for the said retransmissions. We do not consider retransmissions in this study. This mechanism will be investigated in detail in this study, but two additional mechanisms were considered in this study that

were suggested in the 3GPP working groups.

Ericsson proposed short-term reservations (STR) [20], also referred to as short-term sensing, making use of two selection windows between T_1-T_2 and T_2-T_3 as shown in Fig. 2.8a. It works by sending a reservation signal in the first selection window, at T_{res1} to reserve resources in the second selection window on which to transmit the SCI and TB pair (T_n). In this mechanism, the resource on which to send the reservation signal is chosen at random, discounting the resources reserved for the SCI and TB pairs. Until the reservation signal is sent, the VUE continues to listen in case the reservation slot becomes reserved by another VUE, e.g.T_{res2}. At the time of sending the reservation signal, a future free resource is selected if it has not already been reserved in T_2-T_3. This reduces contention for the transmission of SCI and TB pairs.

The second approach discussed is a counter mechanism [21], this approach makes use of a simple counter to increase the randomness of the resource selection process as shown in Fig. 2.8b. When receiving an application layer packet, a counter is randomly selected between {1, 40}. In each subframe, the counter is decremented by the number of free subchannels. Once the counter reaches 0, a free subchannel is chosen at random in the next available subframe where the SCI and TB pair will be transmitted.

Outside of these mechanisms, there is one key SB-SPS parameter that is associated with aperiodic traffic patterns, namely *sl-reselectAfter*. This is a highly important parameter, as it controls how SB-SPS reacts to a missed transmission (MT), and by this specifically this occurs when a grant is configured and the time for transmission arrives but the application has not generated a new packet. In this case, as no packet is available for transmission, no transmission is made and an MT occurs. This is disruptive to the SB-SPS scheduling mechanism due to the fact that an SCI has already been sent reserving the resource, as such the resource will go unused, but neighbouring vehicles will have determined it to be reserved and have avoided selecting it. Further returning to the *sl-reselectAfter* parameter at this point the SB-SPS mechanism will refer to this parameter; it specifically controls how SB-SPS reacts to these missed transmissions. *Sl-reselectAfter* is a simple count of the number of MTs that can occur before the grant is relinquished i.e. before the RC reaches zero, ignoring *probResourceKeep*. In this study, two main configurations of this parameter are considered, the first being *sl-reselectAfter* equals one, meaning that any MT will result in the grant being relinquished

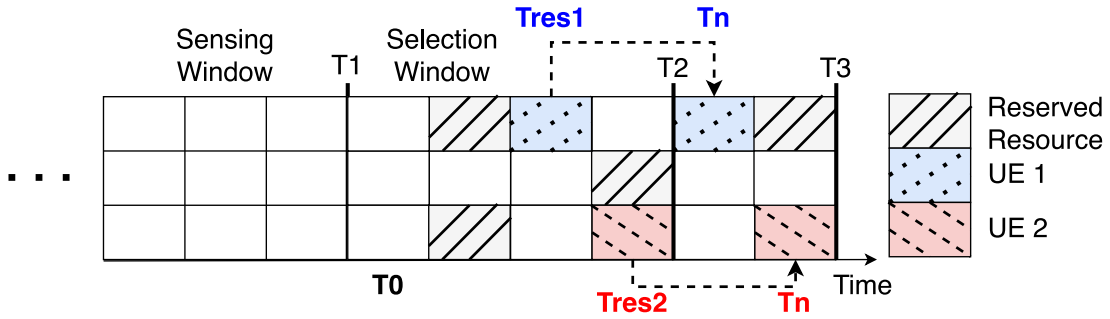

(a) Short-term reservation Mechanism.

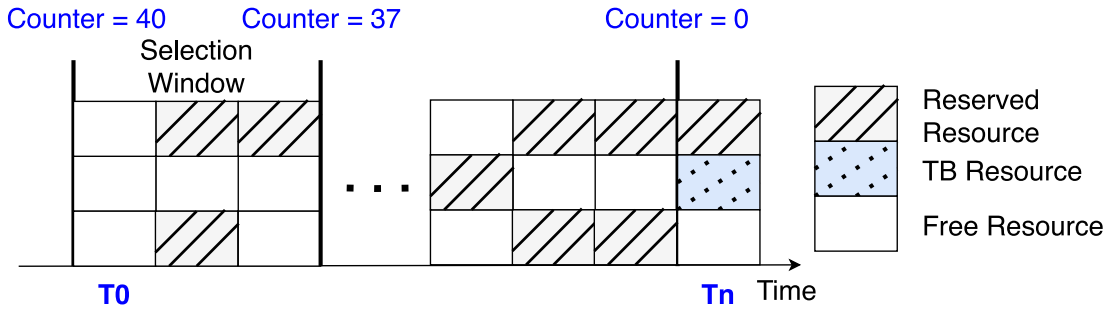

(b) Counter Based Mechanism.

Figure 2.8: Scheduling mechanisms for aperiodic traffic.

and rescheduling will occur on next packet reception. The second configuration is non-grant breaking where *sl-reselectAfter* is greater than RC so it is impossible for a grant break to occur before the grant would naturally end. This is an important parameter and its implications will be fully explored in subsequent sections; it is worth noting that most literature which discussed both C-V2X and NR-V2X overlook this parameter and focus on the non-grant breaking configuration.

2.3.6 Variable packet scheduling and MCS adaptation

The SB-SPS mechanism is designed to accommodate packet size variability with two specific approaches.

The first element focuses on grant generation, based on the configured subchannel structure, i.e., the number of subchannels and RBs per subchannel, as well as the configured allowed subchannel occupation, i.e., the minimum and maximum allowed subchannels, the SB-SPS mechanism will generate a grant based on packet requirements. This will entail the MAC layer determining, based on the current packet size, the number of available subchannels and the MCS options, what the optimal allocation is for this

packet. The MCS determines how efficiently data is encoded in RBs, with a higher-order MCS allowing an RB to hold a greater number of bytes. This comes with a trade-off; higher order MCS' exhibit less error redundancy in the encoded packet, and as such transmitted packets will be more susceptible to errors and ultimately loss. Lower-order MCS allow for more robust transmissions but have large RB occupancy. As such, there is a trade-off between efficiency of resource usage and effective communication. When referring specifically to optimal grant selection, the mechanism will select the fewest number of subchannels required to schedule the packet and, while maintaining the fewest subchannels, use the lowest MCS which allows this. It achieves this by referring to Tables 7.1.7.2.1-1 and 8.6.1-1 provided in [22], these tables map the allocation of MCS and RBs to the maximum bytes that can be encoded. This approach enables the SB-SPS mechanism to accommodate packets of different sizes effectively and efficiently.

The second aspect is the use of an adaptive MCS selection mechanism within the grant on a per-packet basis. As such, once the grant is generated, the MCS does not need to remain fixed and can adapt dynamically for each individual packet to best accommodate that packet based on its size, the granted subchannels, and the available MCS. Thus, SB-SPS determines the size of the grant, i.e. the minimum number of subchannels required, based on the initial packet size and the available MCS'. On initial and subsequent transmissions within this grant, it will then select the optimal MCS, i.e., using the fewest subchannels and the lowest MCS that enables the packet size to be scheduled, balancing resource usage with transmission robustness.

However, while this approach can be effective when packet sizes remain relatively consistent and a broad range of subchannel and MCS are available, it still can be the case that when packet sizes increase, the current allocation cannot encode the current packet. In this case SB-SPS is forced to relinquish its current grant and begin rescheduling based on this new increased packet size; this, while ultimately able to accommodate a wider range of packet sizes, does result in a grant ending before it was scheduled to, this is referred to as a grant break.

One key element when considering MCS adaptation is its impact on the physical layer. For instance, if the total transmission power is assumed to be fixed, any variation in the number of used RBs directly modifies the transmission bandwidth and, therefore, the Power Spectral Density (PSD) of

each transmission. This must be considered every time the MCS is adapted to accurately compute the power levels in the channel and ensure that the PSD remains within the 23 dBm / MHz limit as specified by the standard [23]. This is often overlooked when analysing the performance of C-V2X [24].

In the subsequent Section 2.3.7 MCS adaptation has also been suggested in the context of congestion control. In this case, a higher order MCS will be adopted to try and reduce the subchannel occupation of a packet. For the sake of clarity and to differentiate the two approaches, when dealing with variable packet sizes MCS adaptation will be used, while in the context of congestion control MCS congestion control will be used. In both cases the mechanisms are the same i.e. adjusting the MCS to reduce packet RB occupation but the goal is different, i.e. for MCS adaptation it is to ensure a packet fits in the same resources as the previous, while for MCS congestion control it is that a packet fits in as few subchannels as possible.

2.3.7 Congestion Control in C-V2X & NR-V2X

The first element to discuss when considering congestion control in C-V2X and NR-V2X is to understand what statistics are considered as part of the standard which measure both the level of congestion and the current resource usage of the vehicle. C-V2X and NR-V2X have two key statistics for congestion control; the first measures the level of congestion and is the Channel Busy Ratio (CBR). CBR is determined again using the sensing window and RSSI measurements, a specific RSSI threshold is configured, and for each subchannel in the last 100 ms that exhibits an RSSI above that threshold, that subchannel contributes to CBR. As such, CBR is purely the measure of busy subchannels (or high RSSI subchannels) over the total number of available subchannels. The second key measurement is the Channel Occupancy Ratio (CR); this measures the total channel usage of a vehicle. It is determined on the basis of both historical usage and the usage in the future based on the current configured grant. It comes out as a proportion of subchannels used and to be used over all available subchannels in a one second window.

C-V2X and NR-V2X will be required to handle a congested radio environment due to limited spectrum, widespread vehicular deployment, and frequent packet exchange. As such, congestion control techniques are very important for managing channel load and radio interference. DCC mechanisms as defined by ETSI have been well studied over the past decade for ITS-G5

[25, 26, 27] but have not been adequately investigated for NR-V2X or C-V2X. The operation of the C-V2X and NR-V2X MAC differs significantly from that of ITS-G5, particularly with respect to the scheduling of the MAC layer. SB-SPS assumes that packets arrive periodically and bases its resource reservation algorithm on this assumption. However, congestion control techniques such as Transmission Rate Control (TRC), will result in traffic arriving aperiodically, which will be shown later in this work to cause significant issues.

Initial standardisation techniques in 3GPP have focused on the use of table-based approaches with 4 mechanisms highlighted as possibly congestion control techniques [28]. It is important to note that the ultimate choice of technique to adopt was left up to those implementing the standard; as such, there is no specific standardised technique to adopt.

1. Elimination of retransmissions.
2. Adaptive Modulation and Coding Schemes to reduce channel occupancy of each packet (MCS congestion control).
3. Transmission Power Control to reduce communication range. (TPC)
4. Transmission Rate Control, i.e. packet dropping to reduce the number of transmissions in the channel. (TRC)

This thesis will study two of the aspects specifically MCS congestion control and TRC based approaches and propose solutions for congestion control in C-V2X and NR-V2X. The focus will be to overcome challenges similar to those faced with variable packet sizes and aperiodic packet arrival rates.

It is also important to recognise that while congestion control mechanisms may improve radio conditions and in turn the measured packet delivery rates, they may not improve the performance achieved at the application layer. If a large number of application layer or cooperative awareness packets are delayed or dropped, this may render the service unsafe or unusable. As such, it is of the utmost importance to consider the impact of DCC on the upper layer application Quality of Service (QoS) requirements. Additionally, fairness amongst vehicles and stability of channel conditions are a key concern with any congestion control technique. The focus of fairness is to ensure that no vehicle is starved of resources, making them an increased collision risk. This will result in neighbour vehicles not receiving adequate updates of their position through the CA service. Stability in turn focuses on ensuring that the

channel conditions are kept at a fixed level to ensure optimal performance and avoid oscillations which may result in increased collisions or wasted channel resources.

2.3.7.1 3GPP techniques

The 3GPP C-V2X standard also does not provide specific congestion control mechanisms, but defines the CBR and CR measurements under which it should be invoked. It also specifies a table shown in Table B.1, as part of a 3GPP working group. As already discussed in the early section, no specific mechanism is defined for use, but four are suggested i.e. elimination of retransmissions, adaptive MCS schemes, TPC and TRC. The basic principle of such approaches is that the vehicle maintains CBR and CR measurements throughout and compares the current recorded results to the table. The current CBR determines which CR-Limit should be applied at any given time. The vehicle will then, if its CR exceeds that limit, apply one of the aforementioned techniques. The goal being that the CR will then revert back to within the limit bounds or that the CBR is reduced so that the CR can stay at a higher level. It will be discussed in the subsequent literature review section, but these table-based approaches often exhibit CBR instability due to constant fluctuations between states in the table.

In recent years, ETSI has set out some details relating to congestion control for C-V2X [2]. Specifically, they describe how CBR and CR are to be measured. In Table B.2 ETSI specifies the maximum CR limit for each vehicle based on the measured CBR.

2.3.7.2 ETSI standardised techniques

ETSI defines a number of DCC mechanisms in [3], originally developed for ITS-G5. In this work, it was deemed worthy of investigation whether these approaches may be adopted by C-V2X/NR-V2X and as such leverage already standardised mechanism into the new technology.

The de facto mechanism is transmission rate control which works by increasing the delay between packet transmissions based on the CBR. The means by which this delay is determined is how ETSI distinguishes between its TRC mechanisms, namely *DCC Reactive* and *DCC Adaptive*.

DCC Reactive is the original approach specified by ETSI based on a state machine, where the state is associated with a particular CBR range. Depending on the CBR, a delay is introduced between consecutive packets to control the transmission rate. This is shown in Table B.3. The maximum allowed transmission rate for a CBR range is enforced using T_{off}, the time period before a new consecutive packet can be transmitted. *DCC Adaptive*, is a rate control mechanism based on the LIMERIC algorithm [25]. It uses an algorithmic approach to calculate the delay rather than a pre-defined lookup table [3]. It is designed to adjust packet rate transmission to converge to a target CBR, with a default of 68%. *DCC Adaptive* better considers factors such as stability and fairness to ensure that no node is starved. Amador et al. provide a detailed analysis of the ETSI *DCC Adaptive* in [29]. The algorithm is shown in alg. B.1 the primary control is the parameter δ, which is the limit of channel usage allowed for a vehicle at any given time. This δ is then used in the equation B.1 to determine the next t_{go} time that is the time at which the next packet is allowed to be transmitted. In this equation, $T_{on_{pp}}$ represents the transmission time of the current packet, representing the usage of resources for the next transmission; thus, the message is delayed by adding this calculation to t_{pg} which is the current time. The second equation B.2 represents an update function, where if a delay has occurred and the channel has become less congested, an early release of the packet can occur and this equation represents how the new release time is calculated.

2.4 Literature Review: C-V2X Mode 4 & NR-V2X Mode 2

This section will review relevant literature related to performance evaluations of the C-V2X and NR-V2X standards. Subsequent sections will then focus specifically on the literature related directly to the thesis contributions. Specifically aperiodic application patterns in Section 2.4.1, variable packet sizes and MCS adaptation in Section 2.4.2, congestion control in Section 2.4.3 and finally, improved CA performance through machine learning approaches in Section 2.4.4.

One of the first studies that was beneficial in developing an understanding of the performance of the C-V2X standard was by Molina et al. [28]. This work describes both the C-V2X standard scheduling mechanism, discusses some of

the congestion control mechanisms suggested in the standard, and provides an initial evaluation of the C-V2X standard using an internal model not made available to the academic community. The fact that the evaluation is completed in an unpublished model is one of the motivating factors for the choice of implementing the OpenCV2X model to enable other academics both to replicate work and investigate other areas of the standard. The same group later released an analytical model of the C-V2X Mode 4 standard, which we use to validate the OpenCV2X model [30]. They also completed a number of different studies of the C-V2X standard [31, 32]. These focused on the parameterisation of SB-SPS e.g. RSRP thresholds, probResourceKeep, and sensing window length and provided some guidance into their configuration. Furthermore, they investigate the performance of SB-SPS in comparison to a random scheduling mechanism and when retransmissions are introduced.

The most comprehensive of the early studies to investigate the performance of C-V2X Mode 4 was from Bazzi et al. [33]. This discussed in detail the primary parameters for SB-SPS across the MAC and PHY layers similar to the study previously outlined by Molina et al. This work is more comprehensive and analises several realistic environments, as well as providing a comparison with ITS-G5 and random scheduling. The main outcome of this work is a suggested configuration for the SB-SPS parameters such as the RC, probResourceKeep, RSRP thresholds, sensing window length, and the proportion of resources reported to the MAC layer after RSSI filtering. Providing an analysis of the tradeoffs in selecting any configuration.

As NR-V2X is a relatively recent standard, there are less studies of its performance. However, a key study that has been highly useful and provides a very comprehensive analysis of the Release 16 standard can be found in [18]. This study provides a highly comprehensive analysis of the changes made in NR-V2X Mode 2 from C-V2X Mode 4 with a particular focus on the scheduling perspective, as well as other aspects such as unicast and groupcast introductions, including changes to the SCI format to enable these communication paradigms.

Another useful study that provides a comprehensive review of the literature in the area is again from Bazzi et al. [34]. This study highlights the different modelling environments, the main challenges facing C-V2X and NR-V2X, the challenge of congestion control and aperiodic traffic patterns, and highlights some of the changes made from C-V2X to NR-V2X.

2.4.1 Aperiodic application patterns

The existing literature has not quantitatively evaluated whether Rel. 16 NR-V2X, including dynamic grant, can adequately support aperiodic traffic, with evaluations largely focusing on Rel. 14 C-V2X SB-SPS standard. Such a study is vital as ETSI CAM and Decentralised Environmental Notification Message (DENM) transmissions are fully aperiodic, and these will form the basis for future e-V2X services to support increased autonomy.

Most studies have stated that SB-SPS has been designed to better facilitate periodic traffic [35, 30]. Molina-Masegosa et al. [36] conducted a study contrasting C-V2X Mode 4 with 802.11p for periodic and aperiodic application models, as well as variable packet sizes. They further study adapting the resource reservation interval by balancing reselections against occurrences of wasted resources using three strategies. Notably, the authors suggest that the SB-SPS mechanism is fundamentally counterproductive for aperiodic traffic and highlight that further dedicated schemes are required. They do not consider grant breaking or dynamic grant in their evaluation.

Similarly, Bartoletti et al. [37] also investigated the performance of C-V2X Mode 4 with grant breaking for aperiodic traffic patterns, highlighting its poor performance. The authors conclude that while wasteful of resources, disabling grant breaking performs adequately when dealing with aperiodic traffic. The same authors continue to investigate SB-SPS for the NR-V2X Mode 2 standard in [19]. The findings were interesting in that they showed that the removal of RSSI filtering and RSRP averaging in Mode 2 detrimentally impacts SB-SPS performance for aperiodic traffic. Only the reintroduction of these features results in a similar performance to C-V2X Mode 4. As such, they conclude that using a static RRI within SB-SPS scheduling is not effective. Both studies do not investigate the NR-V2X dynamic grant mechanism, nor are dedicated aperiodic scheduling schemes considered.

Romeo et al. [38] consider aperiodic traffic in the form of DENMs, sent to alert vehicles of hazardous road conditions. The authors examine the impact of tuning SB-SPS parameters to support aperiodic packets, e.g., by reducing sensing windows, selection windows, and selection probability (*RSel*) when providing CSRs to the MAC layer. Periodic CAM transmissions with single DENM packets are considered as opposed to a traffic pattern/model with aperiodic arrival rates. In a later study [39], DENM retransmissions are considered, including the likelihood that two or more vehicles arbitrarily

choose the same resource, with the subsequent collision being maintained for the duration of the grant, thus affecting the successful receipt of the DENM. However, the authors consider a one-time reservation to mimic the Dynamic Grant mechanism but this operates using the C-V2X Mode 4 SB-SPS standard, i.e. includes RSSI filtering and RSRP averaging. As such, it doesn't accurately model NR-V2X Mode 2.

Finally, Lusvarghi et al. [40] investigate a hybrid application model using aperiodic and periodic traffic. The authors do not consider NR-V2X Mode 2 but rather examine C-V2X Mode 4 performance. Their findings show that periodic traffic is prioritised, resulting in poor performance for aperiodic traffic. This disparity is worsened when large aperiodic packet sizes are considered e.g. 1000 Bytes, due to high subchannel occupation.

Several elements have not been adequately explored in the aforementioned studies, and this thesis will seek to address all the challenges faced by C-V2X and NR-V2X in a comprehensive manner as highlighted in **contribution 2** of this thesis. As such, this study describes in detail both the grant breaking and non-grant breaking variants of C-V2X and NR-V2X, highlights the primary issues that both can cause and thoroughly evaluates different packet generation rates to provide a comprehensive evaluation of the scheduling mechanisms employed. Although some existing studies provide some analysis of these challenges, it is our belief that our analysis provides a fully comprehensive analysis of these challenges. In addition, this study provides an analysis of both the dynamic grant mechanism proposed in NR-V2X as well as two other proposals for aperiodic application traffic, and provides a solution to enable the adoption of dynamic grant with the current generation of ETSI standard applications. The challenges for aperiodic performance will be discussed in Chapter 4 as well as the proposed solutions, and a detailed analysis of both these challenges and proposed solutions is provided in Chapter 5.

2.4.2 MCS and variable packet sizes

Adapting the MCS in real time has been proposed mainly as a mechanism for distributed congestion control in C-V2X, as noted by 3GPP and the related literature. To a lesser extent, it has also been proposed as a potential solution for transmitting packets of variable size. There has been limited work focused on MCS adaptation in the literature; most studies in the area of MCS assume

that it is fixed and evaluate the impact of different MCS between simulation runs as opposed to adapting within the same scenario. In the same vein, there has been a somewhat limited focus on variable packet sizes in C-V2X and NR-V2X, with the majority of studies focusing on a single packet size.

However, studies in the area have investigated the impact of different static MCS configurations, which is not in compliance with the ETSI standards [41], or have focused on discussing the potential adaptation of MCS without analysing its operation and performance. For example, the authors in [28] and [42] briefly discuss the use of MCS adaptation for congestion control in C-V2X Mode 4 and analyse how increasing the MCS can potentially reduce congestion by reducing the number of required resources at the expense of a decrease in performance. However, these works do not analyse the implications of implementing this mechanism on the operation of C-V2X SB-SPS, which is crucial for its performance.

The study in [43] presents a detailed analysis of different congestion control techniques in C-V2X and shows the impact of different MCS configurations on the performance of C-V2X when there is vehicular congestion. However, the study assumes that the MCS configurations remain fixed during the simulations and does not analyse the operation of the MCS adaptation for congestion control following the ETSI specifications [41]. Such standards require that MCS adaptation operate dynamically during ongoing transmissions and as a function of the congestion conditions in the channel.

One of the first papers to look at both aperiodic arrival rates and packets of various sizes was by Molina et al. [36]. This study highlights some of the challenges with variable packet sizes and how it can introduce increased reselection events in the MAC scheduler as a result of the change in packet size. This change ultimately results in grants being unable to accommodate changes in packet size, and this can negatively impact performance. However, this study considers a fixed MCS throughout, which ultimately does not reflect how the standard is meant to operate (i.e., with MCS adapting to packet size). Qualcomm published a rebuttal of such approaches arguing that if adaptation is not enabled, the performance of C-V2X is artificially degraded [44], which this study will corroborate. A similar study for NR-V2X was completed in [45], though with the same limitations as in [36].

A recent study by Harri et al. [46] evaluated variable packet sizes in NR-V2X, with realistic packets ranging from 200 to 1500 Bytes. This study ultimately

concludes that the use of the highest-order MCS, specifically 64 QAM MCS 24, is the best choice to schedule these packets. Subsequent results will corroborate this, showing improved performance of higher-order MCS in Section 5.5. One aspect which is not considered is adapting within those higher MCS'. This thesis will ultimately show that while allowing for selection of higher MCS' beyond the current standard specification is better, it is the adaption of the MCS within the scheduler that will ensure the best performance in most cases.

Similarly, the authors in [47] provide a comparable study. This is interesting because it provides results from real-world experimentation involving two vehicles to provide validation against simulated results. The final conclusion is the same as [46], that is, a higher MCS and reduced subchannel usage will improve the performance of C-V2X/NR-V2X. Notably, the MCS remains fixed for each test and, as such, exhibits the same limitation as in the previous study.

The adaptation of MCS can have multiple applications in C-V2X beyond those already discussed in the related C-V2X literature. This is demonstrated by some works utilising the wireless 802.11p based vehicular standard. For example, the works in [48] and [49] show how the adaptation of MCS can be used to improve the performance of broadcast transmissions in variable channel conditions. Similarly, the authors in [50] highlight the relevance of MCS adaptation in the performance of vehicular unicast transmissions, which is a feature of NR-V2X under the 5G Release 16 specification [51].

As highlighted in **C5** this thesis will provide a complete analysis of the performance of SB-SPS with variable packet sizes by enabling MCS adaptation and highlighting its performance for congestion control (MCS congestion control). The primary element missing from the current state of the art in the literature is the fact that all analysis of SB-SPS either assumes a fixed MCS or when investigating varying MCS adopts a fixed MCS for each configuration, as such the MCS does not dynamically change throughout. This is not how the standard is designed where MCS is allowed to vary dynamically to enable the transmission of variable packet sizes. As such, to the authors' knowledge, this work provides the first full analysis of SB-SPS with variable packet sizes, where the MCS dynamically manages the various packet sizes. In addition, this work corroborates the work of Harri et al. [46] when discussing the example of the ETSI standard for C-V2X [52] the limit of MCS at 11 reduces the performance of SB-SPS by increasing the subchannel occupation, and

increasing this can greatly improve performance. The area where this study extends on the state of the art is highlighting the need to adapt within this higher range, specifically to avoid the use of an overly high MCS reducing transmission robustness unnecessarily or a lower MCS increasing subchannel occupation. Allowing MCS to dynamically adapt as SB-SPS is designed to can ensure the best performance for the given sidelink configuration. The final novel element provided in this study is an investigation of different SB-SPS grant generation techniques to overcome challenges faced when dealing with variable packet sizes, which thus far have not been studied.

2.4.3 Congestion Control in C-V2X & NR-V2X

Congestion control is well studied in the wireless or ITS-G5 standard with several different approaches standardised by ETSI [41], as well as years of academic literature on the topic ranging from table-based approaches initially [53] to algorithmic approach in LIMERIC [25]. There has also been an emphasis on the fairness and stability of such solutions [54, 55, 56]. On the cellular side, C-V2X and NR-V2X have had a much more limited investigation. A key aspect of this thesis focused on congestion control techniques with **contribution 4** proposing several variants and importantly focusing on the fairness and stability of those proposed approaches.

The initial academic literature mainly investigated congestion control mechanisms using table-based approaches. One of the first works to evaluate the impact of packet dropping for C-V2X is by Mansouri et al. [57]. It evaluates the performance of packet dropping based on the 3GPP C-V2X CR-limit tables. This paper highlights some of the issues when dealing with packet dropping, such as vehicles mistakenly choosing the same channel resources because SCIs are not being transmitted. Hu et al. [58] provide an interesting real-world study in C-V2X congestion control at a test site in China with numerous LTE-V2X OBUs deployed at an intersection to simulate a dense and congested environment. An adaptive MCS approach is adopted, whereby the MCS is increased when congestion occurs to fit the packets into 2 rather than 3 subchannels. This has been shown to be effective in maintaining good Packet Delivery Rate (PDR) and CBR performance, but the study is limited in that it focuses on a single packet size. This is unlikely in reality given that packet sizes frequently vary for vehicular services, e.g. taking per example the well known empirical Renault and Volkswagen data sets [59], as well as the

3GPP guidelines for packet size distribution [12]. In such cases, using MCS congestion control will only be impactful if the occupation of the subchannel can be reduced, which is strongly dependent on the packet size distribution [60]. Wendland et al. [61] propose an approach based on an SB-SPS reservation splitting technique. The authors split a single SB-SPS grant into multiple subgrants of lower frequency (e.g. 10Hz $\rightarrow$ 2 x 5Hz). When the network is congested, individual subgrants can be disabled without interrupting the SB-SPS grant mechanism. This approach is the closest to the proposed $RRI_{Adaptive}$ mechanism in that turning off a grant is similar to changing the RRI i.e., a single 5Hz grant is the same as a grant with an increased RRI of 200 ms. They compare their scheme with ETSI C-V2X DCC. Although the reservation splitting approach reduces recurring collisions within a single grant, it does not support dynamically reenabling grants as congestion changes. In contrast, the proposed $RRI_{Adaptive}$ approaches allow this to occur and works within the existing SB-SPS mechanism without requiring any changes.

Algorithmic congestion control techniques have also been explored. In the SAE standard, an approach combining mechanisms such as TRC, transmission power control (TPC) and MCS adaptation is proposed. The most prevalent of these is the SAE DCC mechanism [11] which uses power and rate control. The rate control algorithm is derived from the LIMERIC algorithm [25] and power control is based on The Stateful Utilization-based Power Adaptation (SUPRA) which is designed to control the communication range [62]. Research in [63, 64, 65, 66] investigates the performance of the SAE standard for congestion control. Generally, these authors have shown performance increases over standard rate control while showing minor improvements from power control, with the need for further study before determining their effectiveness for C-V2X. Sabeeh et al. [67] also built a solution based on TPC. This was shown to be most effective with their own customised version of the SB-SPS scheduler, which added the RC and colliding resource information to the SCI message. Unfortunately, the relevance of this study is diminished due to required modification of the SB-SPS C-V2X/NR-V2X standard. Choi et al. [66] use a deep reinforcement learning approach to determine which mechanism to adopt, i.e., TPC or TRC. This approach showed some marginal improvements over the simple table-based solution suggested by 3GPP, although may be somewhat over-engineered.

Another noteworthy study by Bazzi [35], comprehensively evaluates several

approaches, namely transmission power control, transmission rate control, and MCS congestion control for congestion management. It should be noted that the study also provides a comparison with ITS-G5 when applying each of the techniques. Importantly, when discussing the TRC approach, there is not a dynamic approach proposed; simply different configurations with different rates to show the direct comparison. As such how the rate is changed is not addressed nor the impact of GB or MT should the rate impact the scheduling mechanism. Ultimately, the paper highlights that TRC is the only effective approach for congestion control in C-V2X. TPC is not effective due to the synchronisation of the transmissions. When a collision occurs and interference takes place if all vehicles have the same reduced transmission power, the actual Signal-To-Noise Ratio (SINR) recorded remains the same due to both reduced RSRP and equally reduced interference. Relating to MCS congestion control, the findings showed similar impact as TPC in that it was not an effective approach for congestion control.

There are a number of key areas where this study improves the state of the art as highlighted in **Contribution 4**. The full definition of these will be provided in Chapter 4 and an evaluation of the challenges and proposals given in Chapter 6.

- Firstly, this study provides a comprehensive overview of the main challenges faced by SB-SPS when managing congestion specifically with TRC approaches, the current level of analysis is limited in the literature.
- The thesis introduces a new technique *RRI Adaptation* which ensures that any TRC approach can be effectively adopted in C-V2X and NR-V2X effectively. This includes three new *RRI DCC* mechanisms are proposed:
    1. $\text{RRI}_{\text{Reactive}}$ is proposed based on ETSI table-based *DCC Reactive* mechanism.
    2. $\text{RRI}_{\text{CR\_Limit}}$ is proposed based on 3GPP table-based CBR to CR-Limit approach.
    3. $\text{RRI}_{\text{Adaptive}}$ is proposed based on the ETSI algorithmic *DCC Adaptive* mechanism.

- A key area which has been overlooked in the literature is a focus on fairness and stability. This study is the first in the authors knowledge to focus on these when proposing congestion control approaches for C-V2X and NR-V2X. This is a key topic and has been extensively investigated in the context of ITS-G5/802.11p [54, 55, 56].
- Finally, this thesis provides the first analysis of MCS congestion control which was developed in collaboration with researchers in the Universidad Nacional Autónoma de México (UNAM).

2.4.4 Trajectory prediction and aperiodic performance

Trajectory prediction is a significant area of research, with numerous datasets, modelling approaches, and publications that highlight new approaches and models that might improve the prediction accuracy over any previously defined approach. This work ultimately is not a study exclusively on trajectory prediction and does not seek to propose new more optimal approaches. Instead, it seeks to adopt a mechanism that can with a good level of accuracy predict the short-term trajectory of a vehicle. This is described in Chapter 4. The underlying premise is that by adopting a short-term trajectory prediction mechanism it will be possible to determine the possible triggering times of CAMs. These are specifically based on changes in the vehicles dynamics i.e. speed, heading, position changes all bounded by a one-second maximum time. This approach may be integrated to ensure that SB-SPS can more accurately maintain its sensing history and improve overall scheduling performance. In this vein, the work focuses on Recurrent Neural Network (RNN), these being Long-Short Term Memory (LSTM) or Gated Recurrent Unit (GRU). Such models combine some level of historical information from past data provided and do not focus exclusively on the most recent data point. They have been used extensively as models for trajectory prediction in the past. These RNN-based approaches have been shown to be useful in predicting vehicle trajectories with high levels of accuracy [68, 69, 70].

There have been two papers that have investigated this issue and proposed solutions based on machine learning approaches. Lusvarghi et al. [71] present the first work in this area with a k-Nearest Neighbours model that incorporates a change to the standardised SB-SPS process. The authors propose that the

SB-SPS reselection counter (C_{resel}) be included in the SCI. This additional parameter specifies the remaining transmissions in the grant, i.e. indicating how long a vehicle will maintain its grant. In the grant selection process, C_{resel} is calculated by predicting how long a sequence of the same RRI will be maintained. This is potentially problematic as variable vehicle dynamics within a single grant may result in a highly variable sequence of RRIs leading to short-lived grants. In contrast, the proposed model can handle such cases more easily. The author's solution also relies on accurate knowledge of the speed and position of the preceding vehicle. While comprehensively evaluated, it is not clear whether the proposed model would generalise for unseen scenarios, only a highway scenario that is outside their original dataset is explored. Finally, this model would require retraining for any CAM rule change or for rule changes associated with different services, whereas the proposed model predictions would be applicable to any rule set, due to itss reliance only on mobility prediction and not the specific CAM time trigger.

The second solution was proposed by Seon et al. [72]. They propose a Multilayer Perceptron model, which is a form of neural network commonly used in image classification problems, but also applicable for time series data. The model directly predicts the CAM time by incorporating speed, position, heading, and time changes, as well as the previous CAM trigger condition. The model is trained on two real-world vehicle mobility data sets representing an urban and suburban driving scenario. 3 vehicle traces are in the urban scenario and the remaining 7 traces are in the suburban; this represents a relatively small dataset. The model shows high levels of accuracy across all train, validation, and test datasets. A primary issue with this study is that the results on the improvement of SB-SPS with their model are reported on the test set. The test set is still formed from the original traces and, as such, is likely to be similar to training data; therefore, this cannot show the effectiveness of the model in generalising to unseen data. As such, the model may not prove to be effective outside of these or similar scenarios. Additionally, while aspects of SB-SPS performance improvements, such as reduced waste of resources when employing the model, are described, the paper still does not adequately explore the full implications and possible benefits of their model.

Contribution 3 is the proposed solution in this study and adopts a trajectory-based approach that compares the future movement of the vehicle with the CAM trigger rules. The solution allows for the prediction of the CAM IPT and is described in Chapter 4 Section 4.3.3 and an evaluation is provided

in Chapter 5. The key benefit of the proposed solution is in its generalisation to unseen scenarios showing it can accurately model mobility and improve on scheduling performance.

2.5 Simulation approaches and platforms

There are numerous approaches to adopt when investigating any communication technology, from empirical testing in the real world to lower-layer signal modelling or analytical modelling. A common approach to adopt is simulation, in this a technology is implemented to mimic its real world performance. This can have multiple different levels of concern, from full-stack whole-network simulation giving an overview of full system performance, to purely link level simulation concentrating on PHY and MAC layer performance. These approaches have different requirements and aspects of interest to the person performing the simulation and there are various platforms that can be adopted for simulation. These models are generally specialised in certain types of simulation as mentioned previously, e.g. full-stack network simulation.

As discussed in the previous chapter, large-scale real-world testing is cost prohibitive and until recently commercial C-V2X On-Board Unit and RSU equipment was not available for testing. As such, this study makes extensive use of simulation modelling.

One of the contributions of this thesis is a discrete event simulation model called OpenCV2X [73]. This was the first full-stack open source model of the 3GPP C-V2X Rel. 14 model to be made available to the vehicular research community and is described in further detail in Chapter 3. Discrete event simulation models the state of a system on an event-to-event basis, including how the system changes or reacts to these events. In the field of networking these events typically represent a packet being generated and may model various communication phenomena that impact reception e.g. signal degradation due to pathloss or interference etc.

Discrete event simulators is especially valuabe for studying and analysing vehicular communication system for several reasons as it can allow for the creation of realistic models of the complex interactions between vehicles, infrastructure, and various communication protocols, thereby capturing the discrete events, interactions, and behaviours of vehicles and communication devices accurately. This realism is crucial for understanding the performance

of communication systems in dynamic vehicular environments.

Secondly vehicular communication systems are inherently event-driven, meaning that actions and events occur at specific points in time rather than continuously. Discrete event simulation is well-suited for modelling and analysing systems with discrete events, making it an ideal choice for capturing the sporadic nature of vehicular communications and their complex time related interactions.

The alternative approach for simulation modelling is analytical or mathematical modelling; this will be discussed in the context of Matlab, which is a common platform used for such models. Analytical or mathematical models generally concentrate more on the link layer or communication layers and seek to model them with the greatest level of accuracy possible, mimicking the real world performance as closely as possible. This real-world performance is then mathematically modelled through specific equations, which can model the behaviour of how radio technology propagates its transmitted data. These models are less useful for full-stack investigations due to them generally being more computationally complex and, as such, more costly to run, but high accuracy can also require that the models be suited to very specific environments and not generalisable to other scenarios. An example in this case being a model designed to model C-V2X performance in a highway scenario is not likely to be translatable to an urban scenario, this is more easily done in a discrete-event simulator through changing the channel model employed.

Ultimately both approaches have their place and it is purely a decision based on what the user is most concerned with, if they want the highest possible accuracy on link layer performance an analytical model is best, as generally these will be based on some form of empirical result and directly implements real-world performance. This comes at the cost of generalisation and compute complexity. While if the user is concerned more with overall performance both network and application level and wants to test the technology across various scenarios, then the best approach to adopt is a discrete event simulation environment.

OMNeT++ as will be described in the subsequent Chapter 3 is ultimately the modelling environment chosen for this project. It is a C++ based discrete event simulator and at the time of defining the model required, it was the most commonly used for V2X applications. This was as a result of the popularity of such ITS-G5 models as Veins [74] and Artery [75], these models were

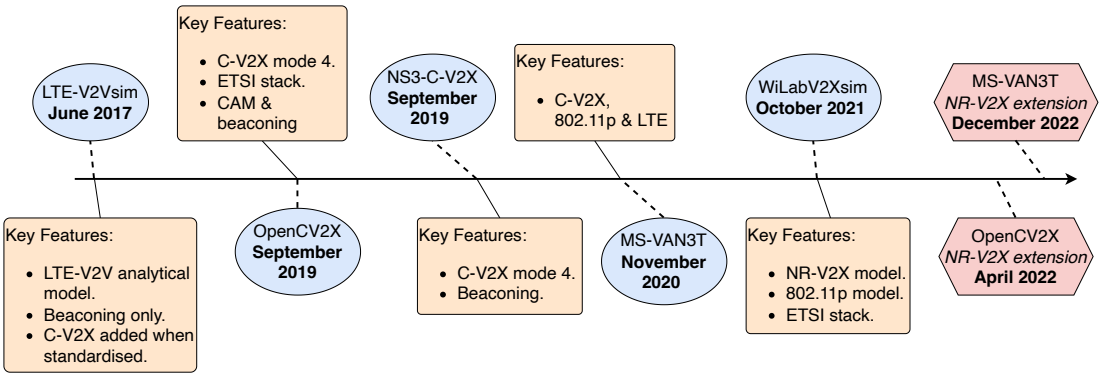

Figure 2.9: Timeline for each model

extensively used to investigate that technology. This and the modular nature of OMNeT++ encouraged the adoption of the platform for implementation of the OpenCV2X model used in this project, as it enabled the OpenCV2X model to integrate with Artery and access ETSI standards for applications.

The subsequent sections will discuss some of the most pertinent models that compare to the implemented OpenCV2X model in OMNeT++ that is presented in Chapter 3. To provide an overview for comparison, the following diagram outlining the key dates and features available in each model are presented in Fig. 2.9.

2.5.1 Analytical models (LTE-V2V-Sim & WiLabV2Xsim)

Matlab is a proprietary environment that allows for many different forms of complex simulation but is most commonly used for analytical or empirical modelling. Through modules such as Simulink, it is possible to model complex wireless networks and numerous other services such as ADAS mechanisms.

In the context of C-V2X there are two models which are available, one of which is used as a benchmark in this thesis for validation purposes. The first model, LTE-V2V-Sim [76], was based on the C-V2X standard and was later extended and renamed to WiLabV2Xsim [77] to include NR-V2X features added in Release 16. This model has been widely adopted and integrates useful features such as enabling ITS-G5 and C-V2X/NR-V2X simulation in the same environment. Several of the academic studies previously highlighted adopted LTE-V2V-Sim or WiLabV2XSim [35, 78].

2.5.2 Discrete event simulation models (NS3-C-V2X & MS-VAN3T)

NS3 is a C++ based discrete event simulator that has integrated several libraries for different networking technologies. At the time of starting this research, it had developed a WAVE model only. Eckermann et al. subsequently released an NS3-based model built on a D2D model available in the environment for C-V2X Mode 4 [79]. This model is the most directly comparable to the implemented OpenCV2X i.e. a C++ based environment based on discrete event simulation. It is worth observing that it lacks some of the features that are available in OpenCV2X, which specifically integrated ETSI standard applications such as CA and DEN. This model was later integrated into a more complete V2X simulation environment called MS-VAN3T which originated in 2020 [80]. This version included ETSI application models such as CA and DEN as well as simulations based on 802.11p and V2X applications based on interactions through the eNodeB within the NS3 LENA module. More recently the model was extended with the NR-V2X feature set and includes emulation features [81].

This model, while released much later, is directly comparable with OpenCV2X and thus it useful to compare the two primary vehicular communication technologies, namely 802.11p and C-V2X/NR-V2X. 802.11p based simulations is also possible with the OpenCV2X model as it is fully integrated with Veins or Artery based simulation models. As such, the feature set between the two models is not significantly different allowing simulation of the same scenarios and application models across both technologies. The primary advantage of MS-VAN3T is in the integrated ability to use the LENA simulation model for LTE based simulation. This is not currently possible with the latest version of OpenCV2X which would require integration with Simu5G. However OpenCV2X is still an important contribution as it was the first to provide the most complete feature set with ETSI based upper layers being available on release in 2019. This provided significant opportunities to investigate the effectiveness of the C-V2X model when using the current generation of ETSI standard applications and congestion control techniques.

2.6 Conclusion

To conclude this chapter it is important to briefly highlight the areas already discussed, those being:

- An introduction into vehicular networking, with a focus on both regulatory spectrum allocations and "Day 1" services.
- A comprehensive analysis of the C-V2X and NR-V2X technologies with a particular focus on the employed scheduling mechanism SB-SPS. As well as how that mechanism changed from C-V2X to NR-V2X and the introduced Dynamic Grant mechanism. Finally, a focus on MCS adaptation and variable packet sizes as well as ETSI and 3GPP congestion control standards.
- The chapter finally provided a literature review focused on a number of key areas:
  - General discussion of the origin of V2X communications.
  - A detailed analysis of C-V2X and NR-V2X with a focus on aperioidc, variable packet size, and congestion control, and how these impact scheduling.
- The final element of the chapter focused on different simulation platforms one of which OMNeT++ will be discussed in greater detail in the next chapter.

This background will be fundamental to the understanding of the subsequent chapters, which will focus particularly on the scheduling mechanism in C-V2X and NR-V2X. The first element of this which will be described subsequently is translating that SB-SPS scheduling mechanism into a simulation environment specifically OpenCV2X, the first contribution of this work.

Chapter 3

OpenCV2X Simulation model - Design, Implementation & Validation

3.1 Introduction

The previous chapter described the operation of the C-V2X and NR-V2X standards and, most importantly, the SB-SPS scheduling mechanism employed by each. This chapter will discuss **contribution 1** (**C1**) and how this is implemented; OpenCV2X[1], the first full-stack open source C-V2X model made available to the V2X research community. This model is in widespread use and directly resulted in a collaboration with researchers at UNAM.

OpenCV2X was first released in 2019 and was the first open source, 3GPP standard C-V2X full-stack simulator. That was fully compliant with 3GPP standards fully integrated into the widely adopted OMNeT++ simulation platform. OpenCV2X was developed on top of the *SimuLTE* model which is used extensively in OMNeT++ and required significant implementation throughout the entire stack. Further, it provided the academic community access to a wide variety of modelling environments and standards, enabling an *Artery* integration which provides ETSI compliant message formats and applications, as well as a *Veins* based solution which enables simpler modelling. This flexibility provides users the ability to concentrate on the aspects most of concern to their research, and we believe enables users to

[1] www.cs.ucc.ie/cv2x/

study a wider variety of issues with the model. It should be noted that each integration is not easily managed especially when dealing with a complex model such as *Artery* which is made up of several sub-components and implements it's own build environment. As such, providing this with little requirement for the end-user to understand the intricacies of the integration and enable them to confidently use a known, validated and compliant model which is fully open source and as such open to scrutiny is a contribution in and of itself.

Specifically this chapter will include:

- An analysis of the model requirements in Section 3.2 and highlights the contribution that the model has made to the vehicular community.
- An overview of the architecture of the model along with implementation details in Sections 3.3 and 3.4.
- Section 3.5 will provide a validation of the OpenCV2X model benchmarked to an analytical model.

3.2 Model requirements and impact

The overarching goal was to implement a simulator that provided a full stack open source implementation of the initial 3GPP C-V2X release 14 standard [82] and integrate it with ETSI message sets and generation rules. This dictated the use of a discrete event simulation platform to best model the implementation of a full-stack simulation environment and the event-driven traffic between network nodes.

At the time (2019) there were no open source implementations of the C-V2X standard. So it was the first fully open-source full-stack implementation of the C-V2X standard available to the community. The implemented model enables the vehicular communications research community to evaluate the performance of their hypotheses, using an end-to-end full-stack approach, thus stimulating further research on this topic. It further allows for full reproducibility of any discussed results, which is key to any scientific endeavour and is lacking in numerous studies in the literature with proprietary or custom models which are not made readily available to the academic community. In addition, the open source approach provides a high quality model as a result of community bug reporting and open implementation.

Overview of application requirements:

- The model should be fully compliant with ETSI message sets and generation rules such as [15, 16].
- The model should be designed to operate purely through the Veins model [74]. This ensures flexibility to the community; if they do not require to focus on the ETSI standards, they are able to use the default application interface of OMNeT++. Veins is an important simulation environment as one of the earliest and most widely adopted 802.11p models available to the vehicular community.

As a result of the two above requirements, there are 2 versions of the model available to the wider vehicular networking community. The first version which follows the ETSI standards is integrated with Artery [75]. Artery is a simulation model that provides ETSI compliant applications specifically CA/DEN as well as other services such as different congestion control mechanisms for ITS-G5. Additionally, there is a second version of the model available to the community that uses Veins for **S**imulation of **U**rban **MO**bility (SUMO)integration and does not require the complexity of using the ETSI applications. This simplified version ensures easier development with a focus more on the access layers and enabled a wider adoption of the model by the vehicular networking community.

Overview of access layer requirements:

- The primary requirement was to develop a model that is compliant with MAC and PHY layer standards of C-V2X Rel 14.
- This implementation would be validated against empirical or analytical models. This was achieved using [30].
- It was the goal of the project to maintain the standard in line with current releases; as such, the current version of the model enables NR-V2X features with a focus on scheduling performance.

The initial publication [73] (*Pub 4*), as well as the validation publication (*Pub 7*) on arxiv, have been referenced 37 times across IEEE, ACM, MDPI, science direct and Arxiv publications, and the associated github page receives 50 unique visitors weekly. Organisations utilising the model include the Mitsubishi Electric R&D centre Europe (MERCE), IRT System X, University of Cantebury Christchurch, University of Clermont Auvergne, among others.

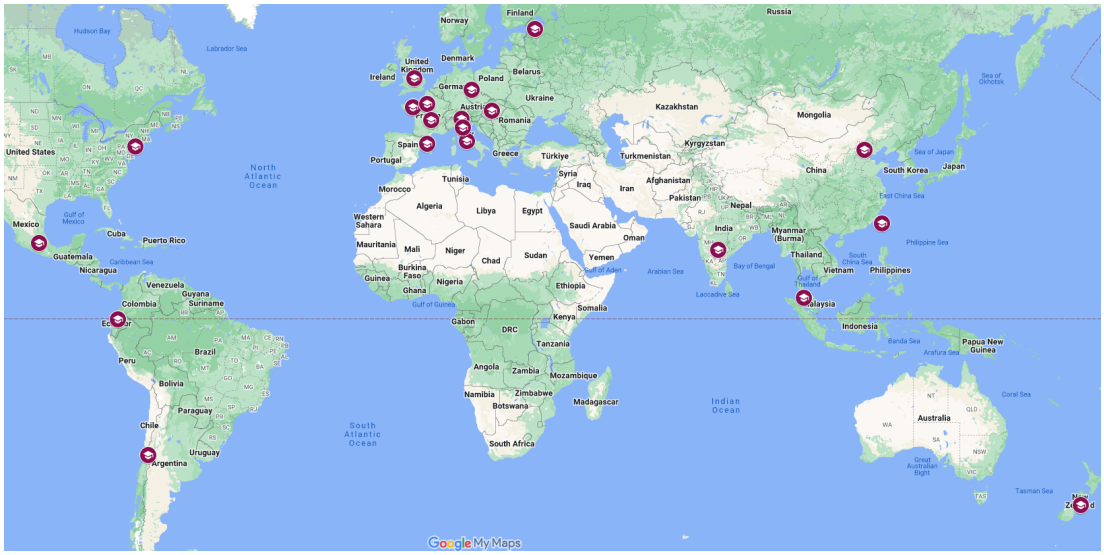

Figure 3.1: Global geographic distribution of OpenCV2X users

Ultimately, the model also led directly to a collaboration with Andres Burbano-Abril and Victor Rangel Licea at the National Autonomous University of Mexico and the Metropolitan Autonomous University-Iztapalapa. This collaboration directly led to the MCS congestion control work with Andres Burbano-Abril, Victor Rangel Licea and Miguel Lopez-Guerrero [60] (*Pub 3*). It further led to engagement with Christoph Sommer of Dresden University of Technology to integrate with veins and finally Giovanni Nardini of the University of Pisa to integrate with Simu5G. Other academics reached out looking for support, which we provided extensively, enabling users to develop their research more efficiently. A personally pleasing aspect of this was an acknowledgement in a masters student's thesis which was based on an investigation of the performance of C-V2X Mode 4 model and for which he used OpenCV2X; As part of this I helped him investigate any issues in their simulations setups as well as discussed analysis of results. Similar help has been provided to a dozen other users who reached out to me seeking advice.

The open source nature of OpenCV2X is a significant factor in its adoption, and it has enabled the improvement of the model over time, providing a ready avenue for users to provide bug reports and enable the model to be improved over time.

3.3 OpenCV2X Architecture

This section describes the architecture and core components of the OpenCV2X model as illustrated in Fig. 3.2. The primary elements are SUMO and OMNeT++ which manage the microscopic vehicle mobility and road topology simulation and the wireless network simulation, respectively. OMNeT++ contains a diverse set of models that provide different network technologies and interfaces that enable complex V2X simulations as well as the latest generation of V2X applications.

The main new features are shown in Fig. 3.2 in red, orange represent components requiring significant adaptation, and then the grey components also had some more minor adaptation. These changes throughout the entire stack ensured that the OpenCV2X standard was compliant with both the 3GPP standards for C-V2X but could consume ETSI standard applications from Artery, as well as operating with Veins if required.

We next describe each component placing particular emphasis on the novel components added.

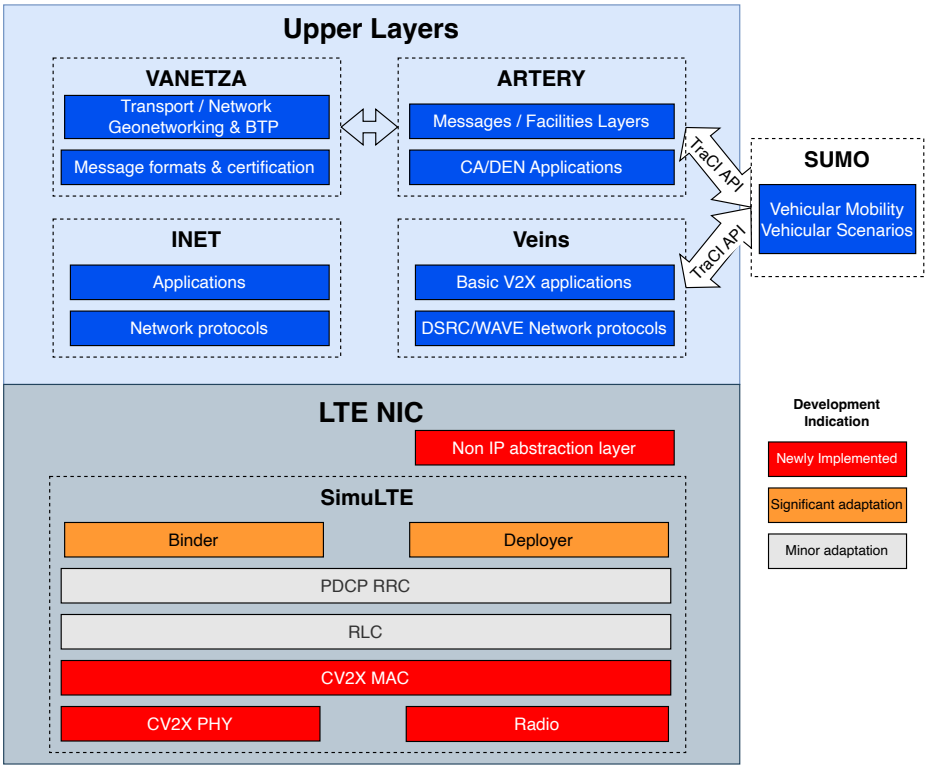

Figure 3.2: OpenCV2X model architecture with Artery.

3.3.1 Microscopic mobility modelling

A key part of any V2X modelling is how the mobility of the vehicles is modelled for this OpenCV2X employs SUMO which is an open source road traffic simulation platform. It is designed to enable the modelling of microscopic vehicle, pedestrian, and cyclist mobility, although our concern is primarily vehicular mobility [83]. It was developed by the German Aerospace Centre and has been extensively studied; its underlying techniques for modelling such as the Krauss car following model, have been shown to be effective in mimicking real driver performance [84]. It has become a de facto standard for vehicular mobility simulation in the academic community of V2X. All major simulation platforms provide an integration with SUMO through its **Tra**ffic **C**ontrol **I**nterface (TRACI) API, and both Artery and Veins described subsequently provide interfaces to SUMO through this API. NS3 simulation models such as ms-van3t [80], which provides similar services to Artery and OpenCV2X, use SUMO as well as Matlab, which also provides integration through the same API. Through this integration, due to its open source nature and relative ease of use in the development of complex road networks, the wider community has adopted it extensively.

The TRACI API is again the primary means by which integration can be made between the road network simulator and any other simulation platform, whether it be event-based or mathematical-based, such as Matlab. Through this API, it is possible to update the position of vehicles in step with the simulator time and mimic the realistic dynamics of the vehicles within the simulation environment itself. The interface further enables the modelling of complex applications, for example, in modelling interactions with emergency vehicles. It is possible through the TRACI API to enable vehicles to react to received messages, in this case a DENM indicating the approach of an emergency vehicle and, as such, to take action such as slowing down and changing lanes to get out of the vehicle path. This interaction provides the simulation environment with both realistic vehicular mobility as well as the ability to model complex reactions to different V2X applications and have that be reflected within the SUMO road traffic simulation. This is a powerful tool, as it not only enables proper simulation of wireless network interactions, but also provides a mechanism by which the performance of the road network can be studied when V2X applications are integrated, which can impact it.

Ultimately, SUMO forms one of the core components of the OpenCV2X model,

providing a comprehensive road traffic simulation environment that can be relied upon for our network simulations. We further made extensive use of its numerous realistic mobility models of different scenarios, such as cities and highways, in developing different aspects of our work, such as the CAM IPT prediction model described in Chapter 4.

3.3.2 Wireless network simulation

The second key element for the project is a network simulation environment as described in Chapter 2 Section 2.5, OMNeT++ is the chosen simulation platform; The choice was driven by the following factors:

1. To Facilitate development of a full-stack simulation model, using an event-based simulation platform.
2. OMNeT++ is the most widely used simulation platform by the V2X community, although in the last 18 months NS3 vehicular models have advanced significantly.
3. OMNeT++ had many modules e.g *Veins*, *SimuLTE* and *Artery* which could be leveraged and extended in development.

OMNeT++ is based on the C++ programming language for implementation, though it uses its own internal Network Description (NED) language for the definition of the different classes and modules that make up a specific device. In our case, taking the *LteNIC* as an example, this module can be made up of several submodules, such as the *LtePhyLayer* which will be defined by its own NED file. These NED files act as the interface between the person defining simulations through parametrisation, as well as results recording, and the direct C++ code implementing the underlying technology, as such the lowest NED files would correspond to the C++ code files. This approach enables the simulator to be highly modular in nature, where connecting different submodules simply requires the use of a NED file describing the relationship and providing the connections between the modules. The connection between modules is through gates, which act as the interface through which messages are shared between the different modules or layers in the case of the *LteNIC* there will be gates connecting the *LtePhyLayer* and the *LteMacLayer* providing a specific approach by which they can communicate information to each other without direct integration in the C++ code files. This approach is powerful as it enables the combining of different models together with minimal additional

C++ coding.

3.3.3 Cellular network simulation

OpenCV2X makes use of and extends *SimuLTE* as the basis of cellular network model. *SimuLTE* is designed to implement the cellular standard in OMNeT++ [85], and was initially designed purely for eNodeB enabled communication between the eNodeB and any cellular devices. The *SimuLTE* model provided a full cellular stack from radio or channel modelling, through to scheduling in the eNodeB, physical and medium access layers, PDCP and RLC layers. All the required stages in cellular communication were developed to be 3GPP standard compliant. With the introduction of D2D communication in 3GPP release 12 [86], this was then incorporated as part of the *SimuLTE* model. D2D communication was a form of cellular scheduling in which a cellular device could, through or with an intermediary, transmit and receive without being directly connected to the eNodeB. At the beginning of the development of OpenCV2X *SimuLTE* had developed V2X scenarios that adopted this D2D style communication to allow vehicles to communicate with each other scheduled by an eNodeB. As such it extended into the V2X area but in a limited fashion, but this provided a basis on which OpenCV2X was built.

SimuLTE as will be described in Section 3.4 is extensively changed to enable the 3GPP standard C-V2X model. Including, but not limited to, the following:

- Decoupling of many key modules from the eNodeB to allow autonomous resource selection.
- Introduction of Non-IP based network traffic necessitating changes throughout the stack.
- Implementation of new MAC and PHY layers to implement the SB-SPS standard, including grant generation and sensing capabilities.
- Finally, extensive PHY layer changes including new channel models, channel measurements, and techniques for MCS adaptation.

3.3.4 Mobility modelling integration

Veins is one of the most widely adopted V2X models [74], its use has been prevalent for more than a decade and it is a well-developed and understood simulation environment. It provides integration with SUMO through the

TRACI API, but it was also designed specifically to implement the IEEE WAVE standard, based on 802.11p, the wireless vehicular technology which was discussed in Chapter 2. These aspects of the model are of less interest to the OpenCV2X model which is focused exclusively on the C-V2X and NR-V2X standards, but having an integration with Veins does provide the wider V2X community with certain advantages. The primary advantage is in providing a single interface from which to compare both technologies; this can easily be done through the modular nature of OMNeT++ so that the same applications can be tested on both technologies without the need for reconfiguration. Additionally, it also provides good scope for investigating multi-RAN approaches by implementing scenarios where both technologies operate on the vehicles. This motivated the integration of our model with this model, as it provides greater scope for investigation to the wider V2X community.

The primary means of integration is the use of the TRACI API which is implemented as part of Veins. This can be easily integrated into any OMNeT++ model through the *Veins Mobility* interface (which is defined as a NED module). This provides a specific interface to TRACI and allows easy integration of OMNeT++ and SUMO without the need for any significant implementation. By providing the model integrated with Veins, this functionality is inherited without any need for user implementation. The only additional step involved is to manually turn on the TRACI server which veins uses in simulation; this is provided by a Python script.

3.3.5 ETSI compliant application layer

Artery is an ETSI ITS-G5 compliant simulation model for OMNeT++ [75]. That provides both ETSI compliant lower layers i.e. MAC and PHY and the full stack including ETSI applications such as CA and DEN. Artery provides these services through Vanetza, which is described in Section 3.3.5.1. Artery is similar to Veins in that it provides its own TRACI implementation and, as such, integrates with SUMO, allowing it to operate without Veins. Integrating with Artery allows the cellular stack to run ETSI compliant applications and message formats. As Artery provides DCC functions, it also enables the investigation of those techniques with C-V2X/NR-V2X which is another topic of interest in this work.

Integration with Artery is a much more complex matter as described in detail in Section 3.4. It provides its own build environment through CMake and thus

makes the typical OMNeT++ IDE based on Eclipse less useful for experimentation, due to it being primarily make-based. This inflexibility is another motivating reason for the use of both veins and Artery based OpenCV2X models. Artery is useful when investigating specific ETSI compliant applications and mechanisms. However, Veins offers much greater flexibility and ease of use if the concern is purely network performance with less focus on standard compliance. Providing both interfaces enables users of the model to selectively choose the environment that best suits their study and should widen the adoption.

As highlighted, the integration of OpenCV2X with Artery is more complex and required some new implementation to operate effectively. This is primarily in the introduction of a new radio driver that enables Artery to communicate with *SimuLTE* this is described in detail in Section 3.4.2.

3.3.5.1 ETSI application definitions (Vanetza)

Vanetza is a standalone but core component of Artery. It operates to implement many of the upper-layer aspects of the ETSI standards, from geonetworking, DCC, and all the security aspects around certification.

For OpenCV2X the primary interactions with Vanetza are to provide ETSI standard message formats based on the ASN.1 syntax. As such, any experiments simulated through Artery have ETSI compliant CAM messages being sent.

3.3.6 Network component modelling (INET)

INET is the largest repository of different networking models in OMNeT++. It provides many of the underlying network technologies that may be required in implementing a full network. For example, in an experiment with a Road-side Unit (RSU) veins may be used for 802.11p-based communication from the RSU to the vehicle. Then, if modelling the RSU to core network performance INET would provide the models for ethernet, IPV6, TCP and router models. Outside of this, it further provides models for wireless propagation and channel modelling and different WiFi variants.

In essence, INET is the first repository to search for any network technology that may be required for experimentation. INET does not require any integration from OpenCV2X, but the primary advantage of adopting

OMNeT++ is that it can be easily added to any experiments. As such users of the OpenCV2X model can leverage INET in their experiments to envisage more complex systems outside of simply the direct V2X communication. This allows for the investigation of more complex experiments and move towards use-cases such as edge experimentation without significant implementation or integration overhead.

3.4 Model Implementation

In the subsequent sections, the OpenCV2X architecture implementation is described. This can be categorised into the impact of autonomous resource selection, integrations with both *Artery* and *Veins*, the MAC and PHY layers, and finally channel modelling.

Fig. 3.2 in Section 3.3 illustrates at a high level the OpenCV2X stack. It can be seen that OpenCV2X models the full cellular stack from application layer right through the radio for the VUE in Mode 4. In the proposed OpenCV2X architecture, the upper layers can leverage either the *Artery* or *Veins* frameworks. Both provide the TRACI API to integrate SUMO road traffic simulation with OMNeT++ network simulation. Both can also leverage INET for different network protocols that may be of use in a simulation campaign. This offers OpenCV2X users greater flexibility in which vehicular applications they wish to model on top of the C-V2X radio. They can adopt the *Artery* framework if they wish to focus on ETSI standard applications and protocols in their investigation. Otherwise, the more flexible *Veins* framework may be used if these specific ETSI standards are not required.

Fig. 3.3 describes the interaction between the layers and the messages exchanged. In this case, the model is configured with Artery as the upper layer.

The most significant areas of implementation were the introduction of autonomous resource selection which required extensive reworking of *SimuLTE* which is described in Section 3.4.1. The next area was the required integration with *Veins* and particularly *Artery* which also required further extending of both *SimuLTE* and the *Artery* frameworks, both are described in Sections 3.4.2 and 3.4.3. The final elements of implementation took place again in *SimuLTE* and focused on modelling SB-SPS scheduling through the MAC and PHY layers and finally accurately modelling the channel, all are described in Section 3.4.4 to 3.4.6.

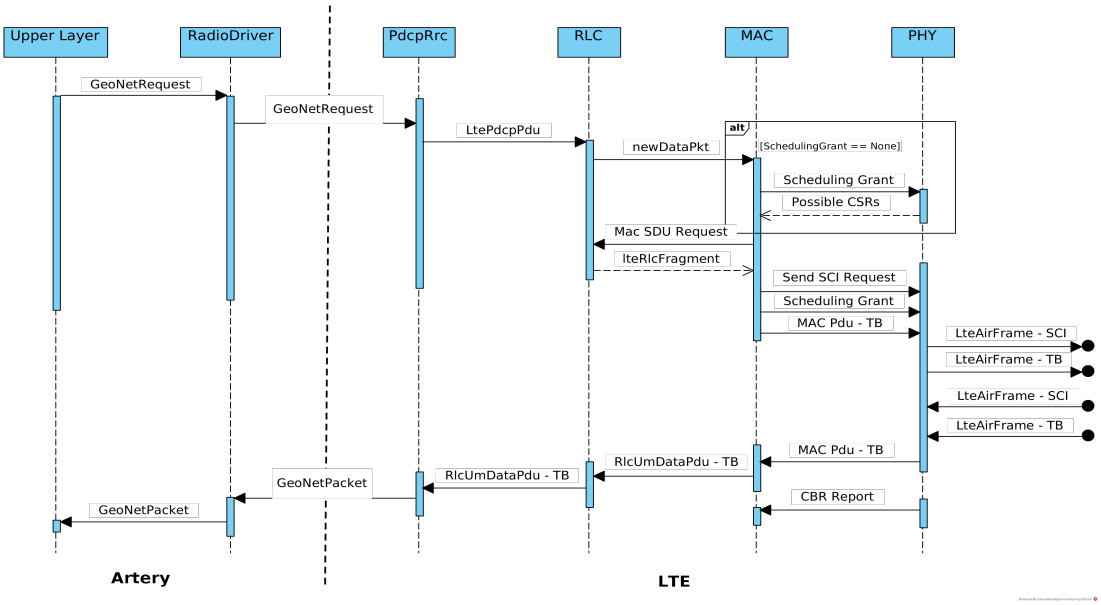

Figure 3.3: Sequence Diagram of packet exchange in OpenCV2X

3.4.1 Autonomous resource selection

This section will introduce the elements of *SimuLTE* which required significant reworking to enable autonomous resource selection including:

- The removal of the eNodeB from the simulation environment, necessitating changes to both the *Binder* and *Deployer* modules which control RB allocation and UE position respectively.
- The introduction of a mehanism by which a grant can be generated by the UE itself, again removing the eNodeB from scheduling.

The first major paradigm change when implementing the C-V2X Mode 4 standard or NR-V2X Mode 2 standard is that the vehicles autonomously select their own radio resources. This is a significant change considering the eNodeB based design of *SimuLTE* and any LTE-based network. All LTE based communication up until this point had been designed assuming an infrastructure assisted system using a eNodeB to manage scheduling. As such, this required changes in how *SimuLTE* operate with this in mind, a *SimuLTE* eNodeB containes two specific components that are key in managing the network.

The first component was the *Binder*. This manages the LTE network by managing which RBs are assigned to each member of the network. This is especially significant in managing interference, as this module is used to

determine whether or not the RBs were occupied by other transmitters at the same time. In OpenCV2X this component was decoupled from the eNodeB to allow it to be deployed independently and to make it accessible to all vehicles in the network.

The other key component was the *Deployer*, which manages how devices enter and exit the network. As such, it is the component that attaches mobility modules and safely detaches them when devices leave the network; in OpenCV2X this is usually dictated by vehicles leaving the SUMO network. It also manages elements of the LTE network such as the number of available RBs, the bandwidth in the network and their disposition, i.e. uplink, downlink, etc. Ultimately, these components are highly coupled, so there is crossover in their responsibilities. However, the primary implementation effort was to ensure that these binder and deployer components were completely decoupled from the eNodeB and made accessible to the vehicles operating in the network.

The final major element of implementation was from the management of SB-SPS grant generation in the MAC scheduler. In the original model, all grants were generated and managed by the eNodeB, with more details on this being described in Section 3.4.4. OpenCV2X can rather than relying on an eNodeB allow individual vehicle to generate new *SchedulingGrants* as such managing RB allocation, as well as specific new SB-SPS based allocations i.e. Subchannels.

3.4.2 Integration with Artery

For the lower layers, OpenCV2X extends the *SimuLTE* framework [85], which provides a system-level model for LTE. *SimuLTE* was designed to model classical internet services which are all based on IP. This is a major challenge when integrating with *Artery* and its V2X-based applications. Such V2X apps are generally not IP based. In the case of CA and DEN services operate using the GeoNetworking protocol. Geonetworking is a routing protocol based on geographical locations rather than a specific IP address that describes the location of the receiver [87]. For OpenCV2X to digest such packets, the first step was to add a non-IP interface so that the *Artery* Geonetworking packets could be consumed. This also required changes in all layers of *SimuLTE* to ensure that they could operate with new non-IP traffic. Specifically, this required the introduction of *LteFlowControlInfoNonIp*. This replaces the *LteFlowControlInfo* which is control information used by *SimuLTE* to manage

how a packet progresses through the layers and manages aspects such as header sizes and source and destination ports, as well as IP addresses. As IP addresses are now removed and replaced with non-IP based communication, this had to be changed to use addresses based on those in Artery or CA/DEN address system. Furthermore, all layers were adapted to ensure that CBR could be reported up the layers from the lower layers of OpenCV2X through the radio driver to *Artery*, this enables the use of DCC mechanisms deployed in *Artery*.

The integration of *Artery* and OpenCV2X also required the implementation of a new radio driver. The radio driver acts as the direct interface between the LTE NIC as shown in Fig. 3.2 and Artery itself. It is in this module that Artery consumes CBR packets from Simulte for use in DCC mechanisms. It also ensures that Artery packets are configured correctly with *LteFlowControlInforNonIp* attached to each packet passed to the NIC.

3.4.3 Veins based simulation

For *veins*-based simulation the process of integration is much simpler. As part of *SimuLTE's* D2D release 12 model, it had already integrated with *Veins* to do V2X-based simulations. These simulations were based on the use of an eNodeB for scheduling and adopted simple applications and message types purely to test network performance. Even so, this integration enabled easy progress with veins, whereby once the simulation model was developed, it was available for use with veins by default. The only additional element was a new Mode 4 application model, since this version would not have the CA/DEN services. This application was modelled to mimic the performance described in 3GPP TR 36.885 [12], which means that the packets were transmitted periodically (every 100ms) and had a size of 190B or 300B with a pattern of 4 190B packets followed by a 300B packet.

3.4.4 MAC layer

In this section, the second major implementation area related to the MAC layer will be discussed, as shown in Fig. 3.2. To implement this required the development of a new MAC class specifically *LteMacVueMode4*. This class replaces the *LteMacUeD2D* class which was introduced in *SimuLTE* based on 3GPP Release 12 D2D Proximity-based Services [86]. This will be split into

three components firstly how the MAC layer manages the Semi-Persistent Scheduling element of SB-SPS, then how congestion control is managed, and finally how adaptive MCS is modelled.

The implementation areas are as follows:

- The most significant area of implementation is the Semi-Persistent element of the SB-SPS algorithm which is managed by the MAC layer and is described in Section 3.4.4.1.
- The next element and a major focus of this study is the implementation of various congestion control mechanisms both 3GPP and ETSI standard approaches outlined in Section 3.4.4.2.
- Finally, a discussion of the impacts of adapting the MCS and the accurate modelling of this in Section 3.4.4.3.

3.4.4.1 Semi-Persistent Scheduling in SB-SPS

The primary objective of the *LteMacVueMode4* class is to manage the Semi-Pesistent element of SB-SPS, which refers to the grant generation and maintenance mechanisms in scheduling. As referred to earlier in Section 3.4.1, the MAC layer needs to be able to generate its own scheduling grant. There are also significant changes in the nature of this scheduling grant, with the introduction of the concept of subchannels rather than simply assigning individual RBs. This change and considering autonomous resource selection required the introduction of a new *LteMode4SchedulingGrant*. This new scheduling grant maintained all the information required for transmission of a packet in Mode 4, such as the selected subchannels, MCS, RRI and retransmission time gap. This class is exchanged between the MAC and PHY layers for the management of both grant generation and packet transmission.

Grant generation occurs when a new packet is received from the RLC layer and there is currently no grant already available as shown in Fig. 3.3 with the *newDataPkt*. The MAC layer will then endeavour to generate a grant capable of transmitting the packet. To do so, the MAC layer will refer to the subchannel and MCS configurations. This is controlled through specific parameters such as *numSubchannels*, *subchannelSize* which dictate the number of subchannels in the scenario and their size in RBs. Grant selection is controlled by the *sidelink_configuration.xml* file contained in the simulation folder. This file dictates the number of subchannels which can be selected, the available MCS',

and finally the RRIs. The MAC layer will then determine, based on the current packet size provided by the RLC layer and the available MCS and subchannels, the optimal grant size. It does so by referring to tables 7.1.7.2.1-1 and 8.6.1-1, provided in [22] and selecting the minimum number of subchannels required to fit the packet based on the MCS currently configured. This ties into MCS adaptation which is part of the standard, but will be discussed in more detail in Section 3.4.4.3. Once a selection is made, a *LteMode4SchedulingGrant* is generated and shared to the PHY layer as shown in Fig. 3.3, details on CSR generation will be discussed in Section 3.4.5. The PHY layer then reports the CSRs available and the MAC layer simply selects one at random and updates the scheduling grant accordingly. The final element of this process is selecting the RC, which determines the length of time the grant is maintained; typically, this is between 5 and 15 transmissions and is selected at random.

This grant is then used in transmission, the transmission times are maintained by it, and the RC is decremented each time. On grant transmission time, the MAC layer requests the packet from the RLC layer, which is the same process as *SimuLTE* would originally operate. Then the MAC layer passes both the packet and the scheduling grant describing the resources the PHY layer should use for transmission. Once the grant has been completed and depending on the *ProbResourceKeep* parameter (which can ensure that the grant is maintained for a further 5-15 transmissions), the MAC layer will inform the PHY of its completion so that the next SCI will inform the neighbouring vehicles of the termination of the grant.

The MAC layer is also able to dynamically manage the RRI throughout the simulation, allowing users to investigate different approaches *RRI Adaptation* for congestion control but also in the management of aperiodic application traffic patterns. The variant related to the CAM IPT prediction, which forms contribution 3, enables the upper layers to provide a new RRI than the currently configured approach through the new *LteFlowControlInfoNonIp*. The MAC layer will then dynamically update the RRI and inform the PHY layer to update the RRI in the sidelink configuration, as well as update the currently configured sidelink grant considering the new RRI. The variant for congestion control will be described in the subsequent section dealing specifically with the different congestion control approaches.

3.4.4.2 Congestion Control

The MAC layer is also a key component in the management of congestion control mechanisms. Importantly, the MAC layer calculates the CR of the vehicle as dictated by 3GPP, which enables it to accurately manage the standard congestion control approaches. Through the same configuration file *sidelink_configuration.xml* it is possible to configure the CBR to CR-Limit tables, which are the 3GPP standard for congestion control. These tables do not dictate a specific mechanism by which the CR-limit may be maintained. The standard highlights a few different mechanisms that can be employed as previously described in Section 2.3.7.1.

The aspects integrated into the MAC layer initially were packet dropping based approaches, where simply depending on the CR-limit and current CR packets were dropped until CR was brought into check again. The second available option is also configurable through the *sidelink_configuration.xml* file as well, specifically being an MCS congestion control approach. The file enables the user to define a reduced number of selectable subchannels when in congested scenarios, as well as a higher MCS to ensure packets fit into reduced allocations.

In later stages of this work, other techniques were introduced that are described in Section 2.3.7.2. Ultimately, three different approaches were implemented, including one based on the LIMERIC algorithm [25], specifically the one stated in the ETSI DCC adaptive standard. This is achieved through the aforementioned *RRI Adaptation* technique, which has been configured in the MAC layer and enables the user to dynamically update the RRI throughout time and still manage the SB-SPS sidelink correctly. As such, depending on the configured scenario, the MAC layer will enable one of the techniques, to deal with each in turn.

The first mechanism is $RRI_{Reactive}$; in this case it is a table based look-up mechanism as described in Section 4.3.2.1. The MAC layer on reception of a CBR update from the PHY layer will lookup the current expected transmission rate from the table and this will inform the RRI. Should congestion result in the current rate of transmission exceeding the expected rate, then the MAC layer will send the current packet with an updated RRI, indicating the new allowed rate. The current RRI is then maintained until the next transmission opportunity. At this point, any application packets received from the upper layers will be kept in the MAC layer prior to transmission. If an updated

application packet is received, then the currently queued packet is dropped, and the new one stored at the MAC layer. This approach ensures that the latest CAMs are transmitted and avoids the transmission of stale information, ensuring the best age of information. Once the next transmission opportunity arises then the MAC layer can again recalculate the RRI based on the current level of congestion and transmit the latest packet with a new RRI based on current channel conditions.

The second mechanism is RRI_{CR_Limit}; in this case it is again a table based lookup mechanism as described in Section 4.3.2.2. This operates the same way as the previous mechanism except that the table is now the CBR to CR-Limit table so the RRI is calculated differently. In this case, the current CBR is received from the PHY layer, the table will be referred to and the MAC layer will determine the correct RRI to adopt based on the associated CR-Limit. As CR is based both on historical and future usage, this requires the MAC layer to investigate based on the current CR which rate will ensure it reduces to below the limit, as such it will check each rate and recalculate CR based on that rate. The remaining elements focusing on queuing of the packets, dropping stale packets and updating the RRI remain the same as previously described.

One key additional element introduced for both of the previous outlined approaches is the introduction of the DCC averaging mechanism, this operates by ensuring the vehicle only changes state depending on a sustained period of congestion. Meaning that for the vehicle to go to a higher level in the table there must be a sustained period of 1 second exhibiting a CBR in above the current level. Then to go to a lower state requires a period of 5 seconds of sustained lower CBR. This approach resolves challenges with any table based approaches, where they exhibit low levels of stability by fluctuating between states in the CBR table due to minor or occasional changes in CBR.

The third mechanism is $RRI_{Adaptive}$; in this case it is an algorithmic approach as described in Section 4.3.2.3. This approach is quite different from the previous two, as it uses the LIMERIC algorithm to determine the current δ or usage that the MAC layer is allowed, this δ is similar to CR but only accounts for the previous usage. The mechanism operates similarly to the previous two in that a CBR update is received from the MAC layer, the current δ is then consulted and an RRI is calculated that ensures the vehicle maintains that limit. The RRI is set accordingly and the mechanism operates the same way, in that it will update the RRI correctly, drop intermediate packets and can update the RRI on

the next transmission opportunity. Another key change with this approach is that the δ parameter is only updated every 200 ms rather than on reception of a new CBR packet or based on the DCC averaging approach, this again is to maintain higher levels of stability and avoid dynamic fluctuations in CBR due to persistent updating.

3.4.4.3 Adaptive MCS

A final important aspect of OpenCV2X is MCS adaptation this aspect of C-V2X was a joint undertaking with researches at UNAM. They provided the PHY layer modelling to ensure accurate calculation of all channel statistics when MCS was varied, while I provided the functionality at the MAC layer for selecting the MCS dynamically based on packet size.

The MAC layer defined in the 3GPP release enabled MCS adaptation, which means that the MAC layer can dynamically change the MCS employed in encoding the packet and was incorporated in the initial model [88]. This is often overlooked in other models, and assuming a fixed MCS can have a significant impact on performance as will be discussed in Chapter 5. The model was designed to ensure efficient subchannel usage. i.e., the minimum number of subchannels that can accommodate a packet is selected. Once this is determined, the most robust i.e. lowest MCS will be used to transmit the packet. The selection of different MCS configurations modifies the occupation of the RB and in some cases the occupation of the subchannel in C-V2X due to its channelisation structure. Even if the number of required subchannels remains unchanged, the selection of MCS modifies the number of RBs required for transmission in most cases. This has implications on the computation of power levels in the channel and consequently on the operation of the SB-SPS mechanism in C-V2X that are important to consider.

The variations of PSD due to MCS adaptation have an impact on the computation of important parameters such as the RSRP, RSSI and Signal-To-Noise Ratio (SINR). RSRP is computed by obtaining the average power contribution of each Resource Element (RE) within the transmitted RBs [89]. Therefore, if the number of RBs changes due to MCS adaptation, the value of RSRP must be computed accordingly.

PSD variations also have an impact on channel interference levels. For example, a transmission that occupies a lower number of RBs in a subchannel produces less interference than a transmission that occupies the same

subchannel, but utilises a higher number of RBs. This variation in interference has an impact on RSSI, which is calculated as the received power plus the interference and thermal noise in the received SC-FDMA symbol [89]. This is also the case for the SINR, which is computed as the ratio between the received power and the interference plus the thermal noise.

Based on this analysis, when MCS adaptation is implemented, the following factors must be accurately estimated and considered: received power, interference, and thermal noise. This is particularly important for parameters such as the RSRP and RSSI, which are used to perform CSR selection, as well as for the SINR, which is used for decoding transmissions.

The C-V2X SB-SPS mechanism filters out CSRs based on the average RSRP and RSSI of the subchannels in the sensing window. Therefore, several scenarios can arise due to variations in RSRP and RSSI as a consequence of the implementation of the MCS adaptation, with further implications for the CSR selection process.

Figure 3.4 illustrates three potential scenarios depending on different reservations and occupation states of subchannels within the sensing window. Scenario A illustrates the case where two subchannels are reserved and fully occupied by a transmission where the implemented MCS requires the use of all the RBs. As all the RBs in both subchannels are used, they have the same levels of RSRP. This is also the case for RSSI, because occupied RBs contribute the same amount of interference to both subchannels.

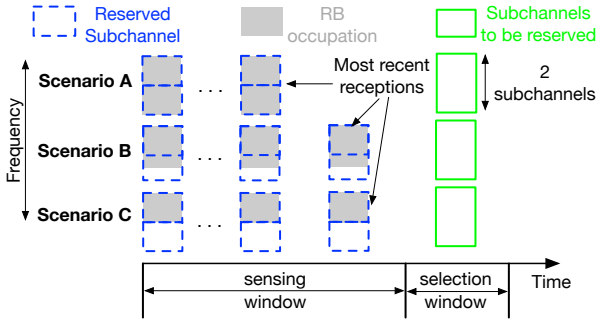

Figure 3.4: Subchannel and RB occupation during CSR selection with MCS adaptation.

However, in the case of Scenarios B and C, the levels of RSRP and RSSI of each subchannel in the sensing window can be different if the MCS is adapted, thus varying the number of occupied RBs. For example, the levels of RSRP in the upper subchannel of Scenario B, where all RBs are occupied, are higher than

those of the lower subchannel. Thus, the upper subchannel has higher levels of RSSI since there is a higher contribution of interference due to the use of a higher number of RBs. This also occurs in Scenario C, where two subchannels have been reserved, but only the upper subchannel is fully occupied due to the implementation of a higher MCS while the reservation is maintained by SB-SPS. In this scenario, the lower channel has effectively zero RSRP since all the transmission power is used in the upper subchannel and also has lower RSSI because there is no interference.

Importantly, these scenarios also have implications during the selection of CSRs, because SB-SPS filters subchannels based on RSRP and RSSI. SB-SPS first discards the subchannels with the highest RSRP measured during the most recent receptions within the sensing window, which is impacted by MCS adaptation as described for scenarios B and C. After this process, SB-SPS selects only 20% of the remaining subchannels with the lowest average RSSI. In this case, the average RSSI of the subchannel depends on the individual RSSI measurements in each subchannel, which can also vary as a function of the selected MCS as shown in scenarios B and C.

Based on this analysis, it is evident that the variations in subchannel RSRP and RSSI due to the adaptation of the MCS have important effects on the selection of CSRs performed by SB-SPS. All these effects have been considered in this study to accurately model the operation of the MCS adaptation in C-V2X.

3.4.5 PHY layer: Sensing in SB-SPS

The third major implementation aspect related to the PHY layer manages the transmission of all packets in the model. Much of this implementation was inherited from the original *SimuLTE* model. Primary implementation aspects are highlighted as follows:

- Sensing window, maintenance, and measurement recording for accurate sensing.
- Management of the CSR selection process including; Selection window generation, all necessary CSR filtering, and reporting to the MAC layer.
- SCI message creation and transmission.
- Introduction of new SB-SPS measurements such as Channel Busy Ratio and it's reporting to upper layers.

The extended implementation is focused on the sensing window and the CSR selection process. The sensing window is involved in many elements of the PHY layer for SB-SPS whether it be Mode 2 or Mode 4 and is a core element of the scheduling mechanism. The sensing window has three primary interactions; the first is on reception of any packet specifically in the *DecodeAirFrame* method. In this method, upon successful reception of any packet, the sensing window will record information in the associated subchannels about the transmission. Specifically, this will be the RRI of the received SCI (for us in CSR selection), as well as the recorded RSRP and RSSI across that subchannel. In the case where a reception is not successfully decoded, the sensing window will still record the received RSSI on the associated subchannel, this is an important aspect of the CSR selection process for C-V2X Mode 4.

The second interaction is based on transmission, whenever a transmission occurs the current subframe in the sensing window i.e. all the subchannels in this section of the sensing window will be updated to log that a transmission occurred. This is due to the fact that the C-V2X and NR-V2X radios experience half-duplex errors. Using both this information and the information from receptions, the CSR selection process can proceed as standardised on the reception of a request from the MAC layer.

The final interaction is in the CSR selection process, which is the new core mechanism implemented at the PHY layer, and forms the basis for the Sensing-Based element of SB-SPS. It is handled by the *computeCSRs* method; on receiving a scheduling grant from the MAC layer describing the desired grant, that is, the number of required subchannels, the PHY layer will go through the steps to generate a list of possible CSRs. This process is fully described in Section 2.3.2 and can be broken down into four main stages. Firstly, the PHY layer will generate a set of all possible CSRs or a selection window based on the MAC layer request. The first filtering stage is based on the half-duplex problem for C-V2X/NR-V2X radios, whereby if a transmission occurs, then reception cannot occur in the same subframe; as such, it is unknown if other subchannels in that subframe are reserved. Due to this, the PHY layer will go through the sensing window and determine the subframes which have been used for transmission, then based on the possible RRIs available in scenario it will eliminate corresponding subframes in the CSR pool.

The second filtering stage is also common to both NR-V2X and C-V2X as well; in this stage, the PHY layer will filter out CSRs based on the SCI received and the RSRP of the subchannels in which they are received. As such, it parses the sensing window, finds reserved subchannels, checks their corresponding subchannels in the selection window (based on the recorded RRI), and removes those subchannels where the RSRP is above the configured RSRP threshold. At this point, depending on the number of filtered resources, there is a retrenchment phase, if too many CSRs have been eliminated, then the RSRP threshold is dynamically lowered and CSRs reintroduced that go under the threshold until at least 20% of all possible CSRs are available for selection. In NR-V2X, this pool of CSRs would be reported to the MAC layer, while C-V2X has one final stage managed by the *selectBestRSSIs* method. In this method, the CSRs are ordered based on having the lowest RSSI; therefore, least interfered subchannels, then the top 20% will be reported to the MAC layer for selection.

The third element is the SCI message, which needed to be introduced to the model. It has its own *SidelinkControlInformation.msg* file which contains all the necessary information as defined in the 3GPP standard [90]. This process requires communication between the MAC and PHY layers, as shown in Fig. 3.3, the MAC layer will share the Scheduling Grant and this will enable the PHY layer to accurately determine the correct subchannels, the RRI to use and the MCS required for transmission. On reception of the packet, it will then send the SCI followed by the packet itself. It also contains the process for managing their reception and recording the necessary information in the sensing window when the packet is received correctly.

3.4.6 Channel Modelling

This section describes the necessary additions for accurate channel modelling in the C-V2X and NR-V2X standards in compliance with standards [12]. The *LtePhyVueMode4* also leverages the *LteRealisticChannelModel* class in *SimuLTE* to model the transmission channel in C-V2X, this is represented by the Radio module of Fig. 3.2. The required changes are outlined as follows:

- The introduction of new PHY layer channel measurements as described in Section 3.4.6.1.
- Section 3.4.6.2 describes the newly introduced channel and fading models, as well as a sensing model from the analytical study used to

benchmark this work [30].

3.4.6.1 PHY layer channel measurements

Due to the significant paradigm shift in moving to autonomous resource selection and removing the eNodeB OpenCV2X needed to incorporate significant extensions in terms of radio modelling.

In the current version of *SimuLTE LteRealisticChannelModel*, packet transmits as groups of RBs and every RB is allocated the total transmission power. This transmission model is generally accurate as long as the transmission bandwidth is maintained fixed during the simulation and all power ratios are clearly defined. Furthermore, in *SimuLTE*, receiver transmissions are evaluated on an RB basis by computing the SINR of each RB and mapping the result to a Block Error Rate (BLER) value. Importantly, all interferers are assumed to use the total transmission power for each RB, and thermal noise is set as a function of the total bandwidth.

However, in C-V2X Mode 4 the transmission bandwidth can be highly variable due to the variety of subchannel configurations, MCS' and packet sizes. As a result, OpenCV2X implements a solution where, rather than every RB being allocated the total transmission power, it distributes all power levels, i.e. the transmission power, interference, and thermal noise, across the number of RBs required to transmit each packet. This is vital to the correct calculation of key parameters of the specific physical layer of C-V2X, namely RSRP and RSSI, since both parameters depend directly on the number of transmitted RBs [91]. Furthermore, it impacts the correct calculation of PHY layer measurements, such as SINR.

The size of the TB defines the number of RBs required to transmit a packet considering the Cyclic Redundancy Check, payload, and redundancy bits. In addition, since we consider the case of adjacent PSCCH+PSSCH transmissions, the 2 RBs needed to transmit the SCI are added to obtain the total number of RBs to be transmitted. At the receiver, the total received power (after accounting for pathloss and shadowing) is divided by the number of transmitted RBs to obtain the received power per RB. This is then divided by the transmission bandwidth to obtain the PSD of each RB. The PSD is the building block on which key C-V2X PHY layer measurements such as RSSI, RSRP, and SINR are computed.

To compute the RSRP, the PSD of the RB is multiplied by the bandwidth of each RE. We assume the PSD to be constant within each RB and, therefore, the power per RE to be the same for all REs. With this assumption, the RSRP, defined as the average power contribution of each RE, is equal to the power per RE. For the RSSI, the power received is obtained per SC-FDMA symbol. To do this, we consider all power sources, i.e. the power, interference, and noise per RB. The power received per SC-FDMA symbol is then calculated as the sum of the total power received per RB multiplied by the total number of RBs used to transmit.

Finally, we compute the SINR of each RB as the ratio between the received power and the power of the interference plus the total noise, all on an RB scale. To compute interference, we account for the PSD per RB from all transmitters using overlapping RBs during an ongoing transmission. By using the PSD per RB and the number of RBs used by each interferer, we can accurately compute the interference. The total noise is also estimated on the basis of the number of RBs used for transmission by adding the noise figure at the receiver to the thermal noise. Thermal noise is computed by multiplying the thermal noise density specified in the LTE standard of -174 dBm/Hz [92] by the number of transmitted RBs and the RB bandwidth.

The second significant change that OpenCV2X incorporates is related to the calculation of error rates. OpenCV2X calculates error rates per packet, in contrast to SimuLTE, where transmissions are evaluated per RB by computing the SINR of each RB and mapping the result to a BLER value, based on the Channel Quality Indicator. SimuLTE computes the SINR by assuming that all interferers use the total transmission power for each RB, and the thermal noise is set as a function of the total bandwidth. This would result in incorrect C-V2X PHY calculations. In contrast, OpenCV2X inputs the mean SINR for the transmitted RBs into the Mode 4 BLER curves defined in [93] to obtain the packet error rate using the curve that corresponds to the MCS of each transmission.

3.4.6.2 Channel Model implementation

To comply with the simulation guidelines specified by 3GPP in [12], we modified *LteRealisticChannelModel* by implementing the WINNER + B1 channel model; see Algorithm 3.1 and the Block Error Rate curves described in [93]. Furthermore, we modified the computation of power levels and

implemented new physical layer parameters, specifically required by C-V2X mode 4 and defined in the 3GPP release 14 standard. The following section describes these modifications in detail.

Algorithm 3.1 WINNER + B1 LOS

1: $dBP = 4 * (heightAntennaA - environmentHeight) * (heightAntennaB - environmentHeight) * carrierFrequency/c$
2: **if** $distance < 3$ **then**
3: $distance = 3$
4: **end if**
5: **if** $distance < dBP$ **then**
6: $pathLoss = 22.7 * std :: log10(distance) + 27 + 20 * std :: log10(carrierFrequency/1e9)$
7: **else**
8: $pathLoss = 40 * std :: log10(distance) + 7.56 - 17.3 * std :: log10(heightAntennaA - environmentHeight) - 17.3 * std :: log10(heightAntennaB - environmentHeight) + 2.7 * std :: log10(carrierFrequency/1e9)$
9: **end if**
10: $pathLossFree = 20 * std :: log10(distance) + 46.4 + 20 * std :: log10(carrierFrequency * 1e-9/5)$
11: $pathLoss = max(pathLoss, pathLossFree);$

Another introduction was the nakagami fading model [94] to complement the Rayleigh and Jakes fading already available in *SimuLTE*.

A final added introduction was a new unsensed error model, this was necessary as part of the validation process. The analytical model [30] used in the validation of OpenCV2X uses an unsensed error model. This represents the scenario where RSRP is below the sensing power threshold of the radio employed and is calculated depending on the received power and the experienced attenuation between the transmitter and receiver; this is shown in alg. 3.2.

3.5 Model validation including Default Simulation Parameters

It is vitally important that a model accurately reflect the technology it is set out to mimic; as such a key process in the development of any model is to benchmark it against a standard or empirical dataset to show that it is reflective of the real world performance of the technology. The subsequent

Algorithm 3.2 Analytical unsensed errors

1: $erfParam = (receivedPower - attenuation - -90.5)/(3 * sqrt(2))$
2: $erfValue = erf(erfParam)$
3: $packetSensingRatio = 0.5 * (1 + erfValue)$
4: $er = dblrand(1)$
5: **if** $er < packetSensingRatio$ **then**
6: $\quad Packet\ Unsensed$
7: **else**
8: $\quad Packet\ Sensed$
9: **end if**

section describes the validation of the OpenCV2X model. This was conducted jointly with UNAM. While I developed, tested, and ran all aspects in the validation stages, as well as analysing results and being the primary writer on this work. The researcher at UNAM advised on the PHY aspects of the validation, and we would have collaboratively overcome such challenges when carrying out the validation stage.

To validate the OpenCV2X implementation, we benchmark against an analytical LTE-V Mode 4 model made available by Gonzalez-Martin et al. [30] in Matlab. We examine the performance of Mode 4 while observing parameters that affect the communication performance, such as vehicular density, transmission power levels, and MCS.

Section A provides a description of how results are generated for each set of results, it can be assumed in all cases that this process is followed for all results. In the case where this is not applied the section will highlight the specific changes made for said experiments. The road environment and vehicular densities considered are summarised in Table 3.1, and can also be seen in Fig. 3.5 to provide a visual comparison of the various densities. Other key simulation parameters are summarised in Table 3.1, and the selection of these is to allow for a direct comparison with the analytical model for validation purposes.

3.5.1 Quantitative Validation

Fig. 3.6 shows the performance of the OpenCV2X Packet Delivery Rate (PDR) (solid lines) compared to the analytical model (dashed lines). This considers P_{Tx}=20dBm, F_{Tx}=10Hz, MCS 7 (QPSK 0.5, 2 subchannels) and four vehicular densities, β=0.06, 0.12, 0.2 and 0.3 veh/m. The lower vehicular densities

Table 3.1: Default simulation parameters.

Parameter	Value
Vehicular scenario	
Vehicular density	0.06, 0.09, 0.12, 0.16, 0.2, 0.25, 0.3 veh/m
Road length	2 km
Number of lanes	3 in each direction (6 in total)
Lane width	4 m
SUMO step-length	1ms
Channel settings	
Carrier frequency	5.9 GHz
Channel bandwidth	10 MHz
No. subchannels & Size	2 & 12 RBs
Adjacency PSCCH-PSSCH	Adjacent
Application layer	
Packet size	190 Bytes
Transmission frequency (F_{Tx})	10 Hz
Artery update interval	1ms
MAC & PHY layer	
Resource keep probability	0
RSRP threshold	-126 dBm
RSSI threshold	-90 dBm
Propagation model	Winner+ B1
MCS	7 (QPSK 0.5)
Transmission power (P_{Tx})	10, 20 and 23 dBm
Noise figure	9 dB
Shadowing variance LOS	3 dB

were chosen to comply with the recommended 3GPP C-V2X simulation guidelines, specifically the Fast and Slow scenarios on the highway [90], while the higher density scenarios are closely compared to the performance of the analytical model [30]. The 3GPP Highway scenarios result in a mean vehicular density of β=0.06 and 0.12 veh/m respectively. We discard the first 500 seconds of the simulation to discount the vehicles entering the simulation.

It can be observed in Fig. 3.6 that the simulated PDR aligns closely with the analytical model with only a small divergence in the PDR, regardless of the vehicular density considered. Table 3.2 considers further densities to validate the precision of the simulation model, by examining the mean absolute deviation in the PDR compared to the analytical estimate. In all cases, it can be seen that it differs by less than 2%.

To ensure that the simulated and analytical models are considering

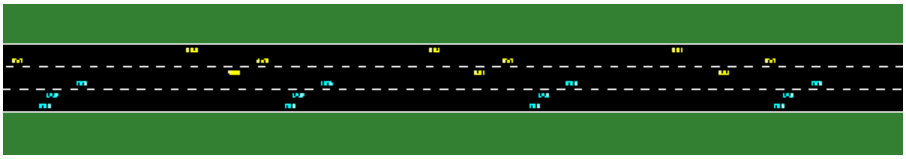

(a) Density of β=0.06 veh/m.

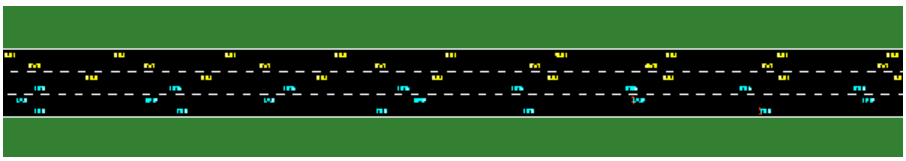

(b) Density of β=0.12 veh/m.

(c) Density of β=0.2 veh/m.

(d) Density of β=0.3 veh/m.

Figure 3.5: Vehicular densities considered in OpenCV2X validation.

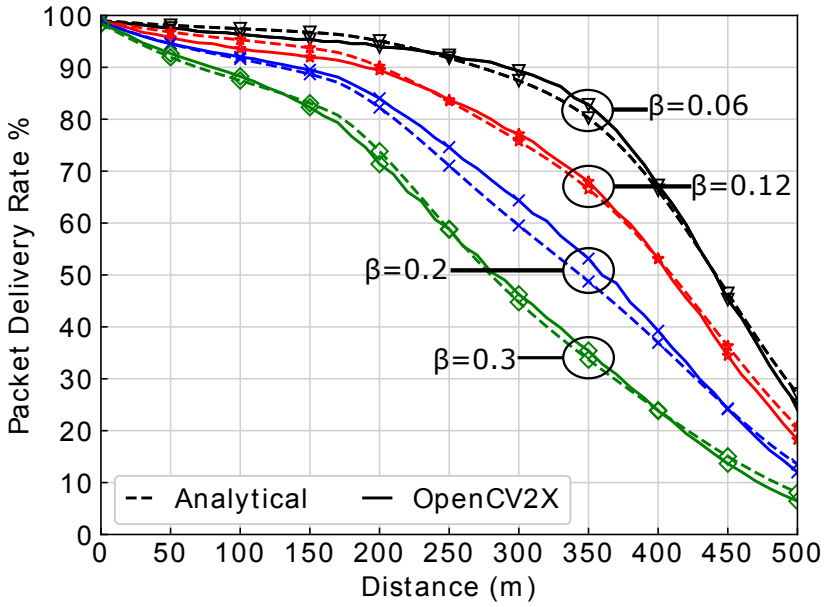

Figure 3.6: PDR as a function of the distance between the transmitter and receiver for varying vehicular densities for OpenCV2X versus analytical model.

Table 3.2: Mean absolute deviation in PDR between OpenCV2X and the analytical model for varying vehicular densities.

P_{TX} (dBm)	F_{TX} (Hz)	β (veh/m)	PDR (%)
20	10	0.06	1.20
		0.09	1.35
		0.12	1.15
		0.16	1.21
		0.2	1.84
		0.25	1.48
		0.3	1.07

comparable channel characteristics and behaviour, we further examine the channel load by measuring the CBR, which is the number of subchannels sensed as occupied, as well as the causes of packet loss. In Fig. 3.7a, we can see that both models consider comparable channel loads with a mean deviation up to 6%. Note that the analytical model produces a single CBR estimate. The occupied subchannels are identified as those filtered by exceeding the *RSRP* threshold with a mean CBR value estimated later for all vehicles in the simulation. In contrast, OpenCV2X calculates CBR according to the 3GPP definition [91], with each vehicle determining the percentage of subchannels in the last 100 ms that have recorded a *RSSI* greater than a pre-configured threshold of -90dBm. This, along with minor differences in the distribution of vehicular density, accounts for the minor divergence in channel load. In Fig. 3.7b, we investigate the causes of packet loss for the scenario $\beta = 0.12$. As outlined in [30], packet loss can be attributed to the following, and for clarity of comparison, we use the same notation:

- Half Duplex Errors, δ_{HD}: Vehicles receive a packet on the same subframe on which they are transmitting and, therefore, are unable to decode the incoming packet.
- Sensing Errors, δ_{SEN}: A received signal does not have sufficient power to meet the sensing power threshold due to attenuation caused by path loss and shadowing.
- Propagation Errors, δ_{PRO}: A propagation error occurs when the received power of the transmitted packet is higher than the sensing power threshold, but the Signal-to-Noise Ratio (SNR) is insufficient to correctly decode the packet.
- Packet Collisions, δ_{COL}: Occurs when another vehicle transmits on the

same subframe and subchannel; this interference affects reception at the receiver by reducing SINR.

In particular, for all cases in Fig. 3.7b, the models aligned with little divergence. Similarly in Table 3.3, the mean absolute deviation of packet loss errors between models is considered, for varying vehicular densities. δ_{HD} errors account for only 1% of all packet loss in both models and incur 0.06% divergence in such cases. δ_{SEN} errors are based on the transmission power and the attenuation incurred by the packet due to path loss. It is a probability-based mechanism that was implemented to match the analytical model, and as such the divergence is less than 0.1%. On the basis of the calculated attenuation, the received packet is only decoded once it exceeds the sensing power threshold of -90.5dbm. As expected, these become the dominant cause of packet loss at greater distances. δ_{PRO} errors are negligible in the simulation and of those that occur there is minimal divergence from the analytical model of less than 0.36% in all cases.

δ_{COL} errors are the other dominant source of packet loss. These occur more frequently as the distance between the transmitter and receiver increases as a result of increased attenuation. At distances exceeding 400m, these are eclipsed by sensing errors, as the incoming packets do not reach the sensing power threshold. A marginal divergence of less than 2% was observed for all evaluated scenarios, as shown in Table 3.3. The divergence is attributed to a slightly lower CBR in OpenCV2X as observed in Fig. 3.7a, resulting in slightly less collisions and a minor improvement in PDR. This is attributable to the density and mobility of vehicles. While the mobility scenarios were parameterised to have a density corresponding to the analytical model, the OpenCV2X model does not have a fixed density and as such it is not completely uniform. Furthermore, we devised these experiments to comply with the 3GPP Highway road network comprising 6 lanes (3 in each direction), whereas Gonzalez-Martin et al. considered a 4 lane highway. As such, the density and subsequently the CBR considered in the analytical model are marginally higher.

Next, we considered the impact of transmission power and MCS on the PDR as shown in Fig. 3.8a. Three transmission powers P_{Tx}=10dBm, 20dBm and 23dBm, have been considered. 10dBm and 23dBm are the minimum and maximum transmission power limits specified by the 3GPP for C-V2X sidelink communications with 23dBm used as the default. 20dBm is considered for

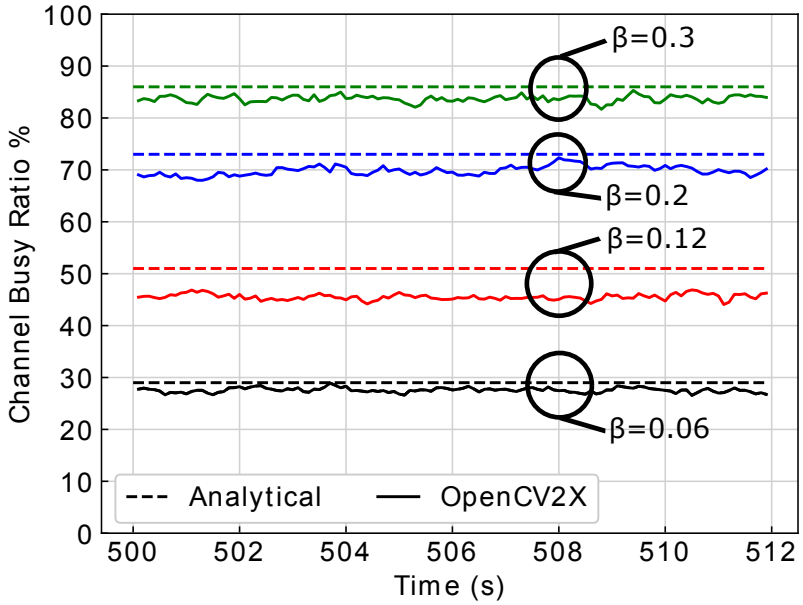

(a) PSSCH CBR for varying vehicular densities.

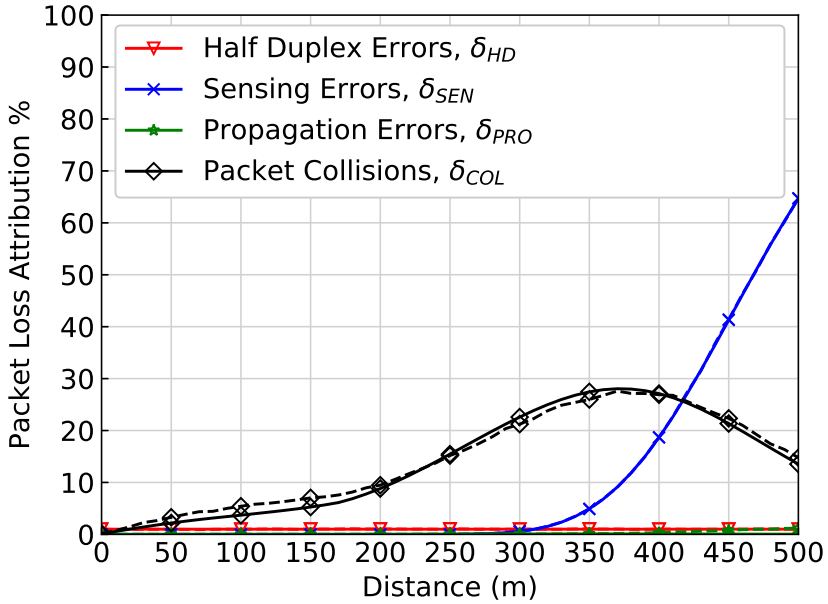

(b) Packet Loss Attribution per type of error (β=0.12 veh/m) for OpenCV2X (solid lines) and the analytical model (dashed lines).

Figure 3.7: Channel conditions including (a) CBR & (b) packet loss attribution.

Table 3.3: Mean absolute deviation in packet loss attribution between OpenCV2X and the analytical model for varying vehicular densities.

P_{TX} (dBm)	F_{TX} (Hz)	β (veh/m)	δ_{HD} (%)	δ_{SEN} (%)	δ_{PRO} (%)	δ_{COL} (%)
20	10	0.06	0.05	0.08	0.36	1.12
		0.09	0.05	0.05	0.22	1.20
		0.12	0.03	0.04	0.19	0.98
		0.16	0.02	0.04	0.16	1.12
		0.2	0.02	0.02	0.13	1.80
		0.25	0.02	0.02	0.09	1.44
		0.3	0.02	0.02	0.07	1.01

comparison purposes with [30]. The ETSI ITS specification [15] includes a minimum transmission power of 10dBm to alleviate interference with CEN DSRC tolling systems. For all P_{Tx} values, the results are closely aligned with only negligible differences. As expected for the lowest value of P_{Tx}, the PDR is significantly affected by the reduced radio range as the distance between the transmitter and receiver increases.

In Fig. 3.8b the impact of MCS is considered, with MCS 7 (2 subchannels of size 14 RBs) and MCS 9 (4 subchannels of size 12 RBs) evaluated. This assumes a vehicular density of $\beta = 0.12$ Minimal divergence is observed between the two models. Using a higher MCS reduces the resources required to send a packet and can be a means of reducing channel congestion, thereby improving PDR. The trade-off is that such transmissions are less robust and more prone to loss.

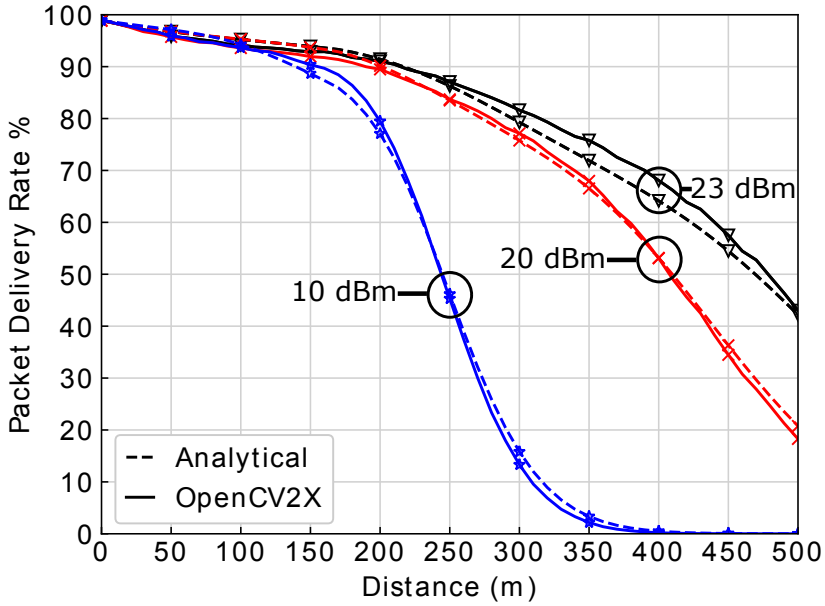

(a) PDR as a function of distance for variable transmission power.

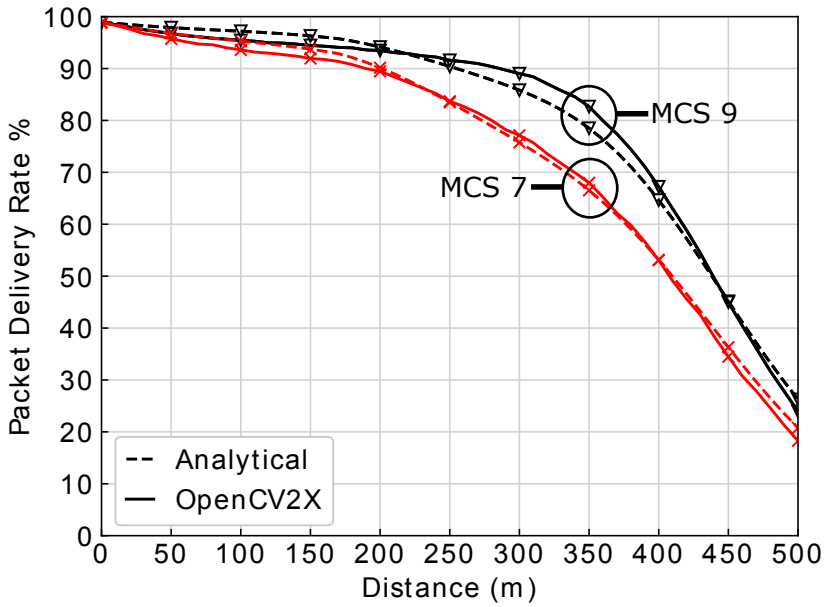

(b) PDR as a function of distance for variable MCS.

Figure 3.8: Impact of varying transmission power and the modulation & coding scheme on PDR.

3.6 Conclusion

This chapter has introduced the **first contribution** (**C1**) of this work, namely OpenCV2X, the first open source full-stack implementation of C-V2X in a discrete event simulation environment. This relates specifically to publications *4* and *7* but was the environment employed in the subsequent investigation.

The following main areas have been highlighted:

- The chapter set out the major requirements of the model, namely to provide a full-stack 3GPP compliant implementation, integrating the current ETSI standard applications, and ensuring the implementation was validated as accurate.
- To this point the model extended a number of different modules and introduced significant new implementation, the following outlines the major areas of change and where they took place:
  - *SimuLTE* required the most significant development:
    - The introduction of Non-IP traffic, requiring changes throughout the entire stack to ensure ETSI applications can operate correctly.
    - The decoupling of significant modules including the Binder and Deployer which control the RB allocation mechanism and UE positions.
    - Removing the eNodeB from the simulation and enabling the MAC layer to generate *LteSchedulingGrants* independently.
    - An entirely new MAC layer, including all the Semi-Persistent elements of SB-SPS, as well as the introduction of the concept of subchannels, different congestion control mechanisms and MCS adaptation.
    - An entirely new PHY layer, including all the Sensing-Based elements of SB-SPS, sensing windows, selection windows, SCI generation, CSR selection.
    - Related to the PHY layer, the introduction of new channel measurements for SB-SPS namely RSSI for RSSI filtering stage of C-V2X, as well as CBR and CR for congestion control.

        * Significant changes in channel modelling, including new channel models, sensing errors and the introduction of new BLER curves.

    - *Artery* to integrate with artery required the introduction of a new Radio Driver ensuring packets could be translated to the new radio technology.
    - *Veins* the final integration required the implementation of Mode 4 specific applications while enabling veins to manage the mobility of the vehicles in OpenCV2X.

- The final element of this chapter was a thorough validation of the model compared with the analytical model in [30]. This ensures OpenCV2X is accurately reflecting the technology itself.

Subsequent chapters will adopt this model in investigating the C-V2X and NR-V2X standards and will enable the remaining contributions of the project.

Chapter 4

Research Problem Analysis: Theorised Discussion, Proposed Solution Overview & Implementation

4.1 Introduction

The previous chapter described OpenCV2X the first full-stack, open source, validated C-V2X model implemented in a discrete event simulator. This model enables a quantitative investigation of the scheduling challenges faced by the C-V2X and NR-V2X standards. Specifically, this chapter investigates the performance challenges that both technologies face when supporting vehicular services linked to their underlying scheduling mechanism, namely SB-SPS. The challenges in terms of scheduling and specifically grant maintenance due to application traffic patterns e.g arrival rates or packet sizes, as well as congestion control mechanisms will be identified. Poor maintenance of the scheduled grant can result in MT and GB, which can significantly degrade the performance of SB-SPS in both technologies. Having explored what both MT and GB are, their causes, and ultimately their impacts, this chapter will then explore solutions to these challenges.

A theoretical discussion of the proposed solutions is provided including, how they overcome the challenges of both MT and GB, and implementation of said solutions in OpenCV2X. The analysis of their performance is explored in

Chapters 5 and 6. Both GB and MT can be caused for numerous reasons, and as such the solutions are categorised according to those causes i.e. aperiodic application traffic arrival rates, packet dropping TRC congestion control mechanisms, which exhibit issues similar to aperiodic traffic patterns, and finally variable packet sizes.

This chapter will:

- Firstly, provide an in-depth theorised analysis of the challenges posed to the C-V2X and NR-V2X SB-SPS scheduler by MT and GB. This is described in Section 4.2.1. It will then investigate the scenarios that cause these in Section 4.2.2 i.e. aperiodic packet arrival rates, congestion control techniques and variable packet sizes. Finally, a quantitative evaluation and analysis of the negative impact on the SB-SPS scheduler is provided in Section 4.2.3.1.
- Secondly, a set of novel solutions is provided with the high-level overview of the solution proposed in Section 4.3.1 and its specific implementation with respect to congestion control techniques and aperiodic arrival rates discussed in Sections 4.3.2 and 4.3.3 respectively.
- The final aspect of this chapter in Section 4.3.4 focuses on the effective management of variable packet sizes.

4.2 Research Problem Analysis

This section theorises and quantitatively investigates the scheduling challenges faced by the C-V2X and NR-V2X standards. These technologies were envisaged after ETSI standardised many of the core "Day 1" V2X services, such as CA and DEN. This means that the new standards will need to be able to operate effectively with such applications but also to manage the diverse scenarios which can occur in vehicular networks. Another challenge they face is congestion; there is limited dedicated spectrum, vehicular networks are highly dynamic in nature, and density can vary significantly over a short period of time. A simple example being merging from a quieter rural road onto a large motorway, the dynamics in congestion may change significantly. This being especially important given the nature of traffic having rush hours with dense road traffic congestion likely causing dense V2X network congestion.

Both the underlying applications and the scenarios in which they are partaking

can have a significant impact on SB-SPS scheduling performance. This section aims to highlight the challenges these applications and scenarios pose to the scheduling performance, focusing on GBs and MTs. These are two issues that can significantly impact the effectiveness of the scheduling mechanism. This work will quantify these phenomena, the possible scenarios in which they may occur, and finally the impact they may have on scheduling performance.

4.2.1 SB-SPS Challenges: Grant Breaking & Missed Transmissions

A primary challenge facing SB-SPS is the impact of GB and MT. As described in Section 2.3, SB-SPS is a grant-based mechanism. The grant is the core element which ensures that the sensing can work correctly; by maintaining the grant through the RRI field of the SCI message, i.e. indicating when the radio resource will next be used. With this, other vehicles can accurately determine which resources are reserved in the future and intelligently select resources.

However, MTs or GBs, have a heavy negative impact on the accurate behaviour of the grant, which will be explored in the subsequent section. The scenarios that cause both GB and MT will be discussed in detail in Section 4.2.2. In short, MT or GB can be as a result of variable packet sizes (in this case, a GB is forced), aperiodic packet arrival rates, and TRC based congestion control techniques (in these cases, a GB or MT can occur). The causes of this are that either the current grant cannot accommodate the packet, which forces a GB, or there is no packet available at the allocated transmission time, which can result in a GB or MT.

4.2.1.1 Missed Transmissions (MT)

The first phenomena to be discussed are missed transmissions and their possible impacts. It is important to note that this is ultimately configurable and that the standard may reduce the allowed number of MTs. What determines if this will occur is the *sl-reselectAfter* parameter. This is configurable in SB-SPS and determines the number of missed transmissions that can occur before reselection or a grant break must take place. This work focused on the scenario where the grant was maintained regardless of missed transmissions (as such *sl-reselectAfter* is set higher than the grant RC and this is referred to as not grant breaking (No-GB)) or when a grant is broken after a single missed

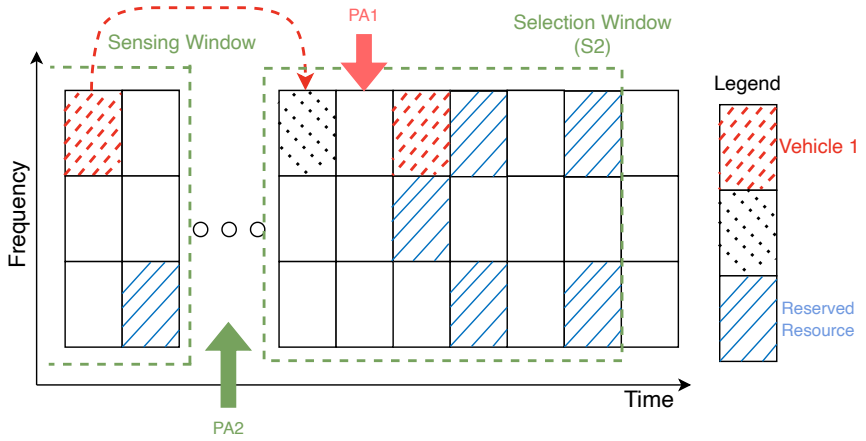

Figure 4.1: Impact of Missed Transmission on SB-SPS Scheduling

transmission (*sl-reselectAfter* = 1 and grant breaking (GB)). In general, in the literature, the not grant breaking configuration is assumed most commonly.

Figure 4.1 shows the scenario in which a transmission is missed. This example assumes two vehicles V1 and V2, and shows the channelisation scheme across frequency and time domains, with the rectangular box representing an individual subchannel in a given time slot. Red indicates the subchannel reserved for V2 as part of its SB-SPS grant. Blue indicates subchannels reserved by vehicles (other than V2), and Clear represents free resources/subchannels. V1 has an SB-SPS grant which is being used, shown in red marking on the individual subchannels. At the current time, V2 receives a packet (PA2) indicated by a green arrow. As it does not have a grant it commences the selection process by initialising a selection window (S2). S2 contains one resource that is ultimately missed by V1 indicated by the black dots on the subchannel. This might be the result of aperiodic arrival. Due to the fact that an RRI was sent indicating the reservation of that resource (red dashed arrow), V2 will not consider that resource as available (though it actually turned out to be free). This is one of the issues with missed transmissions, unutilised reserved resources. This artificially reduces the number of available resources, which increases the likelihood of collisions in selection.

This is further compounded by a subsequent problem. V1 receives a new packet (PA1) indicated by the red arrow. V1 transmits this on its second red-labelled subchannel. However, because V2 did not send on its last reserved resource (black dot) there never was an SCI sent. V2 has no idea that this will be used and may select the second red transmission resource thus causing a

collision, thus increasing interference and incurring a half-duplex error.

Both of these issues result in poor CSR selection due to the inaccuracies in the sensing window. The missed transmissions themselves resulting in a reduced number of selectable CSRs as they appear to be reserved. This results in certain CSRs becoming actual colliding resources which is unknown to neighbouring vehicles due to the lack of an SCI.

In C-V2X this has a more significant impact than in NR-V2X. This is mainly due to the RSSI filtering stage as described in Section 2.3.2. In C-V2X, because only 20% of the CSRs are available for selection, any reduction in available CSRs will increase the likelihood of collision to a greater extent. In contrast in NR-V2X the removal of this stage means that while MTs can reduce the overall number of available CSRs the impact is diluted due to a greater pool being selectable in general.

As such, if MTs occur frequently, then the performance of the scheduler will degrade significantly. Addressing this challenge is one of the core contributions of this work.

4.2.1.2 Grant Breaking (GB)

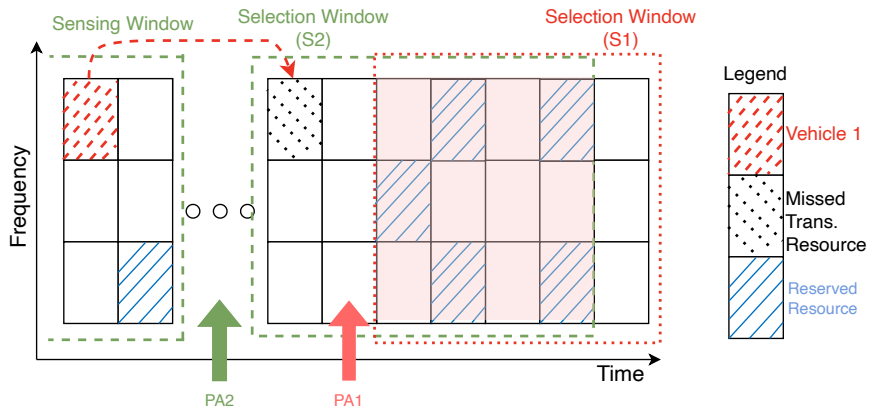

Figure 4.2: Impact of Grant Breaking on SB-SPS Scheduling

Figure 4.2 illustrates the case where grant breaking (GB) is enabled in SB-SPS, i.e a *sl-reselectAfter* parameter of 1. The example is similar to that which is previously outlined except when V1 misses a transmission rather than maintaining the grant, it will instead relinquish the resource and will only generate a new grant when a new packet arrives. V2 receives a packet (PA2) indicated by the green arrow and, starting the grant process, initialises the selection window S2. As V1 did not transmit a packet in its scheduled resource

(black dotted), V1 now relinquishes the grant (due to the MT) and will need to select new resources on reception of a new packet. When this packet arrives (PA1) V1 sets up a new selection window S1 (red dotted box). As V1 and V2 now have overlapping selection windows (S1+S2) shown in the shaded area of Fig. 4.2, this can significantly increase contention in the network as discussed in [95].

Grant breaking in general is a more severe scenario in SB-SPS. This is because it incurs similar issues to MT as described before, i.e. resources reserved that go unused. This can result in valid CSRs being viewed as reserved and reducing the overall number of possible CSRs that a vehicle can select from. The increased number of grant selection processes due to a grant ending before the end of the RC (i.e., after 5-15 transmissions) increases contention in the network. The reduced grant length and the number of MTs result in worse sensing window information, as such, the SB-SPS algorithm performance degrades when the natural grant mechanism breaks down and becomes much less effective. This is more impactful in C-V2X again due to the RSSI filtering stage, as neighbouring vehicles will have similar sensing information and the top 20% of CSRs are likely to be similar pools. The early grant breaking results in frequent reselection processes, and this reduced selection pool results in increased collisions. The elimination of this stage in NR-V2X results means that the CSR pool which can be selected from is much larger, reducing the likelihood of collision but not entirely eliminating it.

4.2.2 Scenarios that cause Grant Breaks & Missed Transmissions

There are a number of applications or scenarios that may cause a GB or a MT to occur, with certain applications or scenarios exhibiting either GBs and or MTs others a GB exclusively. These include:

- Traffic patterns that exhibit aperiodic packet arrival rates, for example, CA where messages are triggered based on vehicle dynamics or cooperative perception (GB or MT).
- Application of a congestion control technique based on TRC e.g. packet dropping or packet delay, when congestion increases (GB or MT).
- When a grant cannot accommodate the packet size based on the subchannels allocated when the grant was generated (GB only).

- If the time to the next transmission based on the current grant exceeds the latency requirement of the packet (GB only).

This study focuses on the first three scenarios and will describe in detail the scenarios under which they might occur and the negative impacts they have on the scheduler.

4.2.2.1 Aperiodic traffic patterns & Congestion Control mechanisms (GB or MT)

The main reason for grant breaking and missed transmissions is due to a scheduled/reserved resource becoming available to a vehicle but there is no packet to transmit on that resource. This can occur for two possible reasons; the first is that the application has not generated a packet and, thus, does not require the transmission. The second is that the channel is congested, a TRC approach is adopted, which requires a packet drop or delay imposed on the packet. Depending on the configuration of the sidelink, either a GB or a MT will occur. The negative impact of this is quantitatively analysed in Section 4.2.3.

This section discusses why an application layer packet may not be available and theorises why this causes an issue for the SB-SPS scheduler. It is important to note that SB-SPS was designed assuming periodic arrival rate of packets, which is why it employs a strict RRI mechanism to identify future reserved resources. This assumption is the primary challenge that the standard faces. A similar assumption was made over a decade ago by the wireless 802.11p standards, e.g. for the CA service, a CAM would be sent every 100 ms. However, it became clear, based on the evaluated research, that these periodic rates led to excess packet transmissions [96, 97]. Consensus was reached that vehicular services should operate in an event-driven manner on the basis of the vehicles dynamics. As such, the vehicular application standards including ETSI CA or DEN and the more recently standardised CP, are all designed to be triggered either by vehicle dynamics, a specific event, or sensed object dynamics. Unfortunately, the 3GPP standard technologies did not design their scheduling mechanism considering that most vehicular applications would exhibit aperiodic behaviour and again assumed a strict rate such as 100 ms. Thus the challenge is that the scheduling grant assumes periodic that does not match the actual application message rate.

This section will consider two different aperiodic traffic patterns. The first is

based on a 3GPP suggested pattern [12], where application traffic is triggered randomly, characterised as a minimum of 50 ms plus an exponential random delay with a mean of 50 ms. The other aperiodic traffic pattern considered and the one most commonly used in this evaluation is the ETSI CA service. CAMs are triggered on the basis of 4 conditions which centre on the vehicle's own kinematics:

- The position of the vehicle changes by > 4 m.
- The heading of the vehicle changes by > 4°.
- The speed of the vehicle changes by > 0.5 m/s.
- The time since the last CAM was triggered > 1s.

Due to these conditions the road network scenario significantly impacts the CAM generation rate.

These two different aperiodic traffic patterns are used as the basis for the aperiodic application models. The subsequent analysis will consider a simple highway scenario [12], with 3 lanes in either direction and different vehicle densities increasing from $\beta = 0.06v/m$ to $\beta = 0.3v/m$. Fig. 4.3 shows the packet arrival rates for these 2 models based on the aforementioned highway.

It can be observed from Fig.4.3a (3GPP) that irrespective of traffic density, the pattern remains the same. Importantly, more than 50% of the traffic exhibits an inter-arrival time of less than or equal to 100 ms, with 95% of application packets arriving within 200 ms. In the case of Fig. 4.3b (ETSI CAMs) which is more variable and based on real world dynamics, this is very different and depends on the density of the scenario. As density increases, the vehicle speed will decrease due to road network congestion; this results in a reduced transmission rate as such as density increases the packet inter-arrival times also increase.

The key interactions between the packet inter-arrival times and the maintenance of the SB-SPS grant mechanisms are described in Fig. 4.4. This figure shows a grant operating as normal with two different varying packet inter-arrival times, both can be characterised as being less than *2n-2* where *n* is the RRI of the grant. The reason why the RRI must be maintained at *2n-2* as opposed to *2n* is due to the fact that the SB-SPS mechanism requires 1 ms for the transmission to be prepared, as such a packet arriving at *2n* is already too late to be transmitted; it also requires 1 ms for the transmission itself,

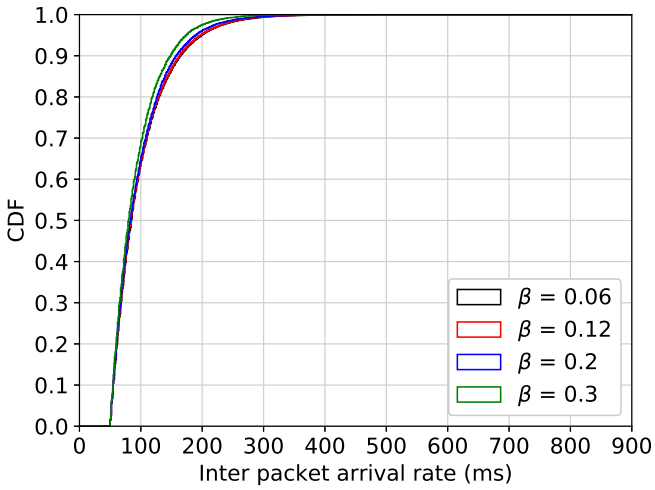

(a) Packet inter-arrival rates for 3GPP.

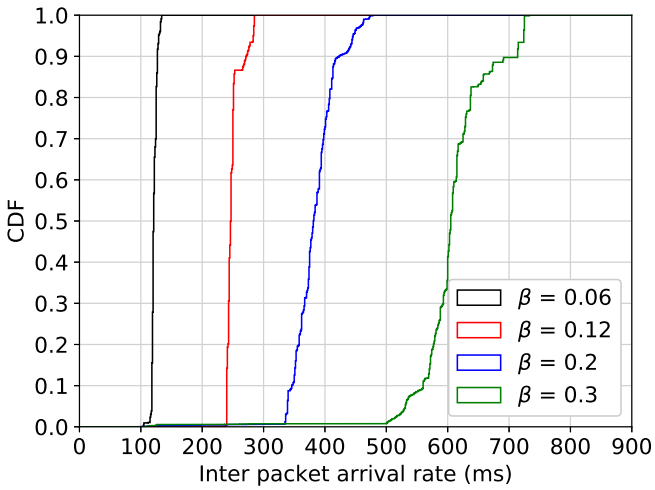

(b) Packet inter-arrival rates for ETSI.

Figure 4.3: CDF of packet inter-arrival rates for aperiodic application models.

therefore 2 ms are required for preparation and transmission.

This can be considered to be typical of the packet inter-arrival times for the 3GPP application model. The first packet inter-arrival time is between the arrival of *Packet A* at *T1* and *Packet X* at *T2n-1*. The second packet inter-arrival time is between *Packet B* at *Tn-1* and *Packet X*. This represents the 3GPP packet inter-arrival rates which are typically 100 ms considering an RRI of 100 therefore $<= n$. There are less frequent rates exceeding n, but less than 10% exceeding the *2n-2* threshold (198 ms). In this scenario, the grant can be effectively maintained, and neighbouring vehicles will see regular

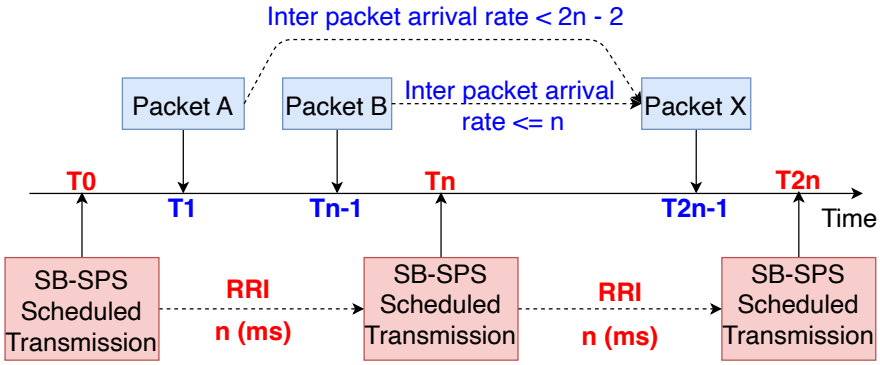

Figure 4.4: Interactions between packet inter-arrival rates and SB-SPS grant maintenance, where packet inter-arrival times are < *2n-2* (3GPP application model).

transmissions at the scheduled 100 ms intervals sign posted by accurate RRIs. Thus the underlying variability does not impact on performance. While it may be the case that an MT may occur due to back-to-back higher inter-arrival times exceeding *n*, in general the grant will be maintained for the majority of transmissions, this will be explored in more detail in Fig. 4.5.

The issue occurs when the inter-arrival rates become higher and more variable, i.e. they exceed *2n-2*, or 198 ms for an RRI of 100 ms. In such cases, the grant will never be maintained correctly, and MTs or GBs will occur constantly. This is representative of the packet inter-arrival rate of ETSI CAMs for densities greater than $\beta = 0.12$ v / m.

Fig. 4.5 goes into further detail on the interactions between inter-arrival rates and the grant maintenance mechanisms. Fig. 4.5a represents the ideal scenario; in this case, the packet inter-arrival rates remain fixed or variable with a maximum time less than *n*. *Packet A* is received at *T1*, another packet is transmitted at *Tn* with the appropriate RRI reserving *T2n*. *Packet X* will be received before *T2n*, specifically *T2n-1* and as such, the scheduling mechanism works effectively as intended. Similarly, to how an application with a forced (periodic) transmission rate will operate.

Fig. 4.5b, represents the highest variability in inter-packet times i,e, this is representative of ETSI the packet inter-arrival rate exceeds *2n-2*. This is the case with all densities above $\beta = 0.12$ v/m, assuming an RRI of 100 ms. In this case *Packet A* arrives at *T0* and is transmitted at *T1*. An SCI is sent with an RRI reserving *Tn*, although the next packet will not be available until *Tn+1* when *Packet X* arrives. This is the worst possible configuration due to the fact that MTs will occur after each transmission due to the complete incompatibility of

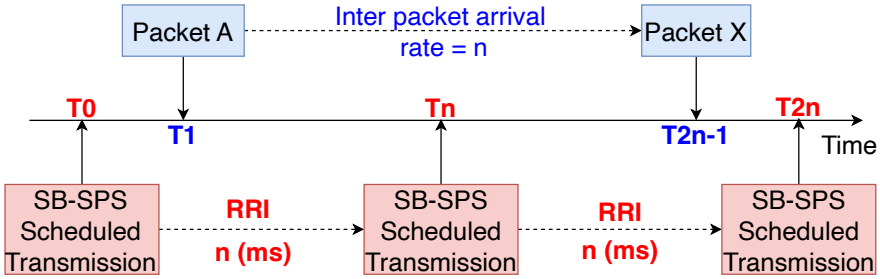

(a) Inter-arrival rate less than n and full grant maintenance.

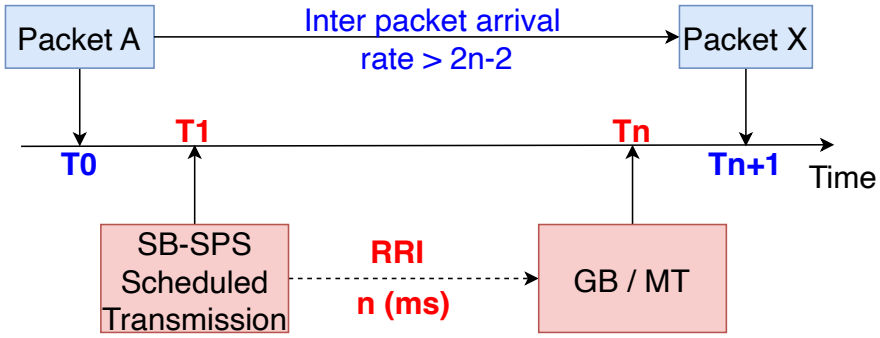

(b) High inter-arrival rate and constant grant interruption (ETSI).

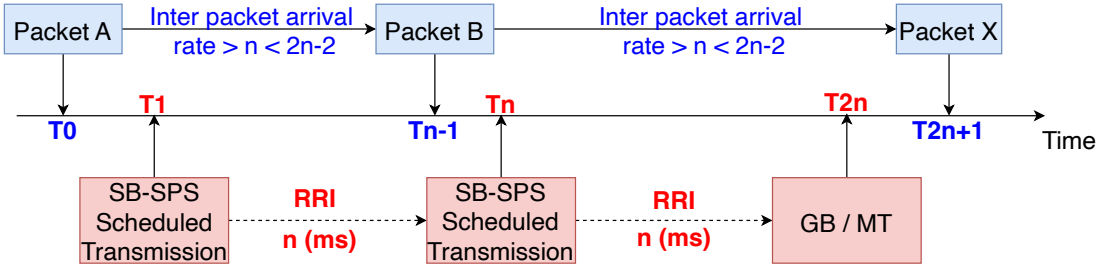

(c) Variable inter-arrival rate and occasional grant interruption (3GPP).

Figure 4.5: Grant maintenance and packet inter-arrival rate interactions.

the RRI with the packet inter-arrival time. Thus depending on the configuration every transmission may result in rescheduling with significant numbers of MTs further disrupting the scheduling mechanism. This scenario is ultimately what the solutions proposed in this thesis will seek to avoid but it is challenging to manage when dealing with diverse applications; it is not simply a matter of tuning the RRI configuration due to the fact that CA/DEN/CP services will have diverse inter-arrival rates depending on the vehicular environment and cannot be easily categorised simply as selecting 200 ms as the typical RRI rather than 100 ms.

Finally, Fig. 4.5c represents what might typically happen with a packet

inter-arrival rate similar to that of the 3GPP application model. In this case, the packet inter-arrival rate is less than *2n-2*. This is typical for the 3GPP scenario shown in Fig.4.3a where the 200 ms packet inter-arrival rate is relatively uncommon (less than 10%). However, the inter-arrival rate can commonly range in the 100 - 200 ms range (40% of packets) which is firmly in the *2n-2* considering *n* as 100 ms. In such scenarios, most packets will typically be accommodated correctly with the RRI indicating the correct resource on which the next transmission will occur. However, the case depicted in Fig. 4.5c can occur and will result in an MT. In this case *Packet A* will be received at *T0*. Transmitted at *T1*. The SCI will reserve *Tn* in the meantime *Packet B* will be received at *Tn-1* and a further transmission will occur at *Tn* reserving *T2n*. Thus far the mechanism is operating as expected and effectively, but due to the underlying random nature of the application traffic, or possibly considering a CA scenario, i.e. a subtle change in vehicle dynamics, such as a slow down that does not exceed the 0.5 m/s rule, the packet inter-arrival rate may increase. This increase does not necessarily exceed *2n-2* it simply must exceed *n* in the case where there is limited additional time, such as in this case where the packet arrives just in time for transmission, that is, *Tn-1* for transmission on *Tn*. In this case, what happens is that at *T2n* no packet is available and an MT or GB occurs depending on the configuration and subsequently at *T2n+1* *Packet X* arrives and will need to be scheduled. This is an unavoidable problem with all applications exhibiting packet inter-arrival rates exceeding *n*, it is just the case that they are not so severe due to the fact that unless all packet inter-arrival rates exceed *2n-2* the scheduling mechanism is able to accommodate most application traffic effectively reducing MTs/GBs.

The conclusion from this analysis is that applications exhibiting highly variable packet inter-arrival rates that exceed the configured RRI will result in some MTs and GBs depending on configuration. Once the packet inter-arrival rate exceeds *2n-2*, the grant mechanism will break down entirely as there is no means for any sequence of packets to be transmitted in the grant without incurring an MT or GB.

So far the discussion has related to the negative impact from applications with variable packet inter-arrival times. A similar phenomena can be found with congestion control mechanisms. One such technique is *DCC Reactive*, which is a table-based mechanism standardised by ETSI. Table B.3 shows how *DCC reactive* works. It employs a delay on packet transmission based on a table look up approach i.e., if current CBR exceeds 30% then a delay of 200 ms is put on

a packet being released by the application layer. This is a mechanism that will be discussed in great detail in Section 6.3, but for now it highlights at a high-level what the mechanism is and what it does. The delay would have a similar impact as a packet dropping approach.

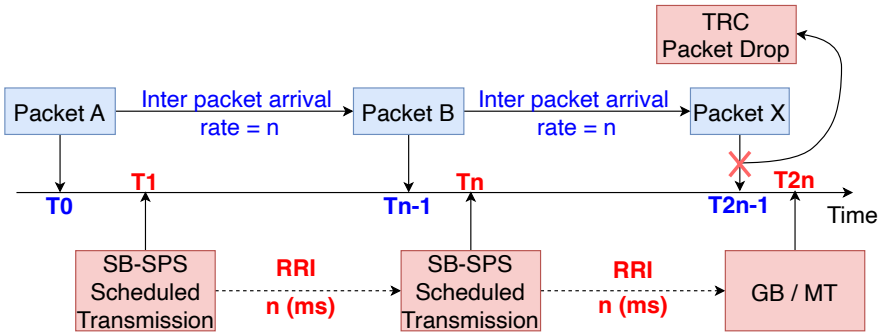

Figure 4.6: Impact of TRC congestion control

Fig. 4.6 shows how this TRC based congestion control approach can negatively impact SB-SPS scheduling and result in an MT or GB, dependant on configuration. In this case it is assumed that the transmission rate is fixed at 100 ms and *n* (RRI) is set to 100 ms also. *Packet A* arrives at *T0* is transmitted at *T1*, reserves *Tn* and *Packet B* arrives in time for transmission at *Tn-1*. However, between *Tn* and *T2n-1*, when *Packet X* should be released to the lower layers, congestion in the network increases. In the case of *DCC Reactive* this may be from 20% to 35% CBR, and as such the mechanism triggers. Now at *T2n-1* when *Packet X* should be released for transmission, *DCC Reactive* instead delays it or a packet drop occurs. As such, at time *T2n* no packet is available for transmission and the MT or GB occurs. It should be noted that packet dropping approaches would have the same impact bar in such a scenario rather than delaying the release of *Packet X* it would instead be dropped.

It is clear that this challenge is similar in nature to those faced when dealing with aperiodic arrival rates, in that a scheduled transmission opportunity becomes available within grant but the packet itself has either not been generated in the aperiodic case or has been delayed or dropped when congestion control is applied.

4.2.2.2 Variable packet sizes (GB Only)

The second key cause of GB that this study will thoroughly investigate is the impact of variable packet sizes. There is one key difference between MTs as a

result of aperiodic traffic patterns or congestion control mechanisms and those that are as a result of variable packet sizes. That is, in the case where a packet cannot be accommodated in the current grant, then an MT will occur and the grant will be broken automatically, as it requires regenerating for the new larger packet size. As an example, the empirical CAM dataset [98] is considered. It will be used in subsequent chapters to discuss MCS adaptation and published in [59]. The distribution of the CAM packet sizes is shown in Fig. 4.7a. Assuming an MCS of 7, common in many academic studies and in the 3GPP simulation standard [12], a bandwidth of 10 MHz and a subchannel configuration of 5 subchannels with 10 RBs per subchannel according to the ETSI standard for C-V2X [52]. The RB and subchannel occupation of each packet size assuming MCS 7 and the ETSI subchannel configuration, is determined according to Tables 7.1.7.2.1-1 and 8.6.1-1 provided in [22].

In Fig. 4.7b it can be observed that subchannel occupation changes when packet size increases, considering the lowest and most common packet size of 200 B requiring two subchannels, 300 and 360 B packets requiring three subchannels, and then 400 B requiring four subchannels. It is worth considering the distribution of these packets as well, such that 38% of packets require two subchannels, 47% require three subchannels and the remainder require 4 subchannels. SB-SPS is configured by default to allocate the minimum number of subchannels to each grant based on the packet size which triggers grant generation; as such if the first packet is 200 B, then a grant of two subchannels will be allocated. If the next packet to arrive is larger, e.g. 300 B packet, SB-SPS cannot accommodate the packet in the current grant. Thus a GB occurs, which also results in an MT, and a grant reselection takes place considering the larger packet size. Once the new grant is configured, it will operate as normal again until the largest packet size of 455 B arrives and results in the same problem. At this point, the grant will have four subchannels and, as such, will not be broken again, but it does represent a waste of resources as the grant will still maintain the four subchannels for transmission of the smaller packets.

A concrete example of the impact of variable packet sizes is illustrated in Fig. 4.8, assuming the same as the above i.e. 5 subchannels, 10 RBs each and an MCS of 7. In this case the inter packet arrival rate is equal to n or the RRI as such that it is not possible to have an MT or GB as a result of no packet available. Instead in this scenario at *T0 Packet A* of size *200 Bytes* arrives and a new scheduling grant is generated, reserving 2 subchannels, occupying 15 RBs

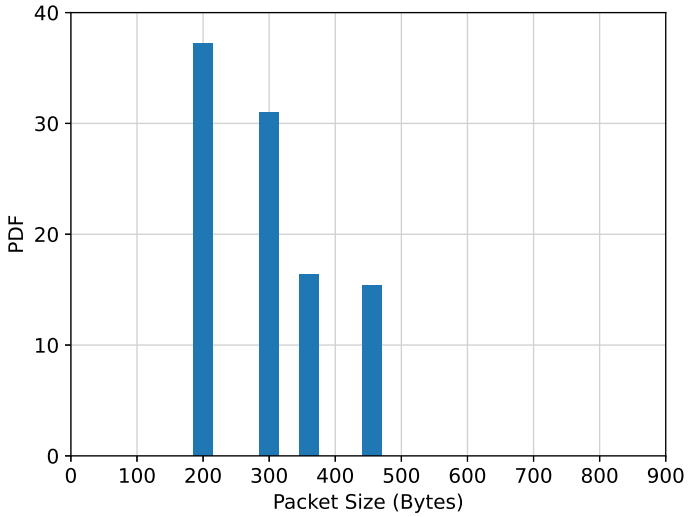

(a) Volkswagen empirical CAM sizes.

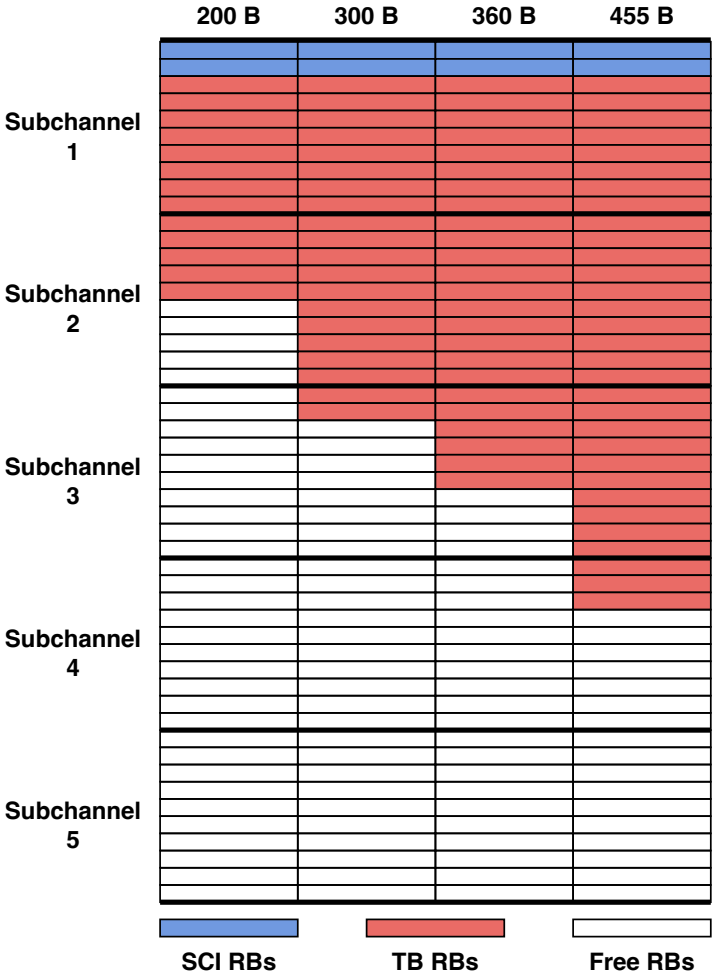

(b) CAM RB occupation.

Figure 4.7: Sizes and RB occupation of the Volkswagen CAMs.

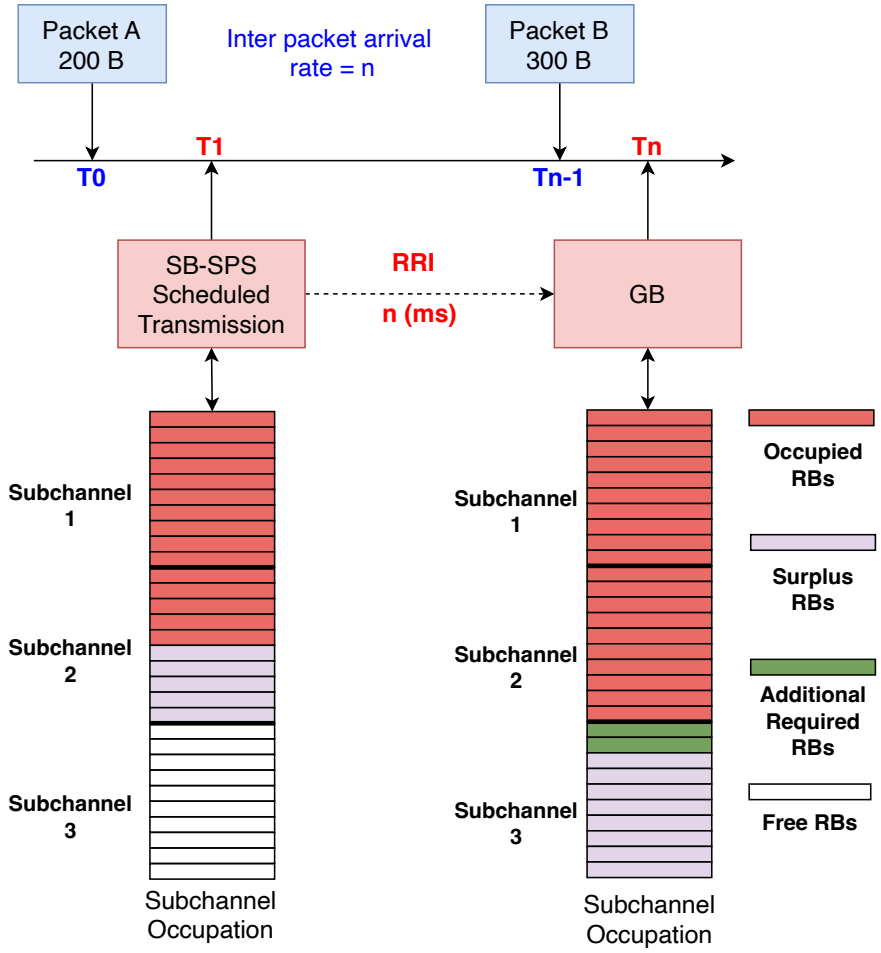

Figure 4.8: Impact of variable packet sizes

with 5 RBs unused. Then at *Tn-1 Packet B* arrives with a larger packet of *300 Bytes*. This is where the challenge occurs; the current scheduling grant allows for 2 subchannels, but 3 are required in this case as 300 Bytes requires 22 RBs with an MCS of 7. As such, the current grant cannot accommodate the new packet and a GB must occur, where SB-SPS will be forced to reschedule for the increased subchannel occupation.

Ultimately, variable packet sizes are a special case for SB-SPS as they result in a GB regardless of configuration, though all three causes represent significant challenges for SB-SPS to manage.

4.2.3 Quantitative Summary of the negative impact of Grant Breaks & Missed Transmissions

This section will evaluate how GBs and MTs impact the performance of SB-SPS scheduling, based on the theorised discussion from Section 4.2.1 on the

negative impact on the SB-SPS sensing history resulting in increased channel collisions due to unnecessary rescheduling events and reduced CSR pools. Also highlighted is the fact that there are many scenarios which can cause both MT and GB to occur, and three have been thoroughly examined; aperiodic applications, TRC based congestion control techniques, and finally variable packet sizes. It is important to remind the reader that for aperiodic and variable packet sizes this study focuses exclusively on the CA service which is one of the backbone "Day 1" services and core to any large-scale V2X deployments. Although the same principles apply for any event-driven application exhibiting aperiodic patterns.

Given that aperiodic patterns and variable packet sizes can have different impacts, a GB or MT as opposed to a GB respectively, they will be discussed independently. First, an analysis of the impact of aperiodic traffic patterns is provided. It should be noted that mechanisms such as *DCC Reactive* which was discussed in Section 4.2.1.1 will have similar impacts to aperiodic patterns.

4.2.3.1 Aperiodic application traffic & congestion control

In this section, the causative effects of declined SB-SPS performance when faced with aperiodic application traffic are explored quantitatively; this analysis is applicable to congestion control and will be highlighted in Chapter 6.

In all subsequent results, as with all results in this study, the simulations have been completed using OpenCV2X, which was introduced in Chapter 3. The common simulation parameters have already been highlighted in section 3.5.1 in Table 3.1, those parameters which have changed from this are shown in Table 4.1. The main significant change is the introduction of an increased number of subchannels and allocated RBs; this corresponds to the standard simulation parameters defined by 3GPP [12]. The results are generated in the same way as outlined in Section A.

The application models employed are as described:

- **Periodic:** Vehicles transmit at a mean rate of 4Hz. This rate is deliberately chosen to allow for comparable analysis with the ETSI model with respect to channel load. Half of all vehicles transmit every 200ms with the remainder transmitting every 300ms, selection of 200ms or 300ms rate is done at random for each simulation.

Table 4.1: Scenario Parameters.

Parameter	Value
Vehicular scenario	
Vehicular density (β)	0.12 veh/m
Number of lanes	3/6 in each direction (6/12 total)
Vehicle Speed	70km/h
Vehicle Mobility	SUMO (step-length = 10ms)
Channel settings	
No. subchannels & Size	3 & 16 RBs
Application layer	
Transmission frequency	Periodic (4 Hz) or ETSI model
MAC & PHY layer	
Transmission power	23 dBm

- **Aperiodic (ETSI):** Literature often assumes a periodic CAM transmission rate. This does not align with the aperiodic transmission of CAMs according to the ETSI specification [15] where CAMs are triggered based on a vehicle's dynamics i.e. deviation in heading (>4°), position (>4m) and speed (>0.5m/s) or at 1s intervals if these conditions are not satisfied.

We now provide a quantitative summary on explore why SB-SPS exhibits declined performance when application traffic follows aperiodic arrival rates. Both C-V2X (Rel. 14) and NR-V2X (Rel. 16) are considered. The distinction between these, as described fully in Section 2.3.4, is that NR-V2X removes the RSSI filtering stage and only uses the RSRP of the most recent transmission when filtering candidate resources.

Ultimately, the results show that both SB-SPS standards exhibit declined performance when handling aperiodic traffic. This is due to inefficient resource use as a consequence of inaccurate knowledge, with the reasons differing depending on whether grant breaking is enabled or not. This leads to a rise in collisions and decreased packet delivery rates. Subsequent results prove this to be the case.

Fig. 4.9 shows that SB-SPS with grant breaking (denoted as GB) performs worst for both NR-V2X and C-V2X. A grant break occurs as no TB exists to send in a scheduled slot, thus the resource goes unused as SB-SPS ends the grant when a missed transmission occurs. When vehicles frequently break their grants, they must schedule new resources. Each time resources are reserved

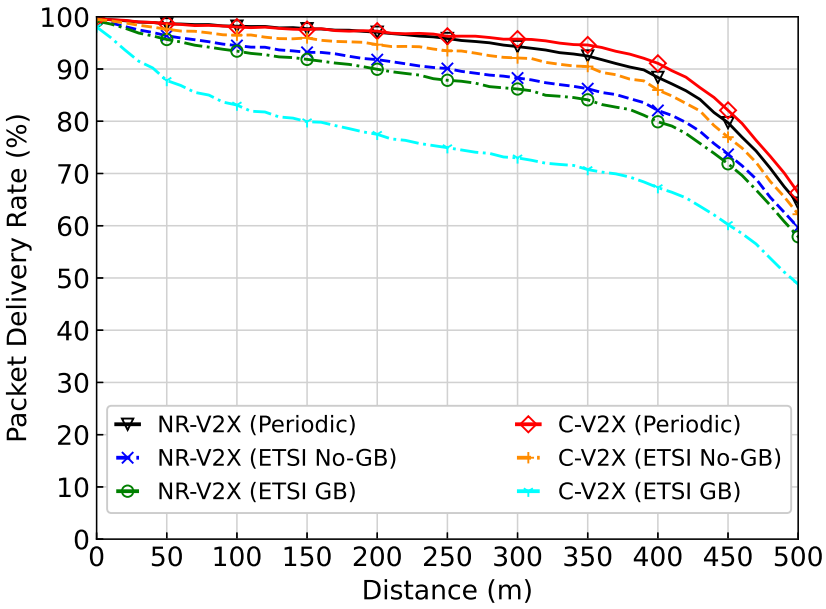

Figure 4.9: SB-SPS performance of NR-V2X and C-V2X for periodic and ETSI application models (β = 0.12 veh/m, with and without grant breaking (GB & No-GB)).

Table 4.2: Absolute number of colliding grants.

Scheduling Mechanism	γ_{TSim}
C-V2X (Periodic)	64
NR-V2X (Periodic)	20
C-V2X (ETSI GB)	4234
NR-V2X (ETSI GB)	1789
C-V2X (ETSI No-GB)	320
NR-V2X (ETSI No-GB)	86

there is a probability that a neighbouring vehicle will choose the same set of resources, particularly if inaccuracies are introduced to the sensing history. Furthermore, frequent grant breaking leads to under-utilisation of reserved resources which causes other vehicles to incorrectly consider increasingly small and similar resource pools. This increases the probability of vehicles choosing the same resources, which can be exacerbated further as vehicle density increases. This has a severe impact on C-V2X Mode 4 due to it's CSR selection mechanism incorporating RSSI filtering. In contrast, this stage is removed for NR-V2X which results in a larger CSR pool. This is shown by the γTSim metric in Table 4.2 which represents the number of instances where grants collide i.e. where neighbouring vehicles select the same subchannel in the same subframe for transmissions. It can be seen that C-V2X incurs a high number of colliding

grants. While NR-V2X incurs less, for the reasons just outlined, it is still evident that grant breaking is incompatible with aperiodic traffic patterns.

Fig. 4.9 also shows NR-V2X and C-V2X SB-SPS performance when grant breaking is disabled (denoted as No-GB). The grant is maintained irrespective of a missed transmission. It can be seen that while this greatly improves performance for C-V2X when compared to grant breaking, it still comes at a cost. Resources are reserved that may not be used. Furthermore these reserved but unused resources introduce inconsistencies in the CSR selection process where a vehicle does not send an SCI due to a missed transmission and neighbouring vehicles mistakenly believe the resource to be free for use. This can lead to collisions on future transmissions. What is particularly noteworthy is that NR-V2X performance when grant breaking is disabled is worse than the performance of C-V2X. This indicates that when the grant is maintained, RSSI filtering can be a positive feature by discouraging the selection of resources that exhibit high levels of interference. Furthermore, RSRP averaging can result in more accurate sensing history.

Irrespective of whether grant breaking is enabled or not, it has been shown that the performance of SB-SPS for NR-V2X and C-V2X becomes compromised due to inefficient resource scheduling. This leads to increased packet collisions as shown in Fig. 4.10.

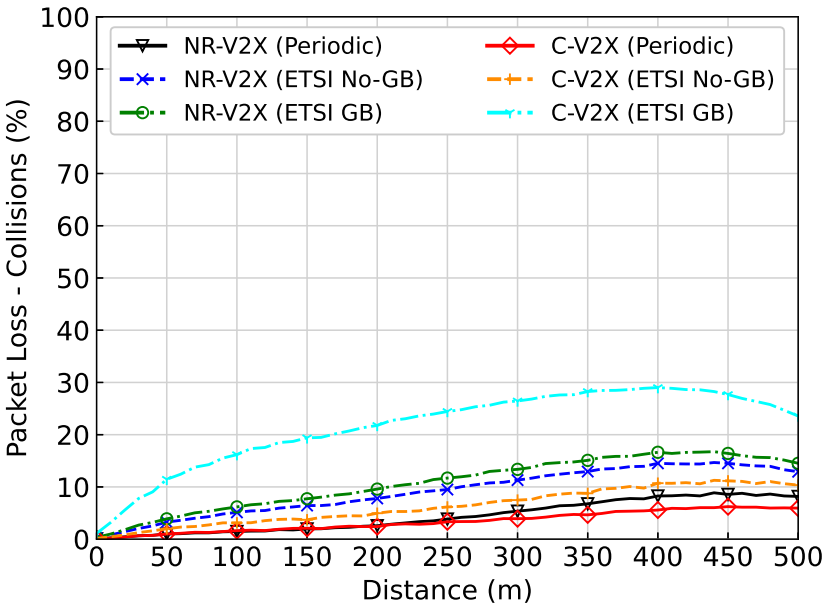

Figure 4.10: Packet collisions for periodic and ETSI application models, (β = 0.12 veh/m)

4.2.3.2 Variable packet sizes

As discussed, in the case of variable packet sizes, if the grant cannot accommodate a new larger packet then a GB must occur. For this reason, the impact of the variable size is not impacted by the configuration of the *sl-reselectAfter* parameter.

The following results are generated with the same parameters as described in Table 4.1, with the following exceptions. To show the impact of variable sizes another 3GPP standard approach is considered, where packet sizes have a sequence of 4 190 B packets followed by a 300 B packet. This is a simpler application model and useful for providing an indicative result of the negative impact that variable pacekt sizes has on SB-SPS can have. In Fig. 4.11, the standard ETSI configuration of 5 subchannels with 10 RBs each and 2 MCS configurations of 7 and 8, is considered. For MCS 7, the 190 B packet will require 15 RBs to accommodate or 2 subchannels, while the 300 B packet will require 22 RBs or 3 subchannels, while MCS 8 will only require 2 subchannels for each packet size (190 B = 13 RBs, 300 B = 20 RBs).

From Fig. 4.11a it is clear that for the MCS 7 configuration the performance is significantly degraded, although the level of congestion shown in Fig. 4.11b is actually higher for MCS 8. The most significant difference between the two configurations is in terms of the number of GBs as a result of packet size, meaning the larger 300 B packet arrived with a grant configured for 190 B. For MCS 7 this occurs a mean number of 2872 times throughout the simulation, whereas MCS 8 entirely eliminates these due to the fact that it only requires 2 subchannels for all packet sizes.

This finding is shown simply to be indicative of the impact of GBs in this context, with further results shown in Chapter 5 to explore this in much greater detail. The primary observation here is that it can significantly degrade SB-SPS scheduling performance. Additionally, it should be noted that MCS and subchannel configuration are highly important variables to consider when examining C-V2X and NR-V2X performance.

4.2.4 Problem analysis conclusion

Section 4.2.1 provided a theoretical analysis of the primary challenges facing the SB-SPS scheduling mechanism, namely GBs and MTS. Section 4.2.2 then continued on to describe the specific scenarios under which these can occur,

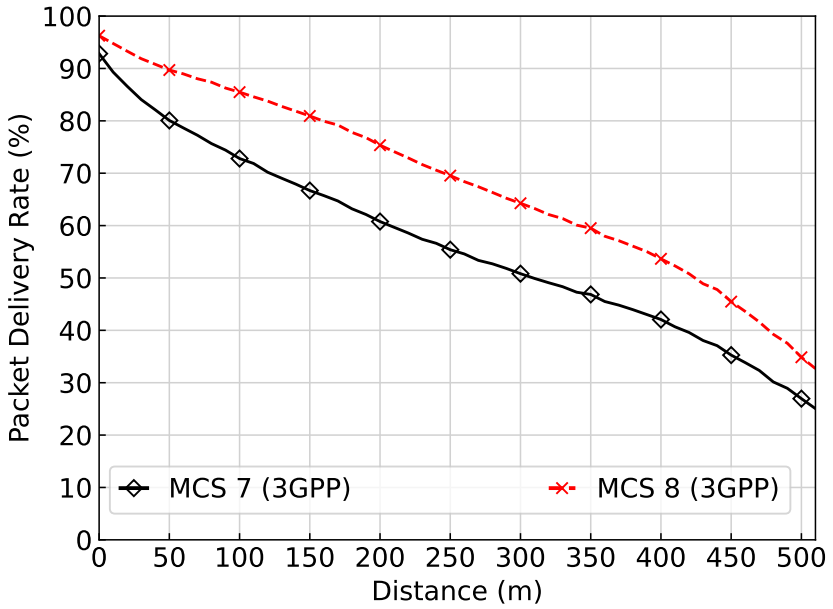

(a) Grant breaking comparison (PDR).

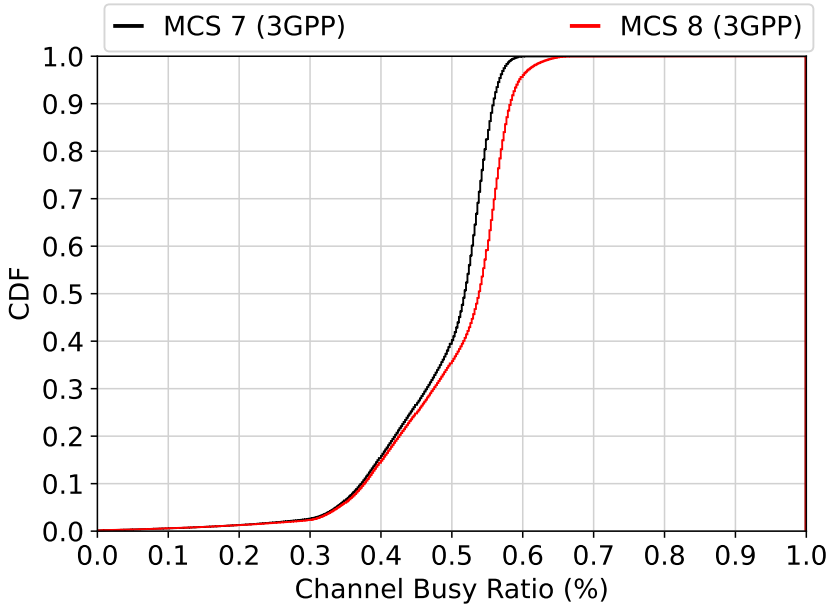

(b) Grant breaking comparison (CBR).

Figure 4.11: Impact of grant breaking as a result of variable packet sizes.

namely aperiodic traffic patterns, TRC based congestion control techniques, and finally as a result of variable packet sizes. The final Section 4.2.3 then provided a quantitative analysis of the impact these scenarios can have which will be supported in subsequent results chapters.

The causes of these MTs and GBs are as a result of the current generation of "Day 1" services such as CA which exhibit aperiodic arrival rates, as well as variable packet sizes, this being as a result of multiple different configurable fields in the CAM itself. Additionally, this is also applicable to any congestion control mechanism that focuses on a TRC approach through packet delay or packet dropping. This is due to both resulting in a packet being unavailable at the time of a transmission being required and there being no application packet available. Dealing with variable packet sizes has been shown to be a special case where, regardless of configuration, any increase that cannot be accommodated in the current grant will result in a GB not just an MT.

Ultimately, through the example results shown, the significant impact that both GB and MT may have on performance has been highlighted; this both corroborates other academic studies which have highlighted these challenges [46], but will be further proven in the deeper analysis which can be found in Chapters 5 and 6. Considering these challenges, the chapter will now focus on the means by which one can reduce or overcome these and ensure that C-V2X and NR-V2X can operate effectively with services exhibiting aperiodic arrival rates, as well as in dealing with high congestion scenarios.

4.3 Proposed Solutions

Now that the challenges are clearly defined and quantitatively analysed, the goal is to propose solutions to overcome these. This section will focus on the theoretical basis of the proposed solutions, and all aspects of their practical implementation. Subsequent chapters 5 and 6 will provide a detailed analysis of the performance of these proposed solutions, as well as a comprehensive review of the impact of both MT and GB. As already outlined, the main problems faced by both SB-SPS standards are GBs and MTs which degrade scheduling performance; our proposed solutions will seek to overcome these challenges in a standards-compliant fashion.

Due to the causes of both GB and MTs being similar in the case of aperiodic application traffic and TRC congestion control mechanisms, and the

overarching solution as applied in Sections 4.3.2 and 4.3.3 will be outlined first in Section 4.3.1. The other main challenge is dealing with variable packet sizes, which cause GBs, and will be discussed in Section 4.3.4.

4.3.1 Overview: RRI adaptation for accurate grant maintenance.

The issue when dealing with both aperiodic application packet arrival rates and any TRC mechanism is the fact that each SB-SPS transmission requires an SCI with an RRI field to reserve the subsequent resource. The challenge then is the fact that in both cases this transmission rate is not a strict period as it is being effected either by the application generation rules or by a congestion control mechanism. As such the RRI need to dynamically change after each transmission and this is what causes the MT or GB in this case.

Our proposed solution is simple in concept; if the RRI can be dynamically changed to correspond to either the congestion control mechanism or the dynamics of event triggered application, this will overcome the issue of MT and GB. With such an approach, the sensing history will be accurately maintained and enable SB-SPS to operate effectively as normal, where the underlying traffic appears periodic in nature, as a result of the accurate RRI but is, in reality, dynamic.

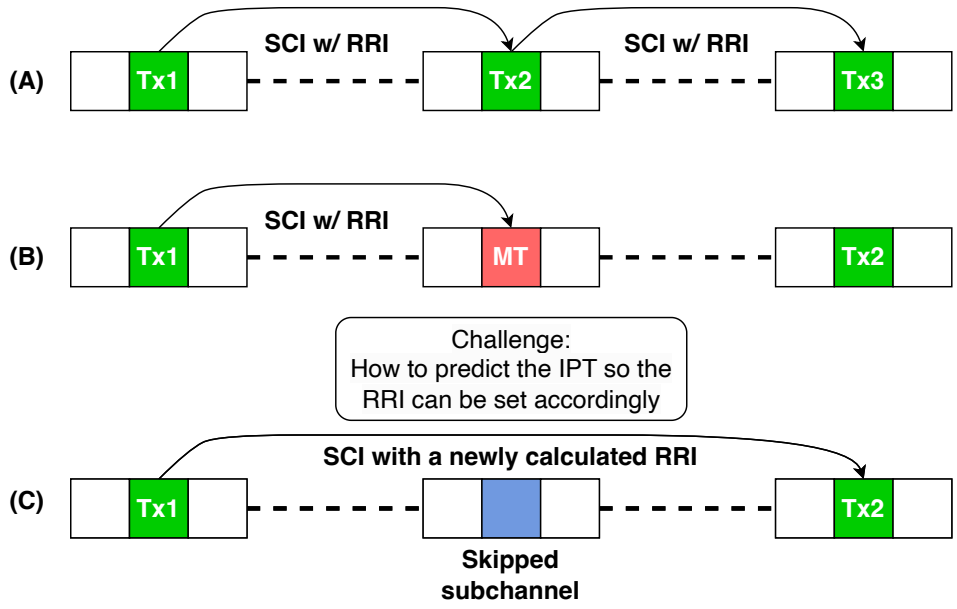

Figure 4.12: Overview of RRI Adaptation.

This is shown in Fig. 4.12, timeline *(A)* in this represents the ideal scenario where the application traffic appears in a periodic sequence and maintaining a single RRI throughout. Timeline *(B)* represents the scenario which the mechanism is designed to overcome. Where at *Tx1* there is a transmission and an SCI is sent with an RRI reserving the intermediate resource. This intermediate red resource becomes a *MT* due to the lack of packet availability for transmission and considering that a no GB configuration *Tx2* is sent without a reserving RRI which may impact scheduling performance. Timeline *(C)* illustrates the proposed solution where at *Tx1* some mechanism exists to predict or calculate what the inter-arrival time of the next packet will be and sets the RRI accordingly. As such the intermediate blue resource is correctly skipped and *Tx2* is accurately reserved. There are two benefits to this. First and most importantly is that the SB-SPS grant mechanism is correctly managed, as such the neighbouring vehicles will have an accurate sensing history and will schedule more effectively. The second benefit is that these intermediate resources, rather than being an MT, become free resources that other vehicles may reserve.

The key challenge with applying this technique is not in the simple act of adapting the RRI, which is easily configurable. Rather it is in the accurate selection of the correct RRI to predict the packet inter-arrival time. This is a different challenge for each of the causes. It requires integrating this adaptation into the mechanism itself, i.e. the congestion control technique takes action in tandem with the scheduling mechanism. This aspect does not overcome additional challenges faced when implementing any congestion control technique i.e. to effectively manage congestion in both a fair and stable manner. It only enables the mechanism to operate correctly with SB-SPS. When handling aperiodic application patterns, it is a different challenge, with a more complex means of predicting the RRI. As such different machine learning-based approaches for predicting CAMs were investigated, which is ultimately a trajectory prediction problem. This required integrating a machine learning approach with our RRI adaptive mechanism, whereby the prediction is used to determine the RRI itself and thus maintain the SB-SPS scheduling mechanism effectively.

A discussion of the specifics for each solution is provided next, with a focus on congestion control in Section 4.3.2 and aperiodic arrival rates in Section 4.3.3.

4.3.2 Solution for congestion control: Resource Reservation Interval Decentralised Congestion Control (RRI DCC)

The primary goal of the proposed congestion control approaches is to integrate with the underlying scheduling mechanism, explicitly taking into account the SB-SPS RRI. The proposed solution is Resource Reservation Interval Decentralised Congestion Control (RRI DCC) **[C4]** as a way of integrating congestion control approaches with the default scheduling mechanism SB-SPS. This operates by adapting the RRI such that intermediate resources can be skipped, thus reducing congestion in the channel while estimating the RRI correctly and thus maintaining the SB-SPS grant mechanism. It is key that any congestion control technique does this, as observed from Fig. 4.9 in Section 4.2.3.1 not considering this seriously degrades performance resulting in the mechanism becoming ineffective.

To address this, three *RRI DCC* mechanisms are proposed. These mechanisms adjust the time between transmissions based on the current CBR measurement, while considering the RRI of the SB-SPS grant. Three variants are proposed to align with existing vehicular congestion control standards as set out in Section 2.3.7 in Chapter 2, specifically to align with ETSI *DCC Reactive*, ETSI *DCC Adaptive* and packet dropping with 3GPP CR limit tables.

They operate as follows. Each mechanism will determine, based on the current CBR, the new RRI that is required based on either an algorithm calculation or a table lookup. At the time of the next transmission from when the new RRI is determined, the MAC layer will update the RRI and send it to the PHY layer for transmission. Any intermediate packets that arrive before the next available grant will be delayed and dropped should a new packet arrive. The transmission will then only occur when the RRI of the first transmission has elapsed. This ensures that the SB-SPS RRI mechanism operates as it should; however, it does cause divergence from the ITS-G5 standards for congestion control. In ITS-G5 *DCC Adaptive* there is scope to end a delay early if congestion reduces. This is not possible in SB-SPS as ending the delay early would result in breaking the RRI mechanism and likely require the grant to be relinquished and a new one being configured for the earlier transmission. This would contribute to rather than resolve the MT issue and, as such, once the RRI is sent should congestion reduce, the RRI can only be changed back once the next transmission occurs.

There are three key contributions in this aspect of the work, the first being an effective TRC mechanism which is standard compatible in 3GPP using a CBR to CR-Limit table being RRI_{CR_Limit}, as well as ETSI standard mechanisms in $RRI_{Adaptive}$ and $RRI_{Reactive}$. All three approaches overcome the issue of introducing MTs as a result of packet dropping or delay not being accounted for in the scheduling approach.

The second key contribution relates to the implemented $RRI_{Adaptive}$ approach, the primary change away from a table-based approach with both $RRI_{Reactive}$ and RRI_{CR_Limit} offers much improved flexibility. Many academic studies have shown how table-based approaches lead to CBR instability and they pose a significant challenge for dealing with highly dynamic networks requiring significant tuning to be effective. Avoiding this altogether by adopting an algorithmic approach ensures much greater flexibility and avoids the overbearing tuning challenges.

The final contribution is that, our scheme is the first to be studied in NR-V2X and C-V2X with a significant focus placed on the fairness and stability of the proposed scheme. Fairness and stability are key measures with any congestion control technique, it is vital especially in a vehicular scenario to avoid vehicles being starved of resources where others are over using the limited channel resources. Such a scenario may pose a significant safety risk considering a CA application, as such managing congestion to ensure that all vehicles have fair access to the channel and that the limited resources available are used most effectively is a major contribution. It will be shown in subsequent results that $RRI_{Adaptive}$ both manages congestion effectively and in a fair manner while maintaining a high level of CBR stability as such, meeting the three challenges for any congestion control technique while eliminating the additional challenges a table-based approach imposes.

4.3.2.1 $RRI_{Reactive}$ ETSI table based mechanism based on *DCC Reactive*

The first proposed congestion mechanism is $RRI_{Reactive}$, which is based on the ETSI table based mechanism *DCC Reactive*. This uses a lookup table to determine the packet delay, that is, the T_{off} parameter, based on the measured CBR. This results in a packet delay that may not be compatible with the SB-SPS RRI i.e. reserved transmission slots. However, $RRI_{Reactive}$ ensures that the delay offset is linked to a multiple of the default RRI, preventing grant breaks and unused resources. Furthermore, an SCI is transmitted indicating the new RRI

to neighbouring vehicles. This controls the rate of packet transmission without missing a scheduled slot. The translated table is shown in Table 4.3.

Table 4.3: $RRI_{Reactive}$ Lookup table.

CBR	State	Packet Rate	RRI
CBR $<$ 0.3	Relaxed	10Hz	100ms
0.3 $\leq$ CBR $\leq$ 0.40	Active 1	5Hz	200ms
0.40 $\leq$ CBR $\leq$ 0.50	Active 2	2.5Hz	400ms
0.50 $\leq$ CBR $\leq$ 0.60	Active 3	2Hz	500ms
CBR $>$ 0.60	Restrictive	1Hz	1000ms

4.3.2.2 RRI_{CR_limit} 3GPP table based mechanism

The second proposed mechanism is RRI_{CR_limit} which is based on the 3GPP table based mechanisms. This approach is based on packet dropping to reduce an individual vehicles' CR-limit. Similarly to the ETSI and 3GPP V2X approaches, it is also based on a CBR to CR limit table. Given a particular CBR, representing the overall measured congestion of the channel, if it is determined that the vehicle's CR exceeds the limit, the RRI will be increased such that the CR is brought below the limit. An SCI is then transmitted indicating the new RRI. When the CBR returns to a lower range and the CR can be increased, a new RRI will be chosen that maintains the new limit. 3GPP does not define a specific table to employ and as such throughout the subsequent results sections different tables will be defined for each scenario. This is one of the primary drawbacks of such table-based approaches in that it cannot be easily ported to a new scenario. An example table based on an ETSI standard is shown in Table B.2. This does not neatly translate to a specific RRI, which will require calculation on the part of the MAC layer. The MAC layer whenever a transmission occurs will refer to the table based on current CBR and its own current CR and then determine the RRI required to ensure that its CR remains below the limit.

The DCC averaging mechanism is employed for $RRI_{Reactive}$ and RRI_{CR_limit} where RRI transitions occur only after 1 second of CBR exceeding a threshold and RRIs decrease after 5 seconds of lower measured CBR. This results in a more stable level of CBR for all vehicles.

4.3.2.3 $RRI_{Adaptive}$ ETSI algorithmic mechanism based on *DCC Adaptive*

The final proposed mechanism is $RRI_{Adaptive}$ which is based on the ETSI algorithmic mechanism *DCC Adaptive*. This approach uses an algorithm to determine the delay of the packet, that is, T_{off} [3] which is based on LIMERIC [25]. In the standard, δ is the proportion of the channel that the vehicle is allowed to use, based on previous usage and current CBR, it's calculation is shown in Alg. B.1. These can be translated for NR-V2X due to the equivalent nature of CBR calculation for ITS-G5 and NR-V2X, i.e., calculating CBR based on subchannel usage vs. time sensed as busy. For channel usage instead of calculating a proportion of the transmission time, as is the case in ITS-G5, the calculated CR is substituted. However, since the use of the 3GPP calculated CR is a combination of future and historical usage, this causes CBR oscillations, hindering convergence on a target CBR. Therefore, only historical CR is used. The most significant change is to translate Equation B.1 in [3], to calculate the final T_{off} parameter. Equation 4.1 shows how δ is used to calculate the delay T_{off}. The S_{used} is calculated as the number of subchannels used in the previous transmission and S_{total} is the total number of subchannels per subframe for the previous second. As it must be ensured that the generated T_{off} time corresponds to a multiple of the default RRI, equation 4.2 provides this translation. Importantly, there is no means to update the RRI after its transmission in the SCI and, as such, no means of using an adapted version of equation B.2 in the standard. The LIMERIC algorithm uses a moving average of the CBR; it does not require the DCC averaging mechanism and, as such, it is not employed.

$$T_{\text{off}} = ((S_{\text{used}}/S_{\text{total}} * 1000)/\delta) * 1000 \tag{4.1}$$

$$rri = round(max(100, min(T_{\text{off}}, 1000)) \tag{4.2}$$

Fig. 4.13 illustrates the concept underpinning all three *RRI DCC* approaches (shown according to timeline *B*) with the default NR-V2X SB-SPS operation shown according to timeline *A*. In *A*, after transmission *T1*, congestion occurs. This results in a delay due to a TRC congestion control mechanism or a packet drop. The consequence is that a packet is not transmitted in the next scheduled resource reservation slot, i.e. a missed transmission (*MT*). As a result, an SCI is not transmitted so neighbouring vehicles believe the resource(s) to be free.

However, as the SB-SPS grant is maintained, and may be utilised in the future, e.g. transmission *T3*, a collision on that resource can occur. In contrast, the proposed *RRI DCC* mechanisms shown in *B*, explicitly consider the RRI in the grant when adding a delay in the case of $RRI_{Reactive}$ and $RRI_{Adaptive}$ or packet dropping in the case of RRI_{CR_limit}. When transmission *T2* occurs, the new RRI is determined and an SCI is broadcast to all neighbouring vehicles. This ensures that neighbouring vehicles are informed of when the vehicle next intends to transmit and can utilise the free resources in the meantime.

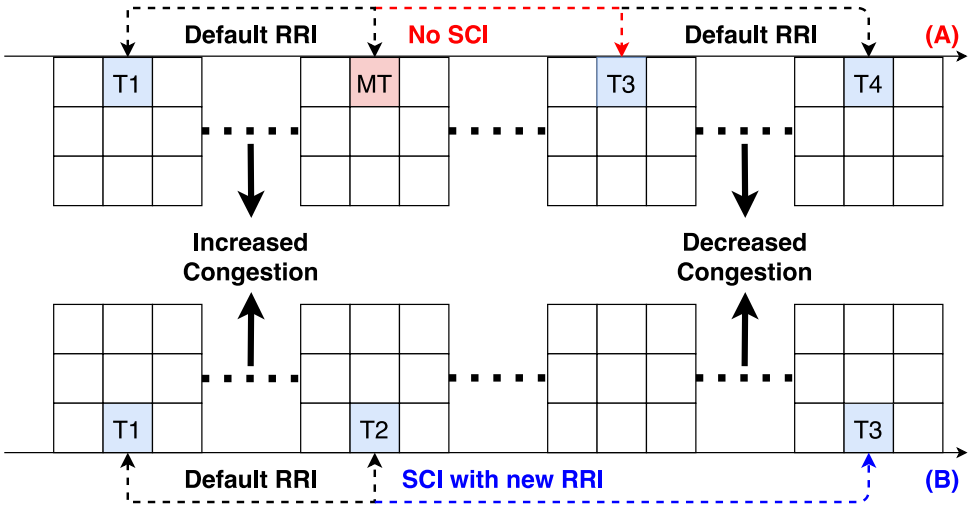

Figure 4.13: RRI DCC approaches versus default NR-V2X SB-SPS approach.

4.3.3 Solution for Aperiodic Traffic Patterns: Enhanced scheduling via CAM IPT prediction

The primary service which forms the core of "Day 1" deployments is the CA service; it is the main service concentrated on collision avoidance and increased vehicle safety. A key trait of this service is that it is triggered based on vehicle dynamics, this choice was determined when initial investigations of ITS-G5 found that with limited spectrum, e.g. 10 MHz, a periodic 100ms packet transmission rate would result in high levels of congestion in the radio network. CA was devised to be triggered on dynamics both to reduce congestion and maintain high levels of position accuracy by informing vehicles of any significant dynamics changes, thus the rules described earlier and will be reiterated shortly in describing our solution.

Another key "Day 1" service is DEN, which is used for all event-based notifications, e.g. emergency braking events or an emergency vehicle approaching. It can also be used for other services such as weather warnings

or slippery surface warnings, generally in a V2I context. It is generally event-based with such events being difficult to predict due to their infrequent occurence. As such it will have a less significant impact on network load but require high reliability to ensure safety. Additionally, CP another service which focuses on ensuring safety for VRUs and vehicles through sensor information sharing, again though it may exhibit aperiodic rates which are difficult to predict.

As already described, these applications, given their aperiodic nature will negatively impact SB-SPS scheduling performance and result in degraded network performance. This is thoroughly investigated in Chapter 5 which highlights in-depth the challenges with MT and GB and additionally focusses on the performance of the new mechanisms employed in NR-V2X to manage aperidoic traffic e.g. Dynamic Grant. Dynamic Grant uses a grantless mechanism where SB-SPS schedules a resource for transmission, and it is used once or a limited number of times, depending on the configured number of retransmissions required. The results section will bear out that this mechanism is not effective, degrading to a pure random access approach, when dealing exclusively with aperiodic traffic, but it can be effective when there is a background level of periodic traffic enabling SB-SPS scheduling performance to improve and manage both aperiodic and periodic traffic more effectively.

This finding led to the conclusion that if the core traffic can be periodic or mimic periodic performance, thus enabling SB-SPS to operate as intended, then dynamic grant can be employed for purely aperiodic traffic patterns, which are less frequent. This is applicable to our two main services; CA forms the majority of traffic, and due to its nature, i.e. being based on vehicle dynamics, is predictable in some form (ultimately a trajectory prediction problem). The other application DEN is much less predictable, and as such is a better candidate for dynamic grant-based scheduling. This motivated the proposal to implement a mechanism for predicting CAM inter-packet time (IPT) and dynamically adapting the RRI to enable it to operate as though it is periodic in nature when, in fact, it is not. To this end a machine learning approach, is implemented which enables the prediction of the next packet IPT through evaluating the trajectory of the vehicle for the next second i.e. the maximum time before a CAM is generated if no other rules are triggered. The RRI is then set to correspond with that predicted IPT as has been shown in Fig. 4.12.

4.3.3.1 Proposed Gated Recurrent Unit Model (GRU) Model for Prediction of CA service

Having described how the scheduler will use the CAM IPT prediction **[C3]**, this section now describes how the prediction occurs. To reiterate the proposal will use trajectory prediction to effectively adapt the RRI of the SB-SPS scheduler to eliminate MT and GB. Several studies have looked at the use of LSTMs to model mobility and trajectory prediction [68, 69, 70]. As the CA application trigger conditions are mobility-based this appeared a promising candidate for an ML based approach. Initial implementations looked at the performance of LSTM based models as well as GRUs:

Both LSTMs and GRUs are recurrent neural networks that operate by maintaining a form of long-term memory throughout the network. This works by persisting information across each step of the network. Through experimentation, it was determined that an GRU-based approach was preferable to that of LSTM in terms of accuracy and speed. The architecture of the GRU model is shown in Fig. 4.14a. The model operates by persisting a form of long-term memory across a series of data. The architecture of a GRU neuron can be broken down into several components, which are highlighted in Fig. 4.14a. The first part being a reset gate that decides whether the previous neuron's state information *ht-1* in combination with the current input data *xt* will be used in the prediction. This is followed by an update gate, which determines the information to be stored in the hidden state, which represents the additional historical data stored by the model and used in subsequent prediction. Then finally there is the hidden state candidate which determines the *ht* or the ultimate hidden state of the neuron. As illustrated in Fig. 4.14b, each neuron feeds into subsequent neurons i.e. taking a specific segment of the data sequence and their predictions feeding forward to later elements in the sequence. An example of this is where the data *xt-1* and prediction *ht-1* influences the prediction of the subsequent data *xt*. This enables the form of long-term memory which allows these RNN models to deal effectively with sequence data.

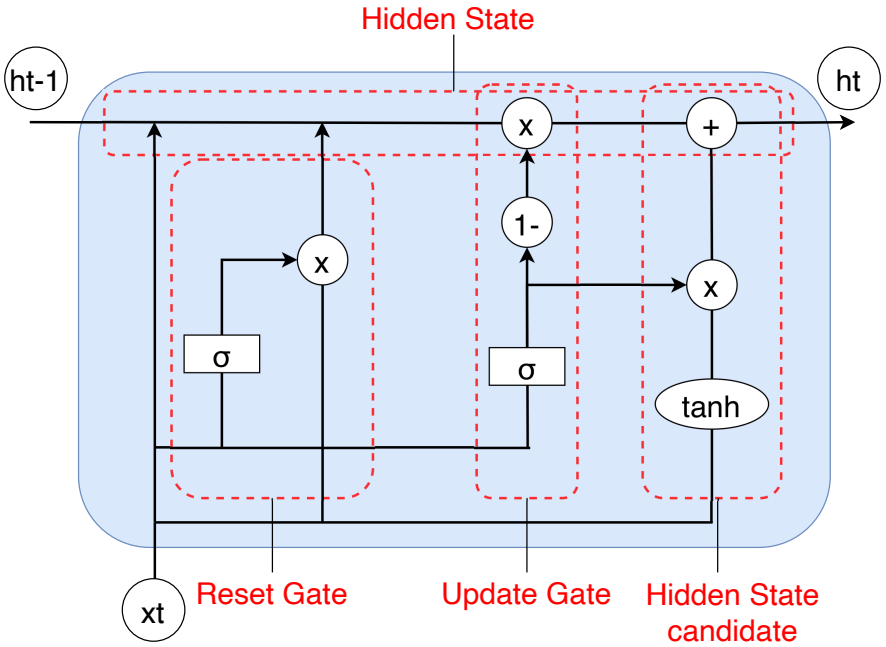

(a) GRU Architecture.

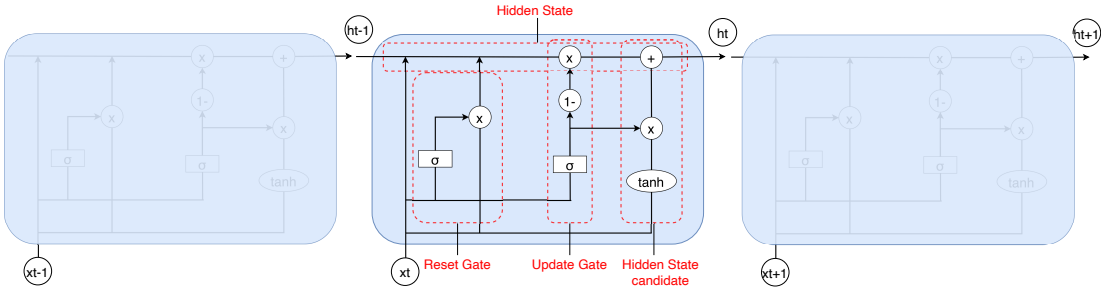

(b) GRU rolled out.

Figure 4.14: GRU (RNN) architecture diagram.

To apply the GRU for CA prediction, several approaches were initially employed; Initial iterations of the model predicted the exact RRI based on the 4 ETSI triggering conditions:

1. Change in position of > 4 m.
2. Change in heading $> 4°$.
3. Change in speed of > 0.5 m/s.
4. The time since the last transmission exceeds 1 s.

This proved to be insufficiently accurate, mostly due to it being difficult to generate a comprehensive dataset with each rule being sufficiently represented, focusing purely on trajectories enables a larger dataset to be developed and generalises the model to other challenges. The proposed

solution implements 3 models which are used to predict the heading (sine and cosine) and speed of the vehicle for a 1 second interval but in 100 ms steps (10 predictions made). The predictions are made by passing the speed, sine and cosine, sampled every 100ms, for the last minute. The predictions are then compared to the ETSI trigger conditions and used to determine the predicted RRI. It was necessary to consider sine and cosine to eliminate inaccurate predictions, e.g., where 355 is nearer to 5 than 30. Considering the cosine and sine of the heading allows for the reconstruction of the heading but avoids issues in the training process. The packet will then be passed to the lower layers with the predicted RRI so that the SCI can be updated with the new RRI.

4.3.3.2 Implementation of GRU CAM IPT Prediction Model

As the accuracy of any proposed model will highly depend on having realistic training data, several different data sets were used as shown in Table 4.4. These scenarios were used to model vehicular dynamics incorporating CA triggering conditions to train the model effectively. The data were split into sequences of 1 minute and 1 second length with speed and heading data sampled every 100 ms. The specific inputs into the model are the ego vehicles current speed in m/s and the cosine and sine of the current heading in degrees. This means there are 600 samples of speed, cosine and sine as input into the model. The output is 10 speed, cosine and sine values, one for each 100ms in the next second. The sine and cosine are combined to determine the predicted heading. The first minute represents the data to use for prediction, with the final 1 second representing the ultimate prediction. Each sequence was then randomly assigned to the training, validation and test sets with a split of 80%:10%:10%.

The details of the implementation of the GRU prediction model is shown in Table 4.5. The structure and parameters used for each of the layers, such as the choice between GRU and LSTM, the number of Recurrent Layers and dense layers, learning rate, neurons per layer, and activation functions used, were determined using a Bayesian search grid on all parameters from which the best performing model was determined and trained on the full data set. These results show good performance for the training of the model with low error across training, validation, and test data showing that the model generalises well to unseen data as well as avoiding overfitting.

Table 4.4: Vehicular Scenarios employed for training data generation.

Scenario	**Type**	**Description**
Digital Twin Geneva Motorway (DT-GM)	Highway	Digital twin of Geneva highways allowing direct link to real-time data, as well as default scenarios representing a typical day.
Dublin (ITSC2020 CAV impact)	Highway/City	Scenario is separated into segments of the M50 motorway around Dublin and internal segments of Dublin city itself.
Ingolstadt (InTas)	City	City scenario based on typical traffic on a day in Ingolstadt.
Monaco SUMO Traffic (MoST)	Highway/City	Detailed scenario with traffic throughout the day in Monaco, including surrounding roads.
Turin SUMO Traffic (TuST)	Highway/City	Similar to MoST scenario is a 24-hour snapshot of Turin including highways surrounding the city.
Cork South Ring Road	Highway	Custom scenario implemented with a 21km scenario with high speeds.

Table 4.5: Structure of GRU Models and performance across data sets (MSE)

Speed Model			**Sine Model**			**Cosine Model**		
Layer	*Func.*	*Neuron*	*Layer*	*Func.*	*Neuron*	*Layer*	*Func.*	*Neuron*
GRU	Tanh	128	GRU	Tanh	128	GRU	Tanh	128
GRU	Tanh	128	GRU	Tanh	128	GRU	Tanh	128
Dense	Lin.	10	GRU	Tanh	128	GRU	Tanh	128
			Dense	Lin.	10	Dense	Relu	32
						Dense	Relu	32
						Dense	Lin.	10
Train	Val	Test	Train	Val	Test	Train	Val	Test
.098	.01	.001	.0147	.015	.0094	.0018	.0015	.008

Several additional models were explored and iterated upon to develop the current model which is used in the subsequent analysis. Key iterations are described in Appendix C.

4.3.3.3 Model practicality

The proposed approach, described in Chapter 5, will ultimately be shown to improve the performance of CAM scheduling through prediction of the IPT and dynamic assignment of the RRI with the changing rate of transmission. In theory this approach is sound but the practically of the proposed approach

needs to be proven through consideration of the computational cost of the prediction and the timeliness of receiving the predicted result.

The currently implemented solution requires significant training time, with the current dataset of each training epoch which represents the time to train across all training samples taking approximately 47 minutes to complete. For the set of experiments completed the number of epochs was 20. Both the time per epoch and the number of epochs is highly dependent on the training data used and the complexity of the proposed model. Regardless, to complete training of a model can take more than 15 hours on the server employed to train and run subsequent experiments. This server was not designed for machine learning using the CPU rather than a dedicated GPU for both training and prediction as such the times taken above can be significantly reduced when a GPU is employed. As shown in [99] training time on a CPU of some 13 hours can be reduced to 2 when using a GPU. Ultimately however it is important to consider that models will be trained in an offline manner as such the trained model is what is deployed in-vehicle and training need not take place locally.

The key in-vehicle concern is inference time which ultimately dictates the practicality of the proposed solution. The current server using a CPU for prediction has a mean inference time of 202.0ms, standard deviation of 74.7ms and a 95% confidence interval reaching to 348.5ms, though the maximum recorded inference time was 329.2ms. As shown in [99] there can be significant improvements in inference time going from CPU to GPU in their case of image based prediction an inference time of 5s is reduced to 2-3s. Therefore swapping to a GPU can result in a halving of inference time which would mean the model performing with a mean inference time of 100ms.

For the proposed solution to be practical it would be required to have an inference time of 100ms or less, this would ensure that even in dynamic scenarios inferences are inline with the lowest possible RRI or highest transmission rate of 100ms/10Hz. As such, a GPU being part of an OBU would ensure the current model already could possibly be used as a practical approach. The key element would be that the model continuously makes inferences such that an up to date prediction is available for each 100ms interval so that when a packet requires transmission the prediction is already available and does not delay the transmission itself. It is also worth noting that the proposed solution was designed for highest accuracy, but it would be possible to implement a model that maintains high prediction accuracy but

also has an upper bound of inference time of 100ms. This would possibly result in degraded PDR performance when compared to the current approach but would ensure that the model could be deployed in an OBU with no risk of predictions delaying the transmission of the CAM packet.

4.3.4 Solutions for variable packet sizes

The final aspect where contributions are made is in variable packet size performance; this is argued through a detailed analysis of the MCS mechanisms in SB-SPS in dealing with variable packet sizes. This is the proposal which forms the basis of **contribution 5**. The aspects considered are as follows:

- First, this work considers MCS outside of the configured range proscribed by ETSI in [52], specifically showing that exceeding the maximum MCS of 11 can show significant performance improvements, especially when dealing with variable packet sizes. This corroborates the work of fellow academics such as Harri et al. [46].
- A key differentiating factor is that this work is the first to allow for the adaptation of MCS within the same simulation. Other studies focus on using either a fixed MCS only or various MCSs across different simulations. The standard and the approach adopted in this study allow the MCS to change dynamically for each transmission, as current studies they do not fully reflect the performance of the technology.
- This will be shown to offer the best overall performance by allowing the MCS to dynamically be set per packet size, ensuring the selection of the most robust MCS that ensures the lowest subchannel occupation.
- Finally, novel approaches are investigated to minimise the number of GBs when dealing with variable packet sizes.

The key element of the approach for dealing with variable packet sizes is based on the standard and specifically the fact that MCS adaptation is assumed in the standard enabling SB-SPS to dynamically adjust the MCS based both on the packet size and for congestion control purposes (MCS congestion control). This applies both at grant generation, where subchannel allocations are determined based on the current packet size, configured available subchannels and finally the configured MCS available. In addition to this the MCS can change dynamically throughout the grant enabling SB-SPS to change the MCS

ensures that for each new packet the best choice of MCS is possible for it. With the aim of both reducing GBs but also ensuring a robust transmission of each packet i.e. avoiding the use of high MCS when a lower MCS will ensure the packet is transmitted efficiently. This is often not enabled by default in the literature either adopting a fixed MCS such as is the case in [100] or enabling multiple different MCS across simulations but not allowing adaptation within the same grant as in [46]. As such, the first aspect explored is enabling this and it will be shown in Chapter 5 that this ensures the best performance of SB-SPS.

In addition to this, this study further investigates novel approaches for managing variable packet sizes. By default SB-SPS is designed to determine the sidelink grant based on the packet which causes grant generation, as such depending on the configuration i.e. subchannel number, allowed selection of subchannles, available MCS, the grant will be setup to optimise the transmission of packets of that size. In addition to this approach, this study further investigates two novel methods of grant generation:

1. First, the grant is generated based on the mean packet size of the current application. Thus, only the larger packets will result in a GB as the majority can be successfully scheduled in the typically generated grant.
2. Second, the grant is generated on the basis of the maximum packet size of the current application. This results in GBs being eliminated as it is not possible for a grant to be generated that cannot accommodate any generated packet size.

The principal behind this idea is that, as has already been discussed and highlighted in Section 4.2.3.2 GBs are a significant cause of the degradation of SB-SPS performance. In this case, the scheduler is designed to minimise them or eliminate them entirely, thus avoiding GBs at the cost of additional subchannel occupation. The goal of this proposed approach is to see if the trade-off ultimately improves SB-SPS scheduling performance or if the overhead of additional channel occupation is such that it is not worthwhile.

Fig. 4.15 illustrates the principal behind the proposed approach, in this case assuming that the generation is designed based on the maximum packet size, eliminating the need for grant breaking. A fixed MCS of 7 is assumed as a means of simplifying the explanation and highlighting the trade-off clearly. In such a scenario SB-SPS will schedule each grant with 4 subchannels, as this is

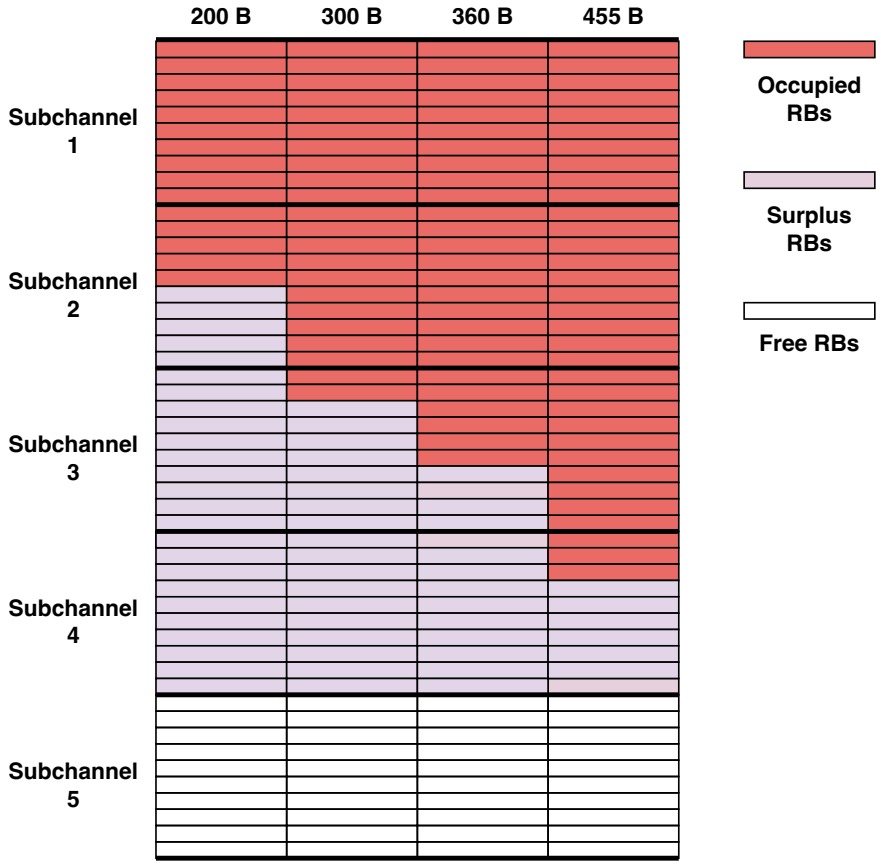

Figure 4.15: Over-Provisioning to avoid GBs

the maximum required number for a packet size of 455 Bytes. This results in the smallest packet being over-provisioned by 2 subchannels, but results in GBs being fully eliminated.

4.3.5 Proposed solutions conclusion

To this point, solutions have been introduced in three significant areas, first in congestion control, secondly in managing aperiodic traffic patterns, and finally in managing variable packet sizes. The quantitative analysis of the efficacies of each solution will be described in Chapters 5 and 6.

On congestion control, three standard compliant mechanisms are proposed, one focused on 3GPP standards and two on ETSI standards. Our solution $RRI_{Adaptive}$ introduces an algorithmic approach to C-V2X and NR-V2X eliminating the need to deal with difficult table tuning issues. Additionally, a comprehensive solution focused both on network performance, CBR stability, and importantly the fairness of the adopted scheme, is proposed. In this vein, this is an area which has limited focus when discussing C-V2X/NR-V2X even

though it is exceedingly important considering a vehicular scenario where unfairness may result in reduced safety due to vehicles being starved of resources and their dynamics remaining unknown to neighbouring vehicles. The quantitative evaluation of the efficacies of these is discussed in Chapter 6.

In dealing with aperiodic traffic patterns, the primary focus has been on the prediction of CAM IPTs, using this to enable more accurate scheduling in SB-SPS, thus turning an aperiodic application such that it mimics a periodic one. This approach will ensure that additional aperiodic applications, such as DEN, will be able to be scheduled with the NR-V2X dynamic grant mechanism and ensure better scheduling performance, in general.

The final topic where solutions have been suggested is the management of variable packet sizes; in this regard, this study focuses on how MCS can impact SB-SPS performance. The main note here is to corroborate fellow academics work where the use of high order MCS can ensure the best performance by reducing subchannel occupation. A second note which has often been overlooked in other academic studies is the point that ensures an adaptive approach rather than a fixed one throughout simulation will result in the best performance possible by avoiding both unnecessary robustness loss from too high an MCS and increased subchannel occupation from too low an MCS. The final note is primarily a note of caution in that adopting the standard approach for scheduling packets, i.e. based on optimal subchannel occupation for the current packet size, is the best approach as overprovisioning often only leads to increased congestion and reduced network performance from increased interference.

4.4 Chapter conclusion

This chapter initially introduced the primary challenges for both the NR-V2X and C-V2X scheduling mechanism SB-SPS those being the introduction of GBs and MTs. The chapter forms the **second contribution** (**C2**) of this work and directly relates to publication *2* though informs the other major contributions of this work and their related publications. The subsequent chapters will provide a comprehensive quantitative analysis that further supports this contribution. The conclusions of this chapter are across three main aspects and are as follows:

1. This chapter provided a theoretical background of the challenges facing

SB-SPS namely MT and GB. The key conclusions from this theoretical analysis are:

- MTs introduce inconsistencies in the sensing history employed in SB-SPS to predict future reserved resources.
- GBs cause consistent rescheduling events to occur, increasing the likelihood that vehicles will select the same resources.
- Importantly, research has not considered that SB-SPS can be configured through the *sl-reselectAfter* parameter which controls whether a GB occurs after an MT.
- Finally, the theoretical background identifies variable packet size as a special case which always results in a GB.

2. This chapter continued on to provide a discussion of the scenarios under which MT and GB occur. Further it provided indicative results as to how scheduling performance degrades in these scenarios. The key conclusions for this element are:
   - The main cause of MT is as a result of aperiodic packet inter arrival rates, which can be as a result of a TRC mechanism or the traffic being event-triggered.
   - Analysis is provided into how variable packet sizes can result in a GB.
3. The final element of this chapter focused on outlining the solutions to these challenges, outlining the mechanisms for contributions **C3**, **C4** and **C5** as well as their implementation. Those proposals are as follows:
   - An overall approach which will be adopted for congestion control and aperiodic packet inter arrival times was provided highlighting the *RRI Adaptation* mechanism.
   - For congestion control, three TRC based mechanisms are outlined all categorised as RRI DCC mechanisms.
   - To manage aperiodic traffic patterns, a CAM IPT prediction model is proposed.

- Finally, for variable packet sizes, several elements for MCS adaptation to manage these patterns are propsed. Namely, adopting an MCS range beyond the current standardised limit, ensuring that MCS adaptation occurs on a per packet basis and a novel grant scheduling mechanism to minimise GBs.

Chapter 5

Quantitative Performance Analysis: C-V2X/NR-V2X and Proposed Solutions applied to Aperiodic Scheduling and Variable Packet Sizes

5.1 Introduction

Chapter 4 theorised why GB and MT cause problems for the C-V2X and NR-V2X schedulers and discussed various reasons why they occur i.e aperiodic arrival rates, congestion control and variable packet sizes. It further provided a short quantitative analysis highlighting the negative impact they have on packet delivery rate and provided a theorised overview of our contribution (**C2**) to overcome these difficulties by adapting the RRI SB-SPS parameter.

This chapter includes:

- An in-depth quantitative analysis of why MT and GB disrupt the scheduler, focusing not only on PDR, but on the disruptive effects it has on SB-SPS in terms of grant usage, reservation outcome and colliding grants in Section 5.3.2. This is especially problematic when considering reliable communication of vehicle dynamics based ETSI standard applications. Furthermore, additional approaches to address this i.e SB-SPS parameterisation, alternative scheduling approaches and NR-V2X

are quantitatively evaluated. This includes an evaluation of the NR-V2X dynamic grant mechanism and its performance with these services in mixed traffic scenarios.

- Based on the observed shortcomings of all proposed solutions, including the very recently proposed NR-V2X dynamic grant mechanism, this chapter continues to provide an in-depth quantitative evaluation of **C3** i.e. the proposed CAM IPT prediction mechanism outlined in Sections 4.3.1 and 4.3.3, which overcomes the SB-SPS scheduling challenges in realistic vehicular environments.
- Based on the shortcomings associated with scheduling packets of variable size, an in-depth quantitative analysis is provided (**C5**) focusing on exceeding the current MCS limits defined in the standards, adapting the MCS dynamically for each transmission rather than a fixed MCS and finally the impact of different scheduling approaches seeking to minimise GBs. In this it will be shown that the current limits on MCS degrade scheduling performance. Additionally, adapting the MCS leads to better performance as it can dynamically change for each packet size rather than maintaining a fixed MCS. Finally the current adopted approach whereby the grant is configured based on the current packet size is the best approach for SB-SPS to adopt.

5.2 Simulation scenario and application patterns

The simulation environment (**C1**) was described in detail in Chapter 3 with common simulation parameters summarised in Table 3.1 in Section 3.5.1. This section highlights any parameter that vary from this is summarised in Table 5.1 and the reason for the change provided. The most notable change is the introduction of an increased number of subchannels and allocated RBs; this corresponds to the standard simulation parameters defined by 3GPP [12].

Although ETSI CA is the application that this chapter mainly focuses on, we also consider the 3GPP proposed aperiodic pattern [12]. As already described, ETSI CAMs are triggered based on vehicle dynamics, and as such the performance is dictated by the scenario employed. The 3GPP mechanism, on the other hand, is a random approach whereby the packet inter-arrival time is 50 ms plus a random exponential delay with a mean also of 50 ms. To characterise both application models, CDFs of the packet inter-arrival rates are

Table 5.1: Scenario Parameters.

Parameter	Value
Vehicular scenario	
Vehicular density (β)	0.6 & 0.12 & 0.2 & 0.3 veh/m
Number of lanes	3/6 in each direction (6/12 total)
Vehicle Speed	70km/h
Channel settings	
No. subchannels & Size	3 & 16 RBs
Application layer	
Transmission frequency	3GPP or ETSI models
MAC & PHY layer	
Transmission power	23 dBm

shown in Fig. 4.3, for a highway scenario where depending on the density the vehicle speed is reduced to mimic a congested traffic scenario. In the 3GPP model, close to 100% of the packet inter-arrival times are below 200ms. The ETSI packet inter-arrival times are dependent on vehicle dynamics and thus increase with the vehicular density.

5.3 Quantitative Analysis of the negative impact of aperiodic application traffic on SB-SPS

This section builds on the theorised discussion of scheduling challenges posed by GB and MT in Sections 4.2.1 and 4.2.3. This showed the negative impact at a high level with respect to PDR. Now we provide:

- We next investigate how the characteristics of the service i.e. IPT has a large impact by considering additional application service models (3GPP and single slot). This quantitatively investigates our theorised discussion in Section 4.2.1.
- An in-depth analysis of the negative impact of aperiodic arrival rates on SB-SPS with a focus on reservation outcome, grant usage and the number of colliding grants, to investigate how radio resources are mismanaged.
- We provide an in-depth analysis of alternative scheduling mechanisms for aperiodic traffic but that are outside of SB-SPS and in particular the dynamic grant method in NR-V2X.

- Finally, we quantitatively investigate the performance of dynamic grant with varying levels of aperiodic and periodic application traffic and conclude that when fully aperiodic the performance is not adequate. This provides the benchmark for the analysis of the proposed solution in Section 5.4.

5.3.1 Impact of different application traffic models

This section investigates how the C-V2X SB-SPS Mode 4 standard performs when application traffic follows an aperiodic arrival rate. Application traffic is modelled as follows:

- **Periodic (Default):** Static CAM transmission every 100ms with a default sensing window of 1000ms.
- **Aperiodic (3GPP):** Application layer packets arrive every 50 ms + an exponential distribution with a mean of 50 ms [101].
- **Aperiodic (ETSI):** Importantly, the previous literature often assumes a periodic CAM transmission rate. This does not align with the aperiodic transmission of CAMs according to the ETSI specification [15] where CAMs are triggered based on a vehicle's dynamics, that is, deviation in heading (>4°), position (>4m) and speed (>0.5m/s) or at 1s intervals if these conditions are not satisfied.
- **Periodic Single Slot Usage:** Application traffic is transmitted every 100 ms but the SB-SPS grant is deliberately broken after a single reservation as this represents the most inefficient use of SB-SPS allocated resources and can be considered as a worst case scenario.

Fig. 5.1 explores the performance of SB-SPS with aperiodic traffic when grant breaking is enabled and disabled, that is, the *sl-reselectAfter* parameter is set to one resulting in a grant break after each transmission or the grant will always be maintained despite missed transmissions. This non grant breaking approach is assumed by Molina-Masegosa et. al [36] but is not compared to a grant breaking scenario. Maintaining the grant is the default behaviour for the SB-SPS algorithm. Fig. 5.1a shows the performance of the 3GPP application model with and without grant breaking, where an average PDR increase of 5% is achieved when grant breaking is disabled. This increase is attributable to a reduction in collisions due to reduced resource rescheduling. In the case of ETSI, as seen in Fig. 5.1b, with PDR improvements up to 10%. This even

outperforms periodic traffic due to much lower CBR, in subsequent results, the periodic rate will be adjusted to mimic the congestion levels of aperiodic results. In particular, the improvement in performance across both application models derived from disabling grant breaking comes at a cost. In both cases, the models exhibit reserved but unused resources. Fig. 5.1c illustrates the challenge of enabling grant breaking, in this figure the outcome of a reservation is highlighted. The elements of this are as follows:

- **Reserved but unused:** This represents the number of reservations that were ultimately unused due to MT and impact sensing history.
- **Advertised and used:** This represents the number of reservations correctly used, i.e. SCI sent and subsequently a TB sent in reserved resource.
- **Not advertised and used:** This represents the number of transmissions where no SCI was sent previously reserving the resource, this occurs after an MT if the grant is maintained and a packet becomes available at an upcoming transmission opportunity.
- **First Transmission:** This represents the first transmission in a reservation, this is a special case as it will be unadvertised but cannot have been reserved as the grant has only been generated.

The 3GPP model has 29% reserved but unused resources when GB is enabled as shown in Fig. 5.1c, which remains the same when GB is disabled (No-GB). A similar outcome occurs for ETSI with the same number of reserved but unused across GB and No-GB with only a 3% difference, 46% and 43% respectively. There is however the introduction of an additional source of error, as discussed by Harri et al. [102] when unintentional collisions are introduced. This occurs when a slot goes unused but the grant is maintained. As an SCI goes unsent neighbouring VUEs may mistakenly believe the resources to be free. This is shown with the *Not advertised and used* (red bar) introduced by the No-GB configurations. There is a clear trade-off between the two configurations, with GB resulting in significant increases in the number of *first transmissions* that occur (blue bar), while No-GB introduces a new source of error. It is also clear that grant breaking will cause the performance to degrade towards the worst-case single slot performance, and the degradation is closely linked to the traffic pattern being employed. This ties back into an aspect that has already been discussed in Chapter 4, specifically in that the traffic pattern and the RRI

are closely linked and SB-SPS has some ability to manage aperiodic traffic as long as it exhibits rates similar to the RRI. This is generally the case with the 3GPP traffic pattern, where the rate is in the vast majority of cases below 200 ms; given that the RRI in this case is set to 100 ms, it therefore fits into the already mentioned 2n-2 rule. In this case, that will be 198 ms being the maximum packet inter-arrival time which can be accommodated, and typically the 3GPP pattern will exhibit an IPT less than 200 ms ensuring it does not always trigger rescheduling or a missed transmission. This is why regardless of which configuration there are significantly greater number of *Advertised and used* reservations (black bar). In the case of ETSI with the $\beta = 0.12vpm$ case, this will always break the 2n-2 rule and as such will incur a significantly increased number of MTs resulting in much degraded performance, especially in the grant breaking case as a result of greatly increased rescheduling events, with *first transmissions* as high as 46%. Only in the case of purely periodic traffic though can it be seen that the mechanism is operating in a fully effective manner with over 90% of packets being transmitted with a reservation.

These experiments formed the early elements of this study when the initial development of OpenCV2X was completed and formed a basis for the investigation to come. Later in the project, 3GPP released the NR-V2X standard, and this introduced several significant changes to the SB-SPS mechanism employed in C-V2X. These have been discussed in Section 2.3.4, but to highlight the most significant changes made. The first change of importance in this aspect of the work is the fact that the RSSI filtering stage was removed from the CSR selection process, as such any CSRs which do not appear reserved, i.e. have an SCI and an RSRP above the preconfigured threshold, will be possible CSRs reported to the MAC layer. The other change is the adjustment of the RSRP filtering stage; this has been changed to only consider the most recent reservation as opposed to averaging across a number of previous reservations. These two changes have a significant bearing on the performance changes between C-V2X and NR-V2X iterations of SB-SPS and the primary goal of their introduction was to improve scheduling performance when these aperiodic patterns were being dealt with. In subsequent results the focus will be primarily on the ETSI traffic pattern, this is both due to the fact that it exhibits the greater challenge to SB-SPS from a scheduling perspective, but also represents a more realistic application traffic pattern being based on an implemented standard.

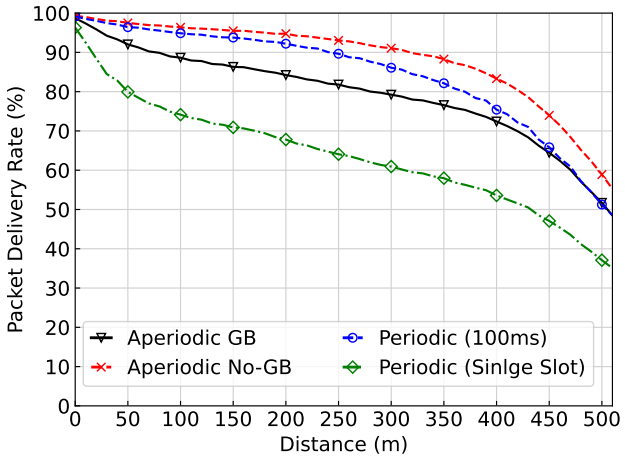

(a) 3GPP application model.

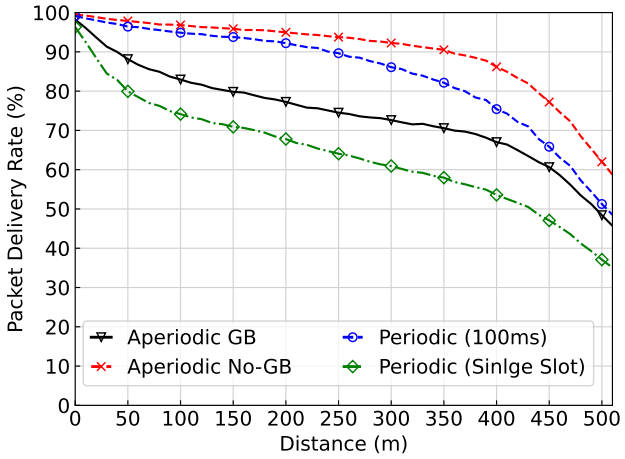

(b) ETSI application model.

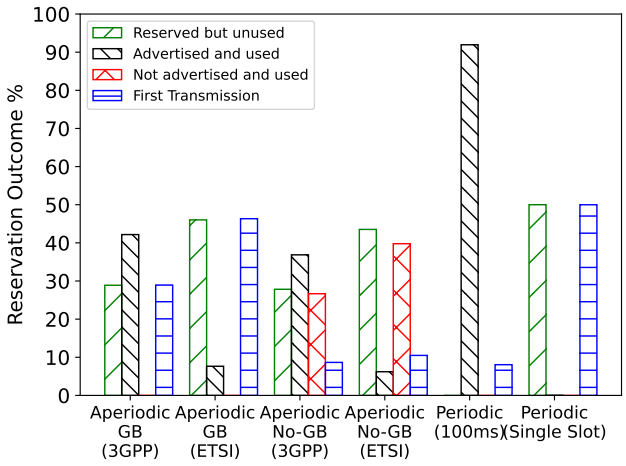

(c) Reservation Outcome.

Figure 5.1: C-V2X SB-SPS performance, (β=0.12 veh/m).

5.3.2 Impact of Grant Breaking (GB) and Non-grant Breaking (No-GB)

Fig. 5.2 highlights performance issues for both C-V2X and NR-V2X when dealing with this aperiodic traffic pattern. This analysis is similar to that provided in Section 4.2.3 with an increase in vehicular density that highlights the degraded SB-SPS scheduling performance more severely. A substantial decline in PDR is observed for C-V2X and, to a lesser extent, for NR-V2X. The cause of performance degradation is as a result of missed transmissions, i.e., when a packet is unavailable to send but the granted resource becomes available.

The trend for C-V2X is as already highlighted in the previous Fig. 5.1 and is unchanged in these results; ultimately, grant breaking has a significant impact on PDR performance. Elimination of grant-breaking, while not so severe, still results in a degradation in performance. Importantly in this set of results, the periodic sequence is configured to exhibit the same transmission rate as the GB and no-GB configurations, as such it represents the expected performance of SB-SPS at this density of network traffic. It is clear, however, in Fig. 5.2b that this does not hold when moving to the NR-V2X standard. Although no-GB still exhibits similarly degraded performance in the case of the GB configuration, performance is greatly improved over C-V2X, though no-GB is still preferable in both cases. This is as a result of the removal of the RSSI filtering stage in NR-V2X, it will be shown in a subsequent section as well as in the C-V2X context. Importantly, the removal of this stage results in a much more diverse CSR selection pool; otherwise, in the case of C-V2X, the pool is automatically reduced to the top 20% of CSRs. Due to the fact that neighbouring vehicles will have similar sensing windows, this compounds the issue of constant reselections, as neighbouring vehicles will be constantly selecting from similar pools, which results in an increased number of collisions.

This increase in collisions is reflected in the outcomes for the reservations as shown in Fig. 5.3. This analyses the outcome of the reservations as outlined inSection 5.3.1. There are 4 categories, namely reservations that go unused, reservations properly used, transmissions sent without a reservation, and finally the first transmission in a grant. The main point of comparison is between the GB and No-GB configurations as the trends remain consistent across both technologies with only marginal differences between both. It is clear that when moving to aperiodic traffic patterns the scheduling encounters

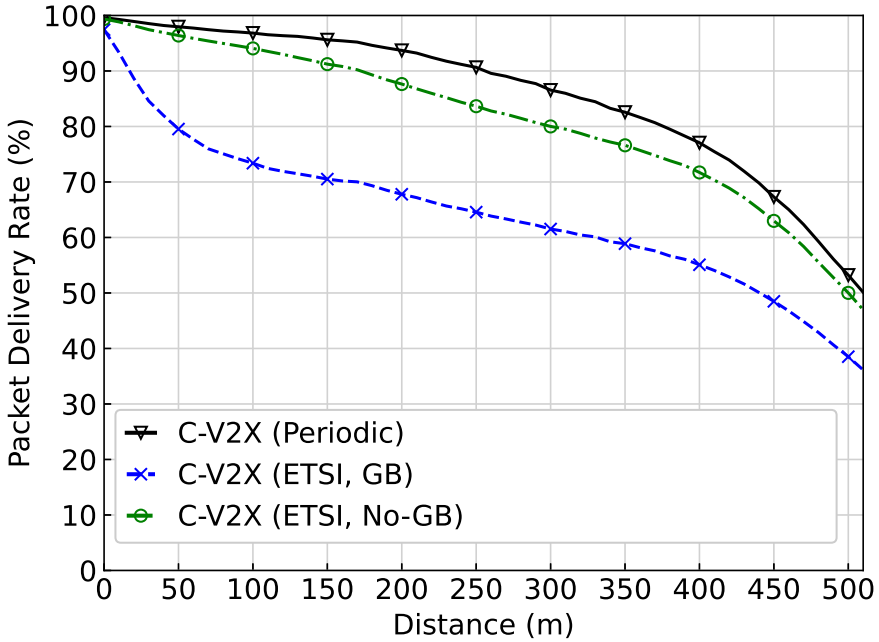

(a) C-V2X SB-SPS.

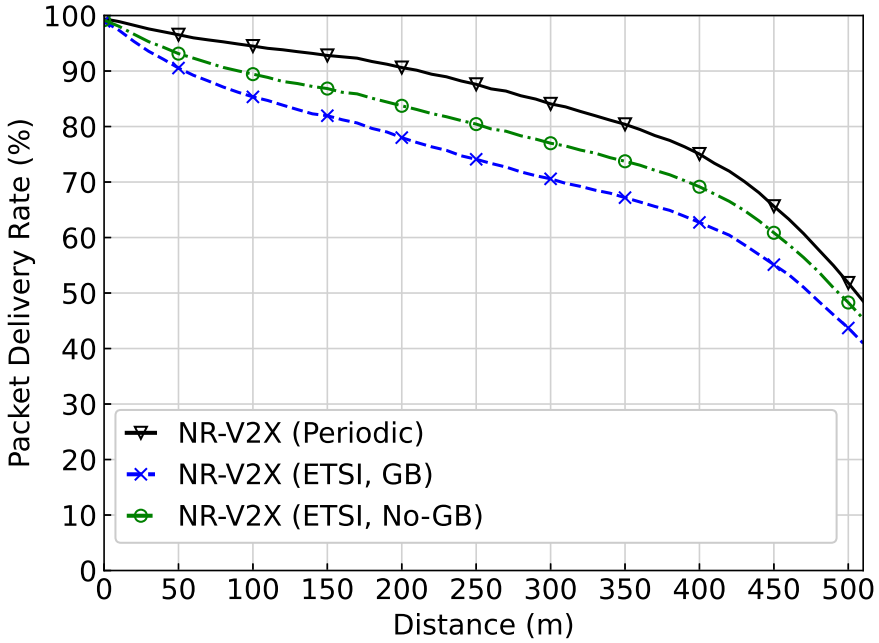

(b) NR-V2X SB-SPS.

Figure 5.2: Impact of ETSI aperiodic inter-packet times (IPT) on Cellular Vehicular Standards, (β=0.3 veh/m).

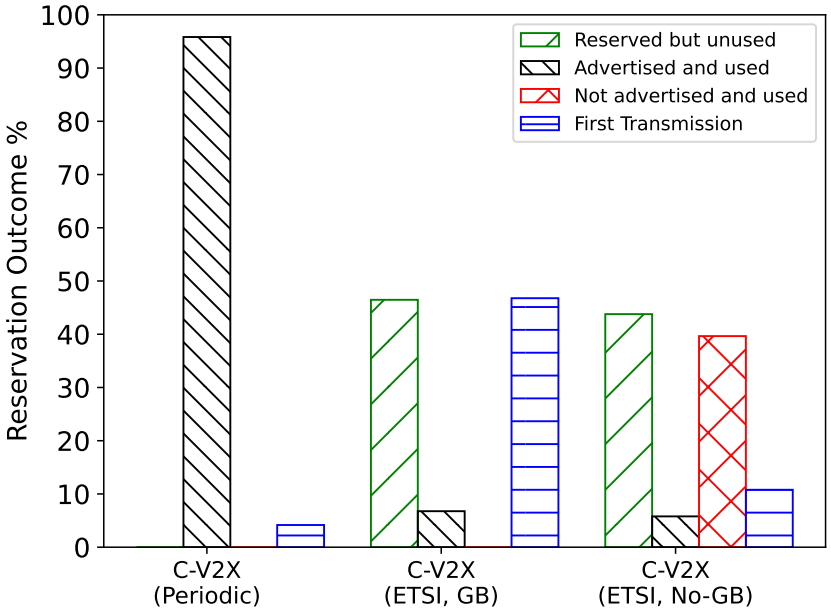

(a) C-V2X SB-SPS.

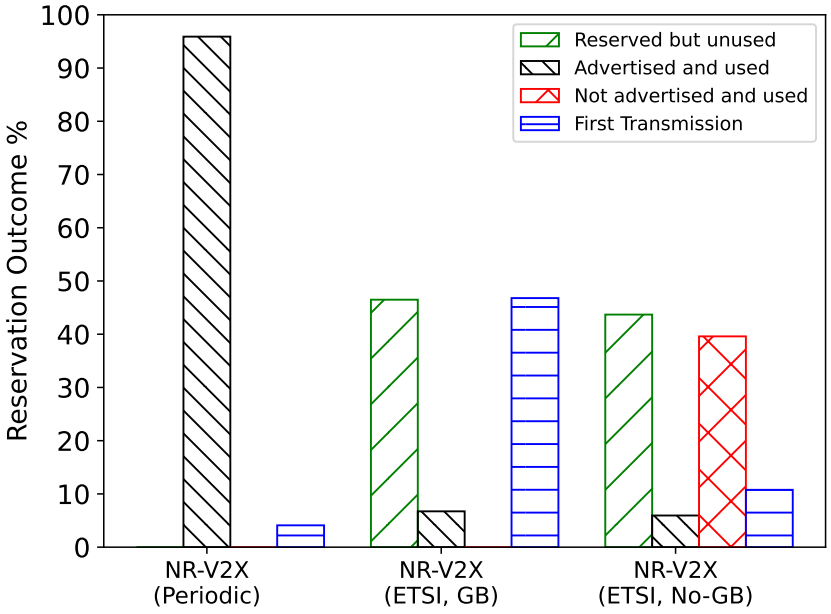

(b) NR-V2X SB-SPS.

Figure 5.3: Impact of ETSI aperiodic IPT on reservation mechanism, (β=0.3 veh/m).

significant challenges with the majority of transmissions not being scheduled correctly. This is shown by the high levels of *reserved but unused*, *not advertised but used* and *first transmissions* reservations with both GB and No-GB. In the case of GB there is a considerable increase in the number of *first transmissions* which is logical given that in this case the vast majority of grants are broken after the first transmission, resulting in rescheduling. The No-GB scenario these *first transmissions* are reduced significantly but then introduce another source of error in transmissions occurring without reservations.

The impact of these reservation issues is reflected in Tab. 5.2, which reports on key grant information, including the number of transmission per grant and scheduled grants that incurred a collision. The collisions considered are γ_{Mt} which represents collisions which occur as a result of a missed transmission, this will be as a direct result of transmissions occurring without a reservation. The second collision is γ_{TSim} which are collisions that occur when multiple vehicles make reservations at the same time, this will occur for the first transmission in a grant. The first element of note is the mean grant transmissions, in the case of periodic across both technologies, as expected, there is a mean of 10 transmissions per grant which corresponds to the RC ranging from 5-15. Both technologies exhibit similar numbers of grant collisions and exclusively due to scheduling at the same time γ_{TSim}. In the grant breaking configurations, the mean grant transmissions drops to a single transmission per grant, in this case there is an extreme increase in the number of γ_{TSim} collisions, which is directly as a result of the continuous rescheduling events that occur as a result of the GB. The final configuration then No-GB, increases the number of transmissions to 3.47 per grant with a high level of variance and reduces the number of γ_{TSim} errors. Importantly, No-GB introduces new γ_{MT} errors. Interestingly, NR-V2X appears to manage the GB scenario more effectively with fewer γ_{TSim} which explains its better overall performance in this mode, while in No-GB configuration both approaches perform similarly.

To show the impact of both the reservation issues and grant collisions, Fig. 5.4, focuses specifically on the packets lost due to interference, which is a direct result of a packet collision and poor scheduling. It is quite evident here that at near distances C-V2X GB (blue dashed) in Fig. 5.4a has many increased collisions. In the case of NR-V2X as shown in Fig. 5.4b this is much reduced with GB incurring while still more collisions not quite as frequent as in C-V2X. Importantly, however, in both cases these changes can have some minor

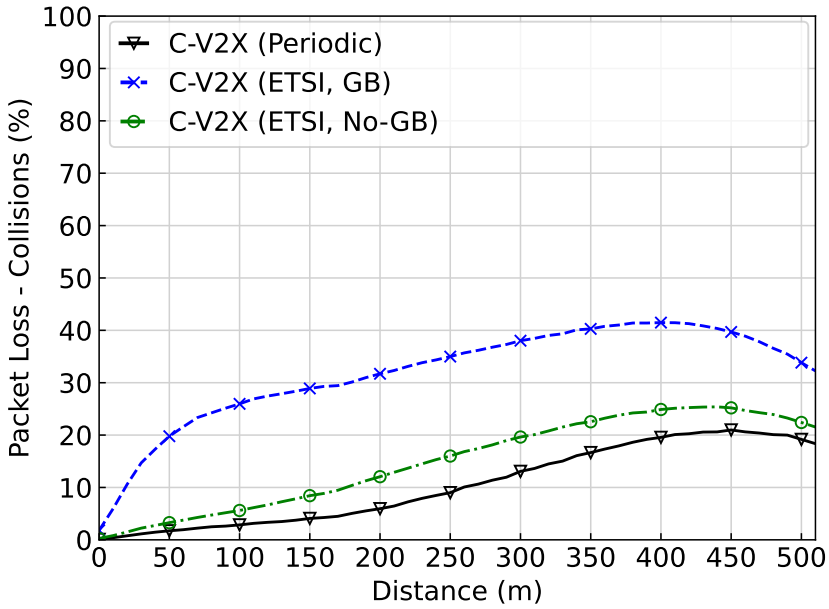

(a) C-V2X SB-SPS.

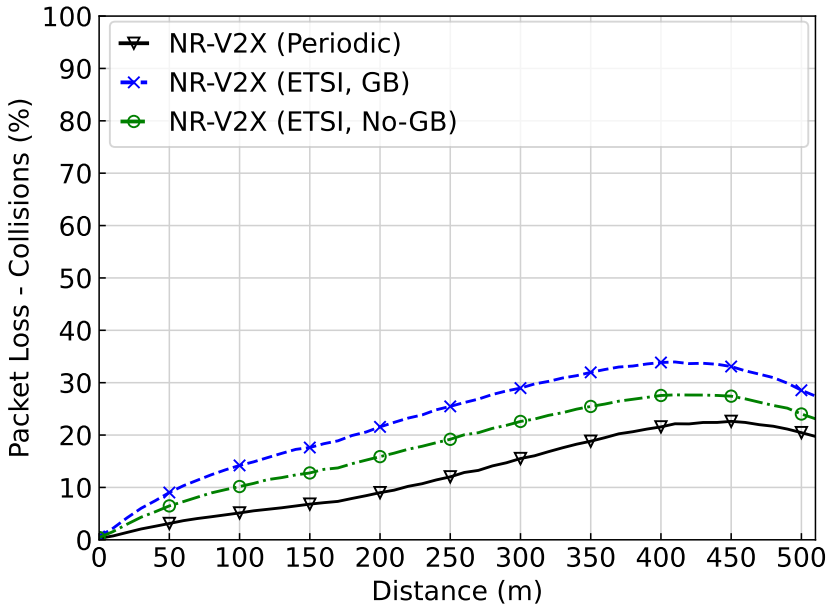

(b) NR-V2X SB-SPS.

Figure 5.4: Aperiodic traffic patterns and impact on interference errors.

Table 5.2: Key grant information (Grant transmissions and colliding grants).

Scheduling Mechanism	Mean Grant Transmissions	γ_{MT}	γ_{TSim}	γ_{Total}
C-V2X (Periodic)	9.21 ± 3.05	0	64	64
NR-V2X (Periodic)	9.24 ± 3.09	0	51	51
C-V2X (ETSI GB)	1.01 ± 0.28	22	5910	5932
NR-V2X (ETSI GB)	1.01 ± 0.28	9	3879	3888
C-V2X (ETSI No-GB)	3.47 ± 3.17	196	495	691
NR-V2X (ETSI No-GB)	3.47 ± 3.18	448	186	634

impacts on overall performance, with the NR-V2X performance in the case of periodic traffic being marginally lower than that of C-V2X at 2%. This is likely an acceptable trade-off from a standard perspective given the propensity of applications that have aperiodic behaviour. The choice of only a minor degradation in performance for periodic to ensure that in a GB configuration the standard does not degrade so significantly is an acceptable one.

Ultimately, the conclusion from this section is that both C-V2X and NR-V2X are not as effective in dealing with aperiodic traffic as they are in dealing with periodic arrival rates. This is a result of their grant-based approach, which requires an accurate sensing history to correctly predict which resources will be used in the future are reserved. Secondly, while NR-V2X has made an improvement especially in the case of a GB configuration, it still incurs similar performance degradation when dealing with aperiodic patterns. This is important in the context of ETSI services as these will generally exhibit aperiodic patterns and incur such performance losses.

5.3.3 Aperiodic scheduling techniques: SB-SPS parameter tuning

The first approach to overcoming these challenges is to focus on tuning the SB-SPS algorithm itself and to optimise its parameters to deal with aperiodic traffic patterns. This section quantitatively evaluates the performance of approaches that either tune the SB-SPS algorithm or deactivate different aspects of the algorithm.

Specifically, the following mechanisms are evaluated:

- **SB-SPS with a reduced sensing window (labeled SW):** The premise of a reduced sensing window is that past reserved sources are not a

relevant predictor for the selection of future resources due to the aperiodic packet arrival rate. This resource allocation mechanism was briefly evaluated in [38], proposed by Bazzi et. al for CAMs in [103] and suggested for DENMs in [104]. In this evaluation we consider shorter sensing windows of T_s=*500ms*, *200ms* and *100ms*.

- **SB-SPS with no RSSI filtering (labelled no RSSI)**[105]**:** It is suggested to remove the RSSI filtering stage which typically filters the best 20% of RSSI-ranked CSRs. This increases the number of CSRs reported to the MAC for potential selection, reducing the potential for collisions due to multiple VUEs selecting the same resource [105]. This ultimately was incorporated into the NR-V2X standard, in this set of results though the RSRP filtering stage remains the same as in C-V2X as such these results are not directly NR-V2X.
- **Random Scheduling:** Resources are chosen at random each time a TB is transmitted. On packet arrival, a single subchannel is selected at random from the subsequent 100 subframes. This is considered the baseline resource allocation mechanism.

Fig. 5.5 shows the performance of these schemes for the 3GPP and ETSI application models, excluding the disabled grant-breaking option, which is explored separately. In the case of 3GPP, both removing RSSI filtering and reducing the sensing window results in improved PDR (or comparable performance in the case of a sensing window of 500ms) with all scenarios outperforming the base case of random resource selection. This is because in aperiodic scenarios, historical RSSI filtering, which is based on averaging all previous transmissions on the available resources, leads to an inaccurate prediction of how future resources will be utilised. Therefore, eliminating RSSI filtering also increases the randomness of possible CSRs, as this stage typically reports only 20% of possible CSRs, based on the lowest RSSI. Removing this filter means that only those packets that are deemed reserved by the RSRP filtering stage will be excluded. In this case, while the proposed CSRs can still be inaccurate as a grant break may have occurred, it is unlikely to cause large amounts of resources to be removed from the CSR pool in lower density scenarios, thus increasing choice. An even larger improvement in performance is shown in the case of the ETSI model (up to a maximum of 14%), attributable to much lower CBR and clearly showing that the current operation of the SB-SPS mechanism cannot properly deal with aperiodic traffic patterns.

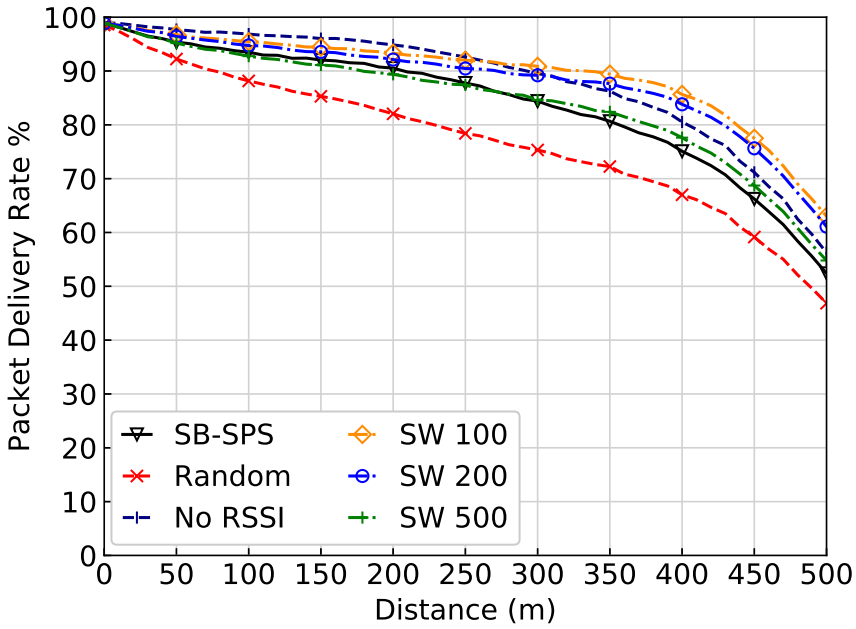

(a) 3GPP application model.

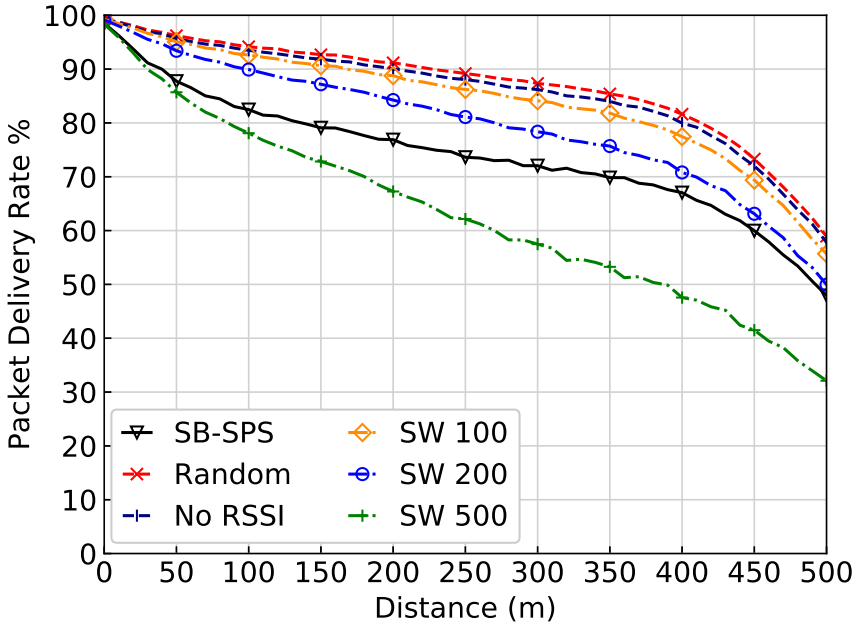

(b) ETSI application model.

Figure 5.5: SB-SPS performance for reduced Sensing Window (SW) duration and removed RSSI filtering (β=0.12 veh/m).

Reducing the sensing window (SW) duration also shows some improvement for both application models, with a maximum gain of 11% and 9% for the 3GPP model, and 12% and 8% for the ETSI model for the sensing windows of 100ms and 200ms respectively. In this case, the performance is improved by minimising the impact of RSSI filtering on the selection of possible CSRs, with only the most relevant information maintained for the RSRP and RSSI filtering stages. Notably, a sensing window duration of 500ms performs significantly worse than the default 1000ms for the ETSI model and is comparable for the 3GPP model. This highlights a deficiency in using a reduced sensing window as a means of accommodating highly aperiodic traffic, as it needs to be heavily parameterised depending on the characteristics of the application model and density of recorded RSSI values. As 500 ms is higher than the lower duration sensing windows, it performs additional filtering based on RSSI, thereby reducing the number of potential CSRs reported to the MAC layer. By continuation the default 1000ms sensing window should perform worse again. However, this is not the case, as there is increased diversity in the CSR selection reported to the MAC by each VUE for 1000ms. For example, for a sensing window of 500ms, the best 20% of the possible resources may not have recorded RSSI values and, as such, are reported directly to the MAC. This results in less diversity in the CSR selection across multiple VUEs and hence an increased likelihood of collisions due to choosing the same resource. However, for a 1000 ms sensing window, only 10% of the resources have no recorded RSSI values, and the remaining 10% have been chosen based on the RSSI values sensed, leading to greater diversity in CSR selection between multiple VUEs. However, by far the most interesting observation for the ETSI application model, which demonstrates highly aperiodic characteristics, is that while reducing the sensing window and removing RSSI filtering does offer performance improvements, the best performance is achieved by simply randomly allocating the resources. This forms a baseline against which to compare the resource allocation mechanisms evaluated in the next section.

The primary result of this section is the fact that tuning, while possible, still comes with caveats and exceptions, with certain traffic patterns it can operate quite effectively as can be seen with results from the 3GPP pattern where sensing window changes can cause somewhat significant improvements. The problem is that, in different patterns, there are different results. In the case of ETSI, while SB-SPS can be improved upon such as through the elimination of RSSI filtering ala NR-V2X or sensing window changes, it is clear that just

randomly selecting a resource in this case is the most effective option. This would obviously point to the fact that SB-SPS in and of itself is not entirely effective at dealing with these traffic patterns and is more a hindrance than a help. Importantly, looking at other options outside of it is the best approach at the current moment in time.

5.3.4 Aperiodic scheduling techniques: Dedicated mechanisms

An alternative approach is that C-V2X and NR-V2X could propose a totally different scheduling mechanism to handle aperiodic applications. Ultimately, a Dynamic Grant mechanism has been standardised to deal with aperiodic traffic patterns, but this is explored in-detail in the next section. As such, it is still informative to evaluate alternate scheduling schemes from the literature. Those considered are:

- **Short-term reservations (labeled STR) [20]:** Proposed by Ericsson, short-term reservations also referred to as short term sensing, makes use of two selection windows. It works by sending a reservation signal in the first selection window, which reserves resources in the second selection window on which to transmit the SCI and TB pair. In this mechanism, the resource on which to send the reservation signal is chosen randomly, discounting the resources reserved for SCI and TB pairs. Until the reservation signal is sent, the vehicle continues to listen in case the reservation slot becomes reserved by another vehicle. At the time of sending the reservation signal, a future free resource is selected if it has not already been reserved in the second sensing window. This reduces contention for the transmission of SCI and TB pairs.
- **Counter Based Mechanism [21]:** This approach makes use of a simple counter to increase the randomness of the resource selection process. When receiving an application layer packet, a counter is randomly selected between {1, 40}. In each subframe, the counter is decremented by the number of free subchannels. Once the counter reaches 0, a free subchannel is chosen at random in the next available subframe where the SCI and TB pair will be transmitted.
- **Random Scheduling:** As per the previous section, resources are chosen at random from the next 100 subframes.

Fig. 5.6a shows that the short-term reservation and the counter-based mechanisms do not improve the performance of the 3GPP application model, which demonstrates little performance degradation in comparison to the periodic application model due to low divergence in the packet inter-arrival rates (staying in the bounds of *[n, 2n-2]* as already defined). At closer proximities, the STR model performs similarly to SB-SPS but performance quickly drops in accordance with the counter mechanism by 400m. This is due to an increase in congestion of 24% in the PSCCH as shown in Fig. 5.6c caused by additional reservation signals.

Fig. 5.6b shows that the short-term reservation and counter based mechanisms considerably improve the PDR performance for ETSI traffic when compared to SB-SPS. This is because these scheduling mechanisms avoid the pitfalls of incorrectly labelling resources as reserved, thereby artificially reducing the available CSRs, with the resource subsequently going unused. It is worth noting that the counter mechanism performs comparably to a simple random allocation approach, so that the additional complexity does not provide much benefit in terms of additional performance. Overall the short-term reservation performs the best, although this is at the expense of additional overhead in the PSCCH due to the reservation signals, increasing load within the network by approximately 16% as shown in Fig. 5.6c.

Ultimately, the performance of the 3GPP application model is considerably worse due to the fact that, while aperiodic, it maintains the SB-SPS grant for longer than the ETSI. As discussed, grants are maintained longer because the 3GPP packet inter-arrival times often do not exceed *2n-2*, where n is the default RRI (100 ms in this case). Thus, the variability in the packet inter-arrival times strongly impacts which scheduling mechanism is effective. As noted earlier in Fig. 4.3 the 3GPP inter-arrival times are mostly confined within a small bounded time duration, whereas it can be observed that for the ETSI model they are distributed more widely, often exceeding *2n-2*.

It is clear from these results that alternative scheduling mechanisms have been proposed and can provide some improvements. There are additional challenges with these approaches specifically that the mechanisms are not standard compliant and require integration with SB-SPS (the proposed approach in Section 5.4 operates within the standard). Regardless of the integration element which can be overcome, the primary challenge is that ultimately the preferable solution in STR does incur a significantly increased

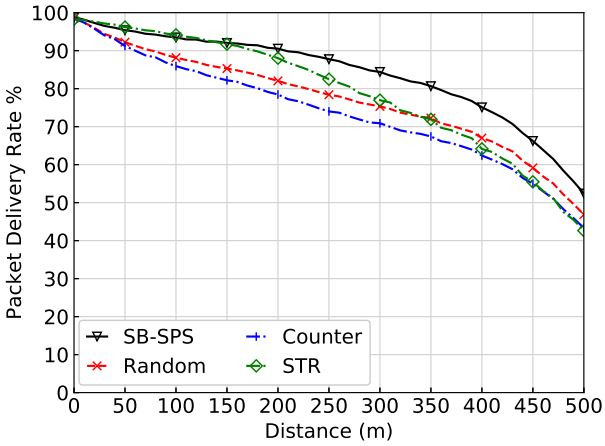

(a) 3GPP Application Model.

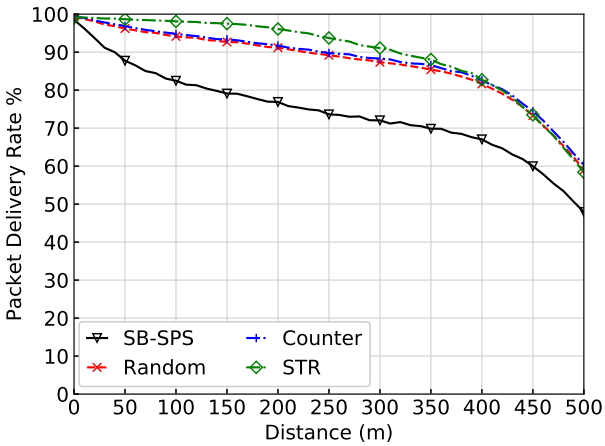

(b) ETSI Application Model.

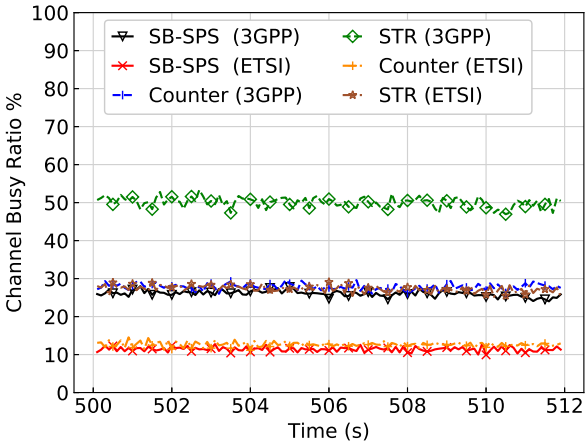

(c) PSCCH (Control Channel) Channel Busy Ratio (CBR) Comparison (β=0.12 veh/m).

Figure 5.6: Performance of dedicated aperiodic scheduling mechanisms for different application models.

CBR in PSCCH, which when congestion increases as a result of increased network traffic, can result in performance falling significantly. As such, these solutions do not outright solve the problem, but may form another approach outside of dynamic grant, and they may be especially relevant with some very high priority traffic where the additional overhead may be warranted.

5.3.5 Aperiodic scheduling techniques: NR-V2X Dynamic Grant

Thus far the analysis has highlighted the limitations of both C-V2X and NR-V2X SB-SPS for two application models (3GPP and ETSI) and having explored tuning parameters within SB-SPS and some alternative scheduling mechanisms outside the scope of SB-SPS. In the NR-V2X standard dynamic grant was introduced (described in Section 2.3.5) to operate using a grant-less approach. In essence, this means that it uses the same underlying scheduling mechanism, i.e. SB-SPS but does not reserve a grant on selecting a CSR (it will maintain a grant only for the number of retransmissions if they are configured).

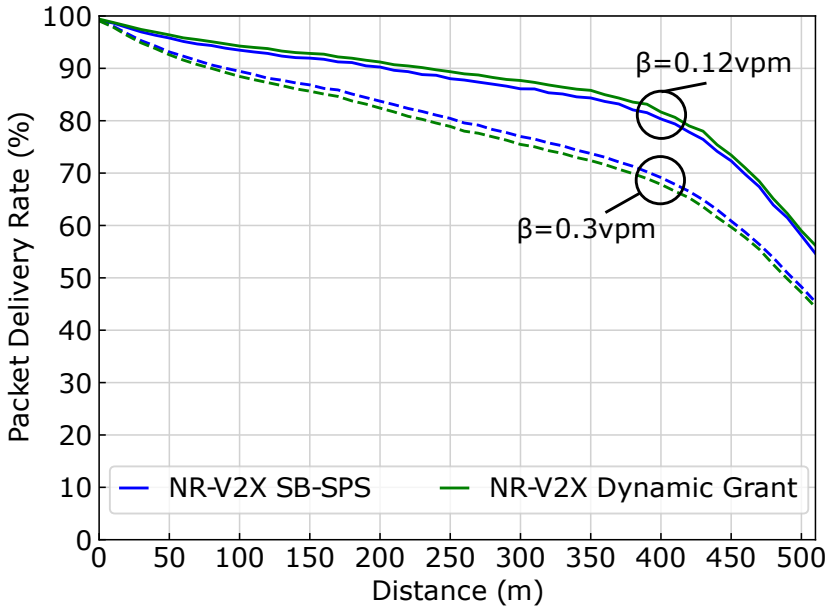

Figure 5.7: NR-V2X dynamic grant performance for fully aperiodic applications traffic. β=0.12 veh/m and β=0.3 veh/m.

Fig. 5.7 shows the performance of NR-V2X dynamic grant versus NR-V2X SB-SPS for the ETSI aperiodic model. Disabled grant breaking is assumed. It is evident that at a lower density of 0.12 veh/m, the dynamic grant mechanism

can marginally outperform SB-SPS. However, when the density increases to 0.3 veh/m, dynamic grant marginally under performs SB-SPS. This is due to the reliance on the SB-SPS sensing window. In the case of fully aperiodic application traffic, such as ETSI, all TBs are scheduled using dynamic grant. As such, all sensing history is irrelevant, as no filtering will be done in the selection window. Ultimately, this results in dynamic grant operating as a simple random selection mechanism where any resource in the selection window is equally likely to be selected.

This is very significant as fully aperiodic ETSI and DENM messages (ignoring retransmissions) are the basis for cooperative vehicular services, and it is widely envisaged that messages associated with increased vehicle autonomy and cooperative sensing will also demonstrate highly aperiodic arrival rates. Given this, it highlights the need to either develop a dedicated scheduling mechanism to better handle traffic with aperiodic arrival rates or alternatively to ensure that application traffic mimics periodic arrivals in order for dynamic grant to function more effectively.

It is worth considering a comparison between dynamic grant and the two alternative scheduling mechanisms already discussed, (counter and STR), given that all three are aimed at resolving the issue of aperiodic traffic patterns.

It is clear from Fig. 5.8 that counter demonstrates negligible improvement over NR-V2X dynamic grant as it essentially resembles a random scheduling mechanism. This highlights that dynamic grant in the absence of periodic traffic ultimately becomes a purely random mechanism. STR demonstrates improved PDR but comes at the cost of significantly higher CBR as shown in Fig. 5.8b. This is attributable to the extra control overhead generated by reservation signals before SCIs and TBs are transmitted. This limitation would result in even worse performance at higher densities, as the reservation signals would become a cause of collisions. While STR has limited applicability for ETSI CAMs, this approach could be applicable for high-priority aperiodic traffic or traffic generated only sporadically, such as DENMs.

5.3.6 Mixed Traffic scenario evaluation

We have now shown that the performance of NR-V2X dynamic grants is poor for fully aperiodic applications, which are characteristic of realistic vehicular

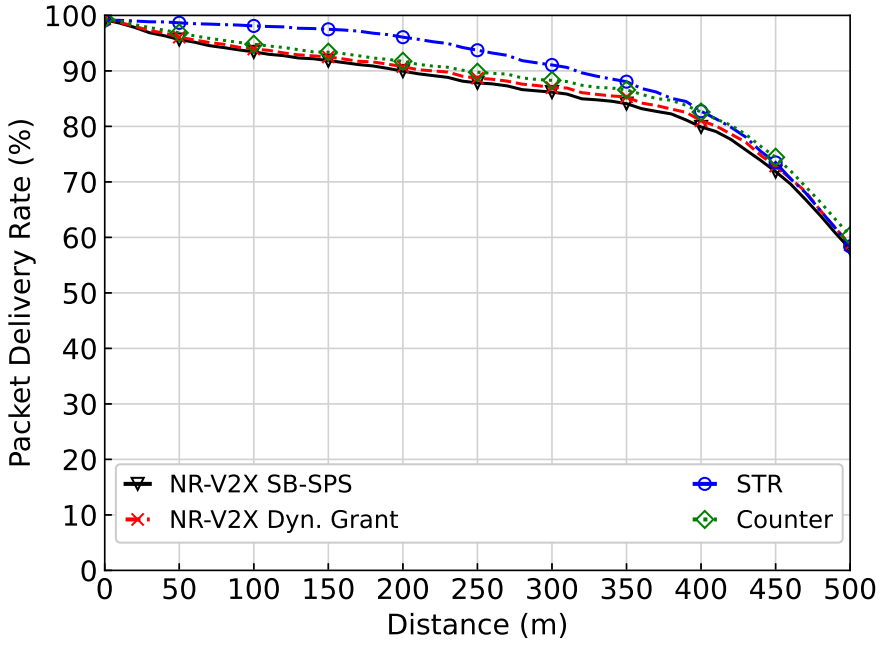

(a) NR-V2X PDR when for SB-SPS, dynamic grant & dedicated aperiodic scheduling.

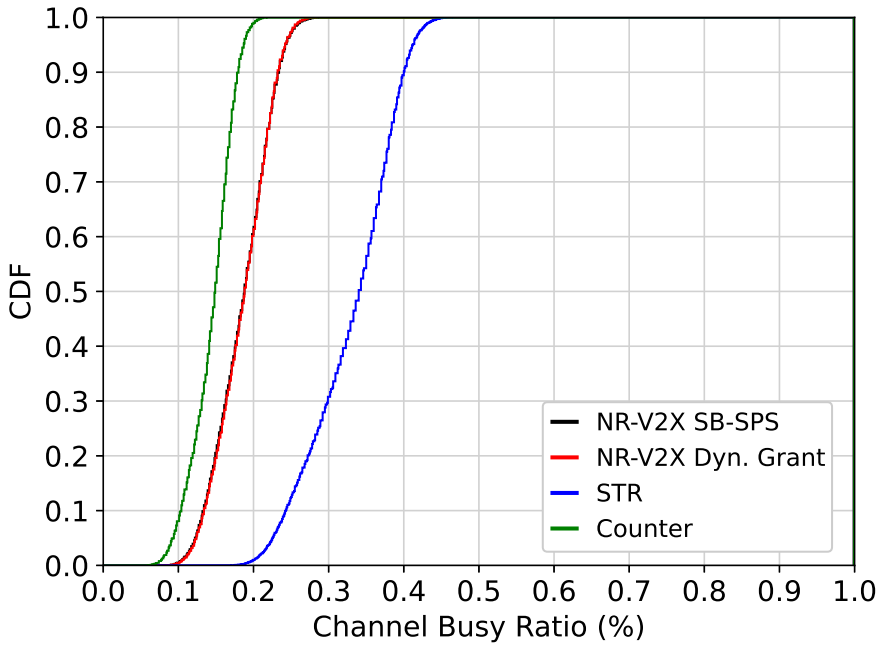

(b) Performance of default NR-V2X SB-SPS versus NR-V2X dynamic grant vs 2 alternative aperiodic scheduling mechanisms.

Figure 5.8: Performance of NR-V2X dyanmic grant versus aperiodic scheduling mechanisms for the ETSI application model.

services. In this section, it is investigated whether NR-V2X dynamic grant can demonstrate improved performance if a proportion of the application traffic exhibits periodic characteristics, thereby leading to a more accurate SB-SPS sensing window.

Fig. 5.9a considers 3 hybrid application models:

- **(90:10):** 90% of the traffic follows an aperiodic arrival rate in line with ETSI vehicle dynamics with the remaining 10% representing periodic traffic with a 10Hz transmission rate.
- **(50:50):** This application model generates a 50:50 aperiodic/periodic arrival pattern.
- **(10:90):** 10% of traffic follows an aperiodic arrival rate as per ETSI, with 90% arriving periodically.

It is clear in Fig. 5.9a that the performance of dynamic grant for hybrid application models is actually worse than the performance of NR-V2X SB-SPS when assuming disabled grant breaking. There is a 5% decline in performance in almost all ranges for ratios up to 50%. Fig. 5.9b highlights a noteworthy aspect with respect to performance. While the CBR increases by over 10% as the models incorporate more periodic arrival rates, the PDR performance does not change significantly. In fact, the PDR performance is best for the highest CBR when the ratio is 90% periodic. Thus, increasing the amount of periodic application traffic improves the performance of both these scheduling mechanisms, as it reduces the overall number of collisions and more effectively utilises the SB-SPS sensing history. For the hybrid model with the highest level of periodic characteristics (10:90), the difference in PDR between the NR-V2X SB-SPS and Dynamic Grant is negligible.

Ultimately, it can be concluded that NR-V2X dynamic grant is not adequate when scheduling highly aperiodic traffic patterns. This is very problematic as there is no guarantee of periodic traffic in the network. Current C-ITS services by their very nature are event driven i.e. aperiodic.

Aperiodic scheduling mechanisms such as counter-based scheduling and short-term reservations can be incorporated to work alongside SB-SPS for mixed traffic scenarios as well. In both cases, this requires that the aperiodic scheduling mechanisms account for SB-SPS grants when selecting a resource. To evaluate whether this approach has any negative impact on the default

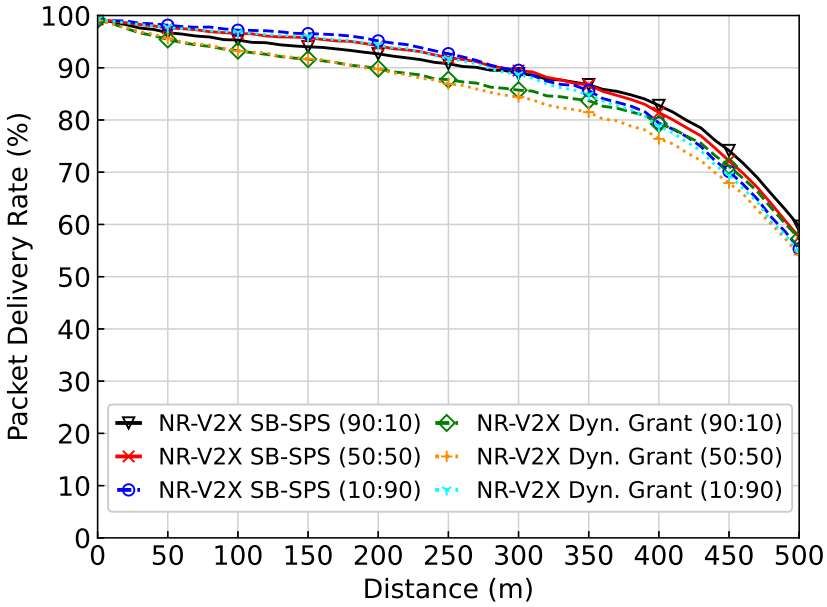

(a) NR-V2X PDR for SB-SPS & dynamic grant.

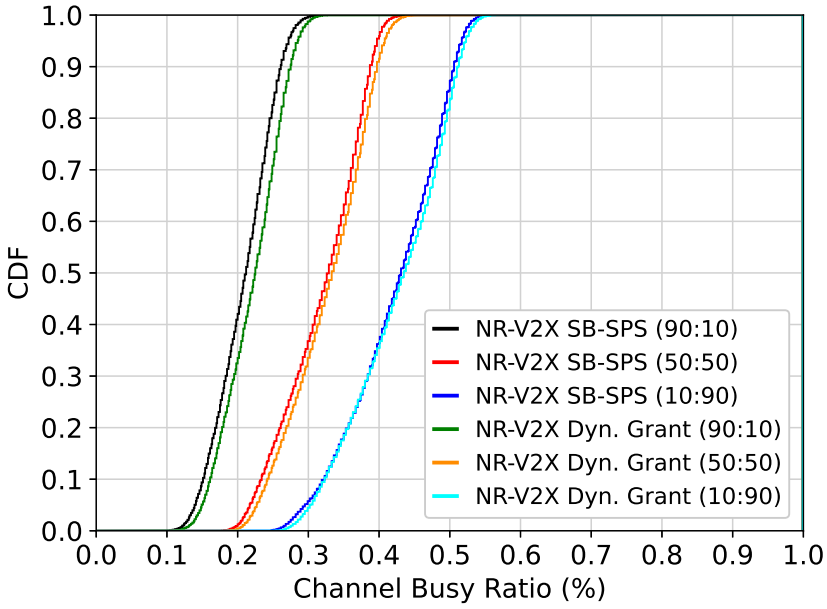

(b) NR-V2X CBR for SB-SPS & dynamic grant.

Figure 5.9: NR-V2X performance when considering hybrid application models (aperiodic:periodic).

SB-SPS scheduling, the following distributions are considered.

- **SB-SPS (Periodic):** Fixed periodic 100 ms.
- **(10:90):** 10% of the traffic follows an aperiodic arrival rate in line with ETSI vehicle dynamics rules with the remaining 90% representing periodic traffic with a transmission rate of 10Hz.
- **(25:75):** 25% of traffic follows an aperiodic arrival rate as per ETSI, with 75% arriving periodically.
- **(50:50):** This application model generates a 50:50 aperiodic/periodic arrival pattern.
- **Aperiodic 100%:** All traffic is generated according to ETSI CAM rules.

Figs. 5.10a and 5.10b show the performance of the counter based and STR mechanisms in these mixed traffic scenarios with the ETSI application model. It is clear that the mechanism works well in conjunction with the default SB-SPS algorithm, and neither traffic pattern significantly impacts the other. In this case, the increase in PDR is due to a lowering of CBR when aperiodic traffic is introduced. When analysing the STR it can be seen that it performs comparably with a purely periodic scenario. This demonstrates that for lower channel loads, the short-term reservation mechanism is beneficial to the performance of the aperiodic traffic. This is likely not to hold at increased densities as already described due to increased congestion in PSCCH.

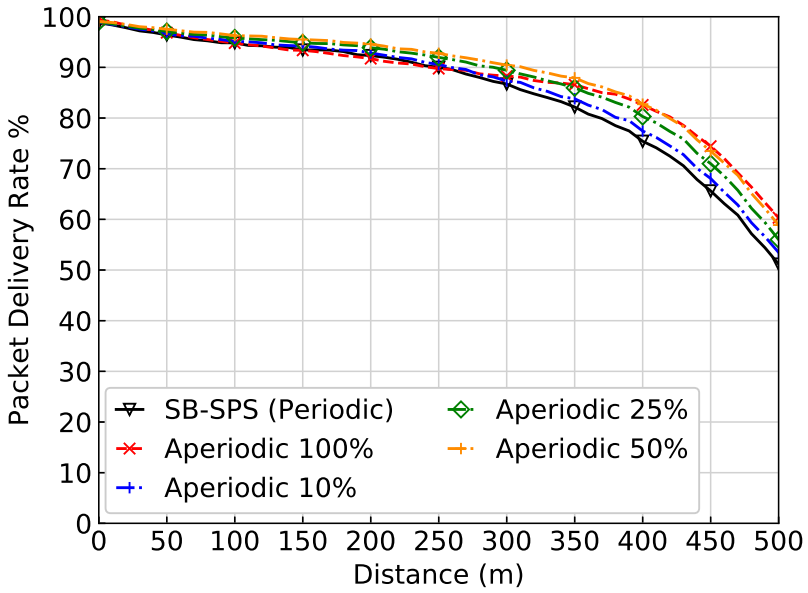

(a) Counter.

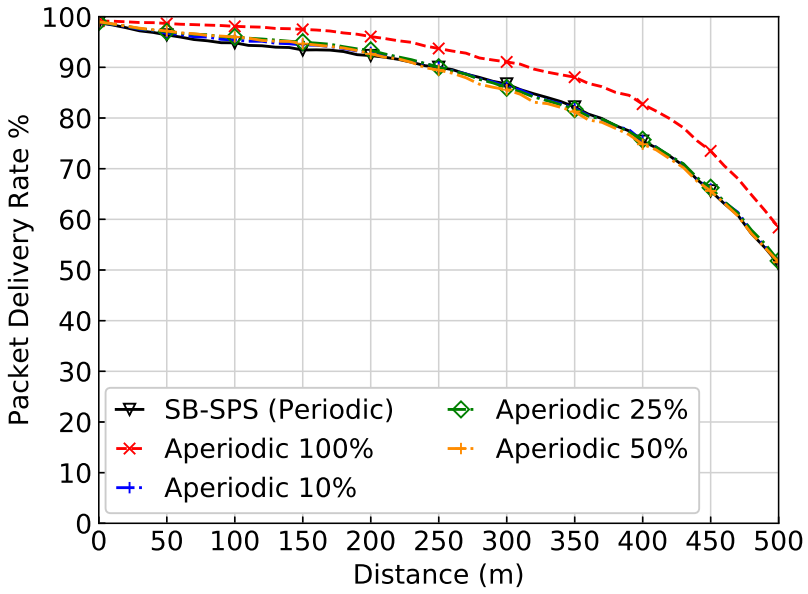

(b) STR.

Figure 5.10: Hybrid application models for the dedicated aperiodic scheduling mechanism (β=0.12 veh/m).

5.3.7 Conclusion on performance of current standardised approaches

In the previous section, the focus was on the aperiodic performance of SB-SPS on C-V2X and NR-V2X. This was a comprehensive evaluation of the primary challenges faced in SB-SPS. With a particular focus on both GB and no-GB configurations, different application traffic patterns, the tuning of SB-SPS parameterisation, alternative scheduling mechanisms outside of SB-SPS and finally a focus on NR-V2X dynamic grant with emphasis on hybrid traffic scenarios. The primary conclusions of this segment are as follows:

- Focusing on the impact of different application traffic patterns, the rate and variability of the pattern is a key concern, as was shown with the 3GPP pattern, which exhibits a rate rarely higher than 200 ms, the SB-SPS mechanism can be effective. While for ETSI which can exceed the *2n-2* rule, the performance of SB-SPS degrades significantly.
- The analysis of SB-SPS management further highlighted the significant performance impact of the *sl-reselectAfter* parameter, with grant breaking and non grant breaking having a major bearing on scheduling performance. For both C-V2X and NR-V2X not grant breaking is preferable, interestingly, NR-V2X is not as severely impacted by a grant breaking configuration due to the changes made in RSRP and RSSI filtering of its SB-SPS mechanism.
- In terms of parameter tuning, the primary outcome of this investigation showed some performance improvements from the default SB-SPS mechanism but ultimately that it is not a solution to the challenge of aperiodic traffic.
- The two alternative mechanisms investigated showed some improvements over the current SB-SPS performance when dealing with ETSI traffic patterns, especially STR with the caveat that increased PSCCH congestion leads to performance degradation when density increases as a result of additional signalling overhead.
- An investigation of the proposed Dynamic Grant mechanism of NR-V2X, which has been introduced to deal specifically with these aperiodic patterns shows that this mechanism does not perform effectively with the lack of any periodic traffic. This is due to its reliance on the underlying SB-SPS mechanism. However, it can be effective when used in a mixed

traffic scenario, where the majority of traffic is periodic in nature and only a smaller proportion of the channel is dedicated to purely aperiodic traffic.

The final conclusion with regards to dynamic grant performance motivated the analysis in the next section. The aim being a scenario underwhich a high level of traffic is periodic in nature and dynamic grant can be effectively used for the remaining aperiodic traffic. This is a possible scenario when considering both the DEN and the CA service, DEN being purely aperiodic in nature as a result of it being event-driven and forming a small proportion of the traffic, while CA is more prevalent. This motivated the proposal to use a machine learning-based approach to enable the prediction of CAM IPT as this would enable the core traffic from CA to appear periodic in nature, thus enabling DEN traffic to be scheduled with dynamic grant more effectively.

5.4 CAM IPT Prediction

Section 5.3 established that:

- Default SB-SPS performs poorly with aperiodic traffic, regardless of how its parameters are tuned.
- Dedicated schemes offer some minor performance benefits but at the cost of much higher control overhead.
- Dynamic grant, proposed in NR-V2X will only work if there is periodic traffic and an accurate sensing history, which cannot be guaranteed.

As such, this section quantitatively evaluates the performance of the CAM prediction models described in Chapter 4 Section 4.3.3 when applied to predict the inter-packet time and use this to determine the RRI for each transmission. The goal of this approach is to achieve a scenario where CAMs mimic a periodic arrival rate when the underlying mechanism is, in fact, aperiodic in nature. The proposed approach adopts a GRU based machine learning approach which enables the vehicle to predict future heading and speed to determine when the CA trigger rules might be broken in the subsequent second, translating this then to a RRI which is sent in the current CAMs associated SCI message.

5.4.1 Simulation scenarios and models.

This section evaluates the performance of the CA prediction model and compares it with 3 different models.

- First, a periodic model that is configured to exhibit the same average packet inter-arrival rate as the ETSI triggered traffic. This represents the best-case performance of SB-SPS, which the ML model should tend toward when it is operating most effectively. This acts as a benchmark for the accuracy and usefulness of the proposed approach and is labelled *Periodic* throughout the analysis.
- The second model is the default operation of SB-SPS (grant breaking is disabled), where if no packet is available for transmission, the grant is maintained, but a missed transmission occurs. This is labelled *SB-SPS (No-GB)*.
- The third model is used as a baseline where the RRI is configured to be the average packet inter-arrival rate (rounded to the closest RRI). This is labelled *Mean IPT* in subsequent discussions.
- Finally, the proposed CA prediction model is labelled *Predicted IPT*.

The simulation parameters are according to Table 5.1, with the exception that mobility is changed as shown in Fig. 5.11 and that the SUMO step-length is increased to 10 ms, this is to align with the SUMO scenarios from which the model was generated. The final significant change is that the CAM triggering conditions are checked every 100 ms, this results in CAM transmissions being in regular 100 ms intervals, this is configurable in the standard but eases the prediction problem considerably.

To establish the performance of the CA prediction model in diverse vehicular scenarios, three scenarios are considered;

1. A simple highway scenario as shown in Fig. 5.11a. In this scenario trigger conditions occur exclusively due to position changes, as vehicles have fixed speed.
2. The second scenario is a curved highway, as shown in Fig. 5.11b. This evaluates the performance of the model when triggering conditions occur due to both positional and heading changes.
3. The final scenario is based on a real road network in Ireland that does

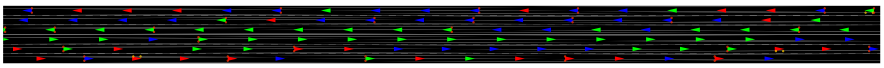

(a) Straight Highway.

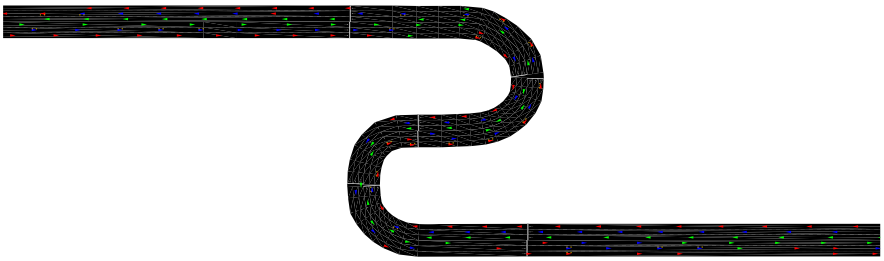

(b) Curved Highway.

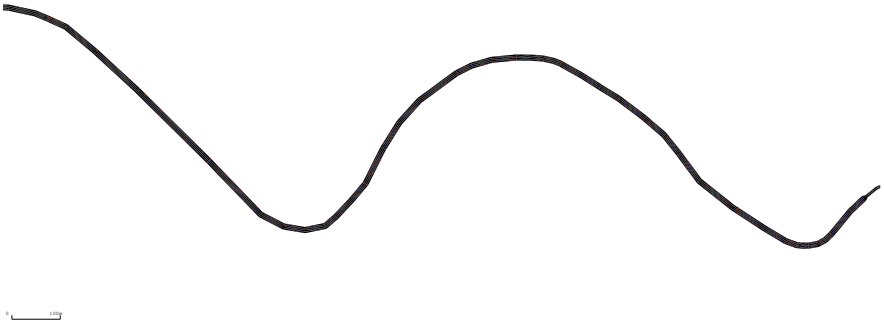

(c) Real road network segment.

Figure 5.11: Simulated road network topologies for CAM IPT Prediction.

not occur in our data sets. This scenario is to show how the model can generalise to other real world environments and is shown in Fig. 5.11c. Generalisation shows that the model is not over-fitted onto training data and can predict on unseen examples.

5.4.2 Performance in a straight highway (fixed speeds).

The first scenario under investigation is a highway scenario, which shows the highest level of performance for the proposed solution *Predicted IPT*. It can clearly be seen from Fig. 5.12 that the predicted performance mimics the performance in a *Periodic* configuration and outperforms the *SB-SPS (No-GB)* and *Mean IPT* models by 5%. Table 5.3 describes important aspects of each models performance, including application and transmission rates, RRI prediction performance and the underlying trajectory predictions. All configurations perform identically in terms of CAM rate (trigger rate) and IPT. The δCol field explains the performance differences, which represents the number of transmissions experiencing interference from neighbouring vehicles (within 500m). It can be observed that *Periodic* and *Predicted IPT* have significantly fewer δCols than the *SB-SPS (No-GB)* scenario, with *Mean IPT* resulting in slightly increased collisions. *Predicted IPT* performs well due to its ability to accurately predict the RRI, thus maintaining the SB-SPS grant mechanism more effectively. The *Predicted IPT* only exhibits 3 errors in prediction in the highway scenario, maintaining a low overall RRI Error. This is the result of a highly accurate prediction with an error of only 0.2 for the trajectory predictions.

5.4.3 Performance in a curved highway (varying headings).

In Fig. 5.13, it can be observed that both *Mean-IPT* and *Predicted-IPT* exhibit similar PDR performance, close to that of *Periodic* while significantly outperforming *SB-SPS (No-GB)*. In Table 5.4 it can be seen that *Periodic* exhibits higher transmission rates while maintaining good performance. *SB-SPS (No-GB)* also exhibits higher transmission rates, which contributes to its poorer performance, highlighted further in its increased δCol value. The *Mean* and *Predicted IPT* models exhibit similar performance in all respects as shown in Table 5.4, except for *Predicted IPT* having a lower number of collisions, which is not reflected in improved PDR. This is due to the low density nature of the scenario. As it exhibits fewer errors in predicting the RRI.

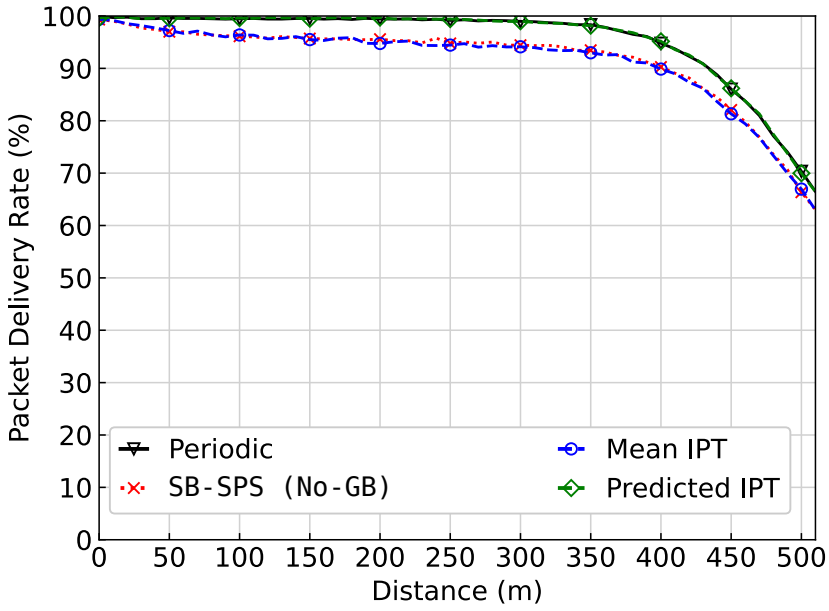

Figure 5.12: Performance of ETSI traffic model in a highway scenario (fixed vehicle speed).

Table 5.3: Application & RRI Performance (Highway)

Scenario	**CAM rate (ms)**	**IPT (ms)**	**δCol**
Periodic	301 ± 14	300 ± 16	99
SB-SPS (No-GB)	300 ± 15	300 ± 25	298
Mean IPT	301 ± 21	302 ± 32	321
Predicted IPT	300 ± 13	300 ± 15	98

Scenario	**RRI Error (ms)**	**Predictions**	**Inaccuracies**
Mean IPT	11.613 ± 81.244	3308	163
Predicted IPT	0.172 ± 5.426	3184	3

Prediction Error	**Heading °**	**Abs. Speed (m/s)**
	0.2 ± 0.2	0.2 ± 0.1

As such, both show similar accuracy in their predictions of the RRI though *Predicted IPT* reduces the number of inaccuracies by over 1000. It should be noted that the performance is not as effective as the highway; it can be seen that the error in terms of speed does not change significantly from the highway, but heading error is much increased at 6.9°. This increase in heading error directly results in more inaccurate predictions, reducing the performance of the *Predicted IPT* configuration. As is a challenge with all machine learning approaches, the primary causes for inaccuracies are somewhat difficult to diagnose because of their opaque nature. Theorising on possible challenges, the primary one is likely that the data used to develop the model does not have many instances with such extreme dynamics. This will result in the model itself having a more difficult time in predicting the outcome as the scenario will not match any of the training data. Otherwise, it is possible the model is not complex enough i.e. too few layers, neurons, and has difficulty in modelling the relationship between the input data and the expected outcome.

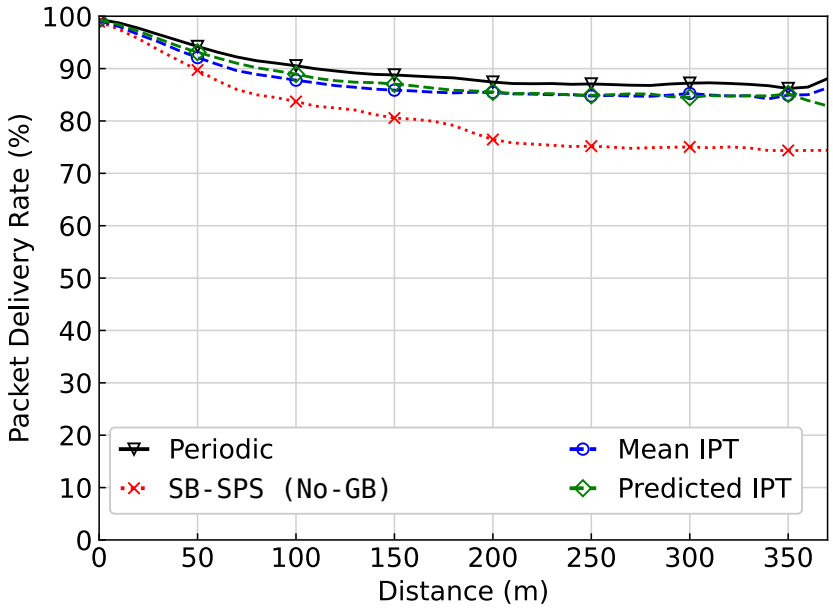

Figure 5.13: Performance of ETSI traffic model in a curved highway (varying vehicle headings).

5.4.4 Performance in an unseen road network (varying speed & heading).

As can be seen in Fig. 5.14a the *Predicted IPT* results increases PDR significantly from the *SB-SPS (No-GB)* and outperforms *Mean IPT* as well. As

Table 5.4: Application & RRI Performance (Curved)

Scenario	CAM rate (ms)	IPT (ms)	δCol
Periodic	266 ± 50	266 ± 51	1862
SB-SPS (No-GB)	299 ± 130	300 ± 131	2551
Mean IPT	349 ± 110	363 ± 100	1920
Predicted IPT	350 ± 111	359 ± 107	1827

Scenario	RRI Error (ms)	Predictions	Inaccuracies
Mean IPT	28.754 ± 126.015	9398	4340
Predicted IPT	22.840 ± 139.954	9354	3279

Prediction Error	Heading °	Abs. Speed (m/s)
	6.9 ± 8.7	0.3 ± 0.4

shown in Table 5.5 there are some differences in terms of IPT for the configurations with *Periodic* exhibiting higher transmission rates, along with *Mean IPT*. Ultimately, though, it can be seen that *Predicted IPT* significantly reduces the number of collisions and closes in on the *Periodic* performance. In this scenario again we see *Predicted IPT* exhibiting lower RRI Error and reducing the number of inaccurate predictions, which explains the performance increase compared to *Mean IPT*. When going through the trajectory predictions it is clear that speed predictions are highly accurate and heading error is lower, although it can still be high for this scenario at 3.1. How this error impacts performance is further depicted in Fig. 5.14b, this shows a confusion matrix for predicted IPT versus actual IPT as a percentage of all predictions. It can be seen that most of the IPTs are concentrated around 200 and 300 ms; overall, the model predicts accurately 88.2% of the time. The primary cause of inaccuracy is due to heading errors in prediction, either underestimating the heading change (100ms IPT predicted as 200ms) or overestimating it (200ms IPT predicted as 100ms). This was highlighted as the error in the heading can exceed the 4 °threshold for CAM triggering.

Ultimately, it can be concluded that the use of the predictive model represents a significant improvement from the current standard. The Cork road network also highlights the models ability to generalise to unseen data, which is key when dealing with machine learning approaches and highlights that the model is not fitted for a single dataset.

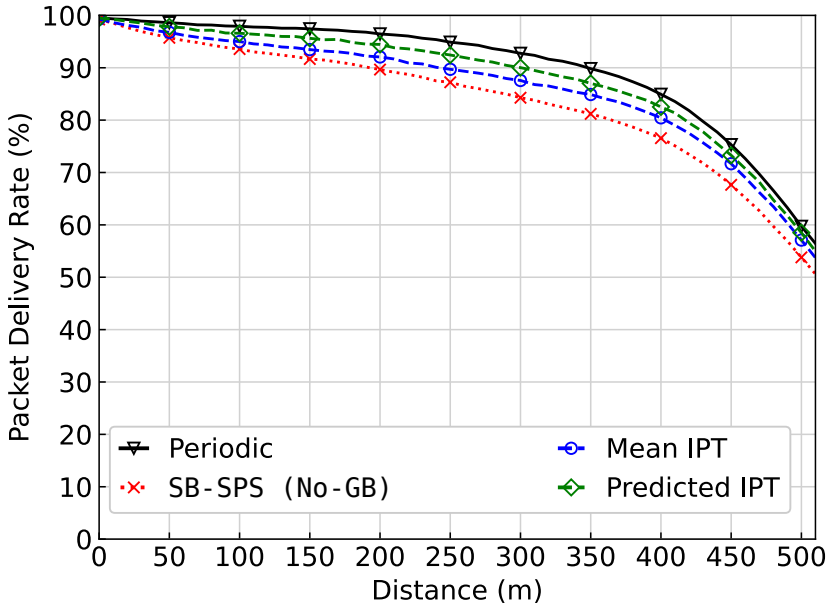

(a) PDR performance

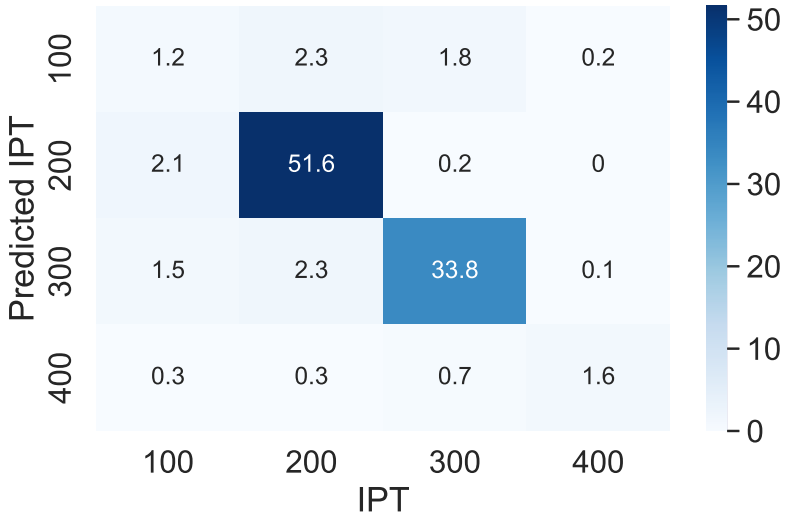

(b) RRI confusion matrix (overall percentage).

Figure 5.14: Performance of ETSI traffic model in unseen real road networking segment (varying speed and heading).

Table 5.5: Application & RRI Performance (Cork)

Scenario	CAM rate (ms)	IPT (ms)	δCol
Periodic	213 ± 36	213 ± 37	813
SB-SPS (No-GB)	238 ± 61	244 ± 64	1836
Mean IPT	227 ± 67	227 ± 68	1441
Predicted IPT	239 ± 60	242 ± 59	1089

Scenario	RRI Error (ms)	Predictions	Inaccuracies
Mean IPT	9.962 ± 58.945	18749	2463
Predicted IPT	2.723 ± 51.731	18988	2285

Prediction Error	Heading °	Abs. Speed (m/s)
	3.1 ± 5.0	0.3 ± 0.4

5.4.5 A comparison with the state of the art

The most directly applicable research that combats the challenge of CAM IPT in SB-SPS is from Lusvarghi et al. [71]. They adopt a similar approach of predicting the CAM IPT through a K-Nearest Neighbour (KNN) model. It is worth considering this study in more detail and offering a fuller comparison with the current proposed solution to highlight the advantages and disadvantages of each approach and prompt other readers on which approach is preferable or should be explored further.

There are a few key differences between the two approaches that can be examined from a number of perspectives. The first perspective is their conformance to the standard SB-SPS. The second is in the machine learning approach adopted including their advantages and disadvantages of each, as well as the means by which training was completed. Each of these perspectives are discussed in the subsequent sub-sections.

5.4.5.1 SB-SPS standards compliance

In the solution proposed in this thesis, the SB-SPS algorithm is deliberately kept completely compliant with the standard. The solution proposed by Lusvarghi et. al. diverges from standardised SB-SPS behaviour in one core element i.e. it includes the RC within each transmitted SCI message. This is motivated by the proposed solution which predicts both the RRI of the next CAM but also the sequence of CAMs with the same RRI. This prediction is then used to set the RC rather than the default.

The advantage of the Lusvarghi solution by introducing this element to the SCI message. It allows vehicles to look ahead at the RRI and the RC can predict subsequent collisions on selected resources with neighbouring vehicles. This will ultimately reduce the number of collisions as long as the predicted RRI and RC are accurate and avoids grant breaking which would likely incur the numerous challenges already highlighted in this work. The primary disadvantage is that by diverging from the standard deployed C-V2X equipment would need to be updated to a new standard and interoperability between versions of SB-SPS. This motivated the decisions for the research proposed in this thesis to ensure compliance with the standard such that any proposed solutions could be incorporated into deployments without the need for fundamental changes to the scheduling approach.

5.4.5.2 The proposed machine learning models

In this section the proposed approached is compared against Lusvarghi et. al from the perspective of the machine learning approaches adopted, their means of training and finally how the predictions are utilised.

The solution proposed in this study has already been described in detail in Chapter 4. It was determined that a GRU based deep learning model is the most effective in trajectory based challenges and hence was the selected approach. In contrast Lusvarghi et. al. adopt a KNN based approach. This machine learning approach compares the prediction data to all the sequences of training data and uses the most similar data to infer the possible future values. Such an approach is often used in regression and classification problems.

KNN models are generally easier to implement and not quite as costly in terms of computation as they generally use simple distance functions for the comparison between prediction data and the training data. In contrast a GRU requires more extensive training time and can be more computationally complex. A key advantage of a GRU based approach is that it maintains the sequence of the input data and as such better accounts for time series data. This is ignored in a KNN model which does place importance on the sequence of the input data. One major disadvantage of KNN models is their cost in terms of memory, as it requires all training data to be used when comparing to new prediction data. This is costly both in terms of prediction time but also in terms of the final model. To have a diverse and comprehensive model, a large

data set will be required which may become quite large and prohibitive to use in an OBU, making the solution impractical. The alternative is a number of smaller KNN models for different scenarios but then this poses problems for identifying when to use which model. The advantage of a GRU model is that while training times are more costly, the final model is not as costly in terms of memory, with only the prediction data and the pre-trained weights required for any prediction. This model can be trained on a diverse set of scenarios as shown in this study and that provides a single model which can generalise to a number of new scenarios without the cost in memory.

Other significant differences exist. Lusvarghi et. al. uses the KNN model to both predict the next CAM IPT to set the RRI but also the subsequent series of CAMs with the same IPT which is used to set the RC. This has the advantage mentioned previously that vehicles can predict subsequent collisions and should improve overall performance. This is not entirely dissimilar from the approach proposed in this thesis with the caveat that a single prediction for each CAM transmission is provided as opposed to multiple predictions. Theroetically, the approach outlined by Lusvarghi et. al is superior with the drawback of diverging from the standard as already highlighted. The key challenge however is in the practicality of the solution. Ultimately for each CAM transmission at the beginning of a new grant a number of predictions need to be made before transmission. In their study they highlight the accuracy up to 10 CAMs in the future, this represents 10 predictions such that the RC can be set which need to be completed before the transmission of the packet. The inference time is not indicated in their study but providing up to 10 predictions in the space of 100ms (the maximum time selection window employed in their study), could prove challenging to achieve in practice. Additionally, should the prediction prove inaccurate this would be a significant loss in terms of computation time for limited benefit. The solution proposed in this thesis provides a single prediction for each CAM which is means the proposed solution or an adapted approach is more likely to be made practical by designing the model with the constraint that inference can occur in less than 100ms.

The approaches also differ in the way by which training occurs. The key factor in any proposed ML solution is the data used in the training as this will influence the ability of the model to be used in a general set of scenarios. In Lusvarghi et. al. their study adopts a single simulation based imported from Modena and simulated for 20 minutes, with some 6400 vehicles and their

associated trajectory and CAM information. This in comparison to our proposed approach which uses a number of different scenarios to ensure that our model will generalise to unseen road networks and traffic conditions. It is likely that the Lusvarghi approach would not be capable of managing a diverse range of scenarios given that it was implemented on a specific road network $6km^2$ road network. Finally, Lusvarghis study has significant draw backs in that results are based on the same vehicular scenario which was used in training. This means that the results cannot be proven to generalise to unseen data and makes it a less practical approach to adopt in a real deployment.

5.5 MCS adaptation for variable packet sizes

This final section of the chapter focuses on the adaptation of MCS for the management of packets of variable size to minimised GBs. This is the basis of **contribution 5** which is a minor contribution of this thesis and is associated with *Pub 2*. With this aim, a number of aspects are considered:

- We provide a quantitative analysis of the performance of both a fixed and adaptive MCS approach, highlighting the shortcomings of assuming a fixed MCS.
- A thorough analysis is made of extending the MCS range beyond the limit of MCS 11 set out in the ETSI standard [52].
- Finally, different grant allocation policies are explored with the goal of reducing or eliminating GBs at the cost of increased subchannel occupation.

In this case again the focus is on CAMs and the parameters used in this study again correspond primarily to Table 3.1 with the following Table 5.6 showing the parameters changed for this aspect of the study.

It is important to note in this set of scenarios the density is kept fixed and different densities have not been investigated in this case. This work sought to determine the impacts of different application packet size models and how MCS configuration might impact performance including the available range and whether MCS could be adapted within the grant. The impact of different densities in such scenarios will be the same across the different configurations. This specifically meaning that in low density scenarios all configurations will have higher PDR, while in high density scenarios the PDR will be lower and in

Table 5.6: Scenario Parameters

Parameter	value
Vehicular Scenario	
Vehicular density (β)	0.2veh/m
Vehicle Speed	40km/h
Channel settings	
No. subchannels, Subchannel size	5, 10 Resource Blocks (RBs)
Application layer	
Packet Size	Dictated by model

all configurations the increase or decrease will be of similar percentages. This is as a result of the fact that density will only impact on the availability of subchannels for selection, it does not impact on the number of grant breaks as a result of larger packets arriving or on how efficiently different MCS settings can schedule the different packet sizes. As such, this analysis for the sake of clarity focuses on a single density of 0.2veh/m.

5.5.1 Modelling of packet size variability

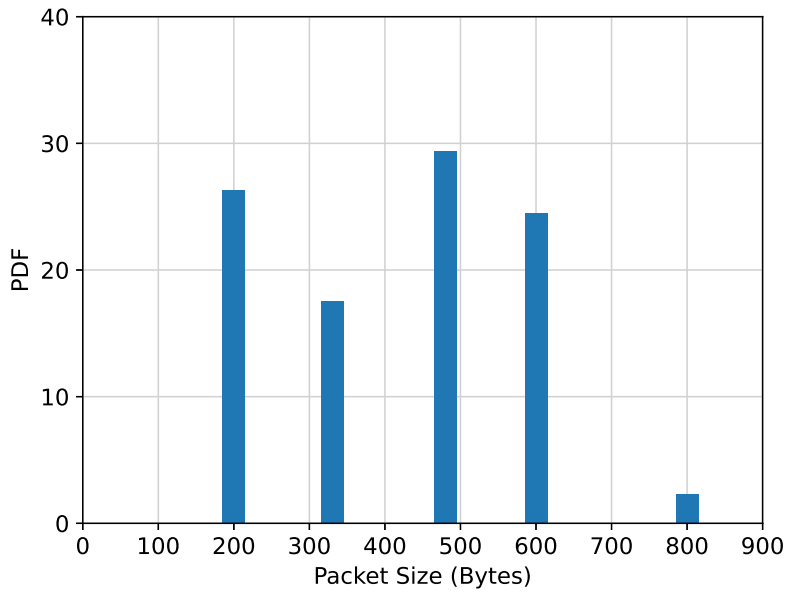

(a) Renault.

Figure 5.15: Distribution of empirical CAM packet sizes for the highway scenario.

Packet sizes are generated by the analytical model provided by

Molina-Masegosa et al. [59], which is based on empirical data gathered from two vehicular manufacturers (Volkswagen & Renault) across three environments (Highway, Urban and Hybrid) [98]. As the study in this paper investigates the impact of MCS adaptation on the scheduling of packets of variable size, a fixed transmission rate of 10Hz is considered to avoid the negative impact of aperiodic transmissions on the scheduler, as studied in [106]. Fig. 5.15 and Fig. 4.8 show the distribution of packet sizes across both vehicle models. For Volkswagen, all generated CAMs range from 200-455 Bytes in size, with more than 85% under 400 Bytes. The Renault model experiences a significantly larger distribution in packet sizes, with packets of up to 800 Bytes occurring. As discussed in [98], CAM sizes can vary depending on the number of optional fields configured and the size of the payload. We also consider a fixed packet size of 190 Bytes in the performance evaluation, which is a commonly assumed packet size in the related literature.

ETSI recommends a default configuration of 5 subchannels, composed of 10 RBs each, assuming a total channel bandwidth of 10 MHz [2]. Given this configuration and the 8 packet sizes collected in the empirical Volkswagen and Renault datasets, Table 5.7 shows the minimum subchannel occupation, the default MCS chosen, and the maximum byte capacity for all the packet sizes that occur in the empirical datasets. It is clear that a significant number of subchannels are required for the larger packet sizes, e.g. 4 subchannels are necessary for packet sizes of 600 or 800 Bytes. All packets for the Volkswagen model can be transmitted using 3 subchannels or less.

Table 5.7: Default MCS & subchannel allocation for the Volkswagen & Renault empirical data sets.

Vehicle Model	**Pkt Size (B)**	**Subchannel Occupation**	**Default MCS**	**Max Capacity (B)**
Fixed	190	1	11	201
Volk & Ren	200	1	11	201
Volkswagen	300	2	8	310
Renault	330	2	9	350
Volkswagen	360	2	10	390
Volkswagen	455	3	8	483
Renault	480	3	8	483
Renault	600	4	8	656
Renault	800	4	10	824

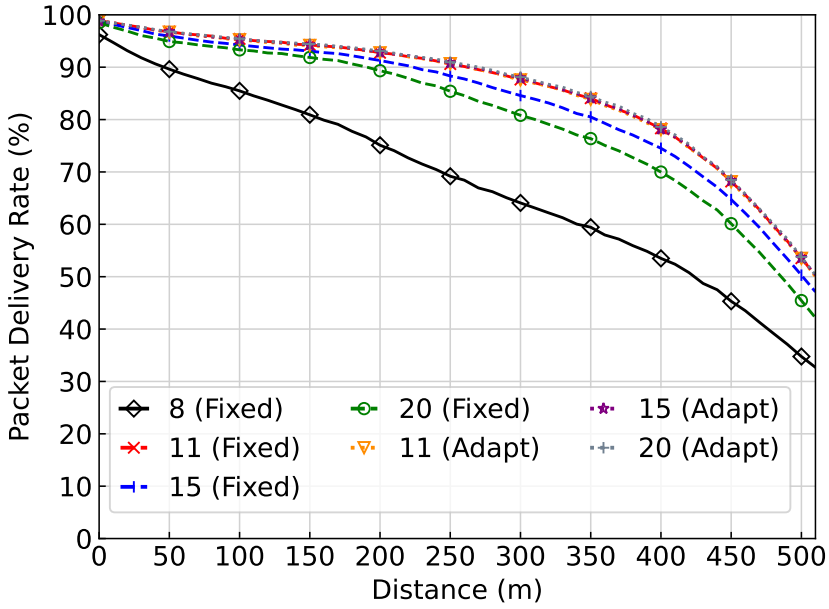

(a) PDR.

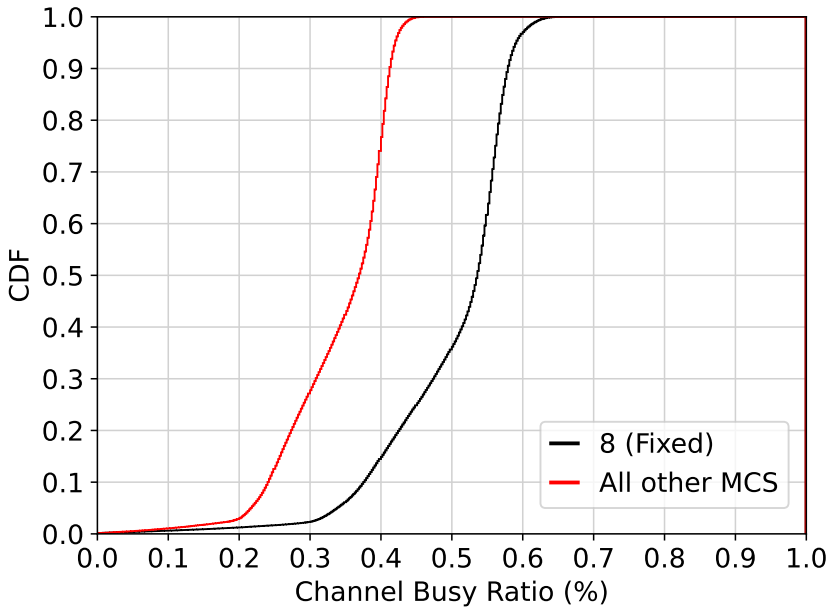

(b) CBR CDF.

Figure 5.16: Scheduling performance of various MCS selection schemes when considering a fixed packet size of 190 B.

5.5.2 Adaptive versus non-adaptive MCS performance

Fig. 5.16a and Fig. 5.16b, set out to evaluate the importance of utilising the adaptive MCS scheme set out by cellular vehicular standards when conducting C-V2X and NR-V2X experiments. A fixed MCS is commonly assumed in related literature. To highlight the significance of sub-optimal MCS choices, a scenario with a 190 B fixed-size packet is initially considered. Based on this, it is known that MCS 11 is the optimal choice based on Tables 7.1.7.2.1-1 and 8.6.1-1 provided in [22]. This is because it allows a packet to be transmitted in a single subchannel. Fig. 5.16a shows that choosing MCS 11 (red dashed line) performs significantly better than configuring a sub-optimal MCS such as MCS 8 (black diamond line) which is used here as a comparison. The reason for this becomes obvious in Fig. 5.16b. MCS 8 exhibits a significantly higher CBR, as it requires two subchannels instead of the single subchannel that MCS 11 and above will occupy. However, it is also not simply a case of arbitrarily using a higher order MCS indiscriminately. While MCS' 15 and 20 (Fixed) also only occupy a single subchannel, the PDR declines as the MCS increases. This is because it cannot further reduce subchannel occupation, yet increasing MCS impacts the robustness of transmissions, resulting in increased packet loss. Swapping to an adaptive approach allows the system to choose the optimal MCS based on the packet size. In this case, as the packet size is fixed, the optimal MCS will always be MCS 11 and as such we can see that MCS 15 (Adaptive, purple star) and MCS 20 (Adaptive, grey bar) both exhibit the same high PDR as MCS 11 (Fixed) and MCS 11 (Adaptive). This is because it automatically conforms to the optimal MCS for the fixed packet size. Thus, this introduces an important tenet of the operation of standardised MCS adaptation; the MCS should be increased such that subchannel occupation is minimised as much as possible to improve channel utilisation while ensuring the MCS is kept as low as possible to not degrade the robustness of transmissions.

The second set of results, shown in Fig. 5.17, focus on the Volkswagen empirical data and aims to explore whether higher-order MCS should be considered beyond the maximum range (MCS 0 - MCS 11) set out by the ETSI standard [2]. Fig. 5.17 provides clear motivation for this, showing that the use of MCS up to 64QAM MCS 20 should be considered, further corroborating research by Harri et al. [46]. As shown in Fig. 5.17a, selecting MCS 20, adaptive or not, shows the most significant performance improvement. This is

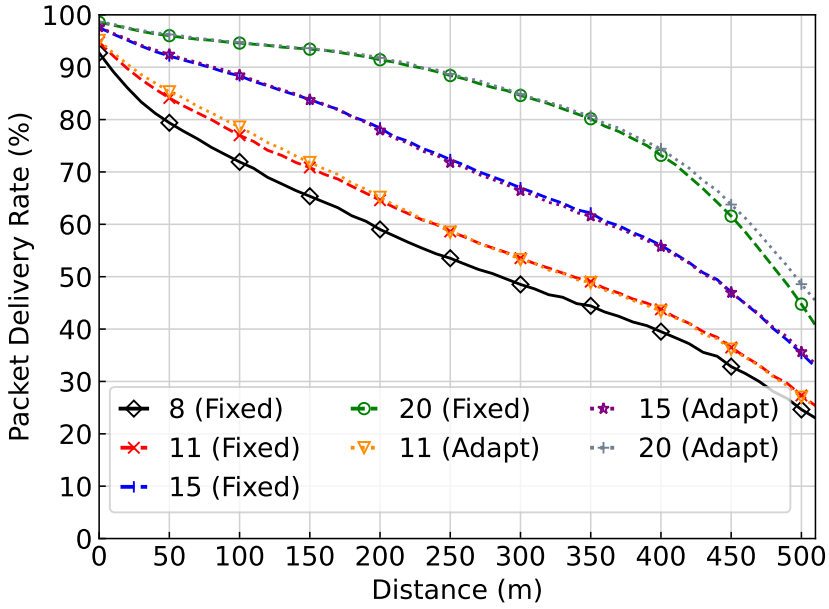

(a) PDR.

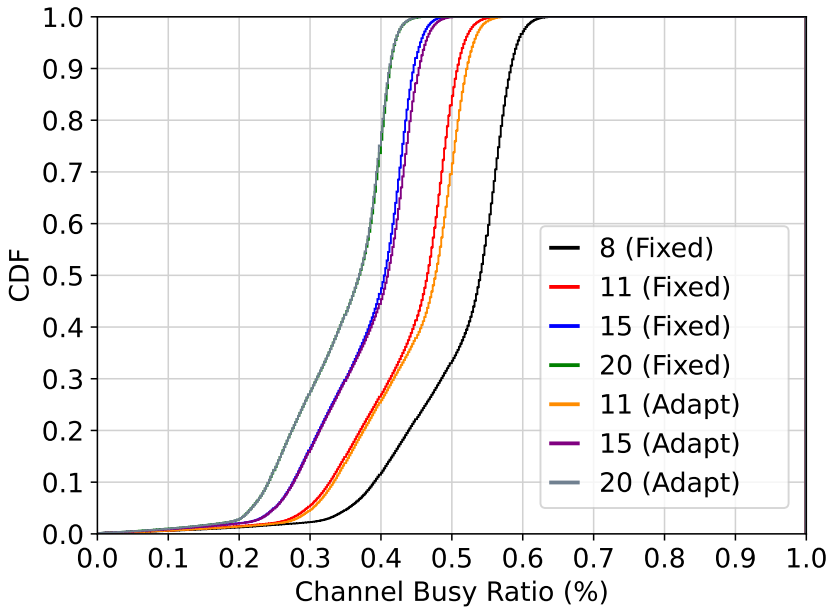

(b) CBR CDF.

Figure 5.17: Scheduling performance of various MCS selection schemes when considering the Volkswagen packet size distribution.

due to the reduction in collisions due to the reduced CBR, which is a direct result of reduced subchannel occupation, as shown in Fig. 5.17b. Even MCS 15, which is 16QAM, results in significant performance improvements.

Fig. 5.18 provides a more in-depth analysis of the performance of each MCS configuration with a focus on the reserved subchannels, the RBs used and the MCS selected given the adaptation for variable packet sizes. In Fig. 5.18a it is worth noting that MCS 8 and 11 (Adapt and Fixed) can use up to 3 subchannels, while MCS 15 will use a maximum of two, and finally MCS 20 will only require one subchannel to transmit each packet size. Fig. 5.18b shows a PDF of the RBs used for each transmission. It can be observed that MCS 20 (Adapt) will always use the full 10 RBs, while in the case of MCS 20 (Fixed), the RB usage can drop below 10. However, this RB reduction does not have a PDR benefit, as we observed in Fig. 5.17a, with a minor drop in performance at distances beyond 400 m as a result of the reduced transmission robustness. To reiterate, this signifies that there is little point in RB reduction if subchannel occupation cannot be further reduced. Fig. 5.18c shows a PDF of how often a particular MCS is used, given the range of MCS that could potentially be used for a particular packet size. Fixed MCS is not shown as no change in MCS will occur. It can be observed that an adaptive MCS scheme can exhibit significant variation in the MCS used, e.g., for MCS 20 (Adapt), the MCS used can range from MCS 10 to MCS 20 itself. This echoes earlier observations that performance can be significantly positively influenced by utilising a wider availability of MCS within an adaptive approach.

The empirical Renault data set, which exhibits much larger differences in the size of the CAM is considered in Fig. 5.19a. It is once again clear that MCS 11 (fixed) limits the performance of the model, and increasing this to MCS 15 allows for significant performance improvements. Furthermore, it can be observed in Fig. 5.19b how the CBR decreases in steps once the MCS is increased. In particular, adaptive MCS 20 is once again the best performing, but performance improvements are more marginal than for the Volkswagen dataset. This highlights the fact that packet sizes, the number of subchannels, and RB configuration have a significant impact on the performance benefits that these mechanisms can gain. Ultimately, the optimal MCS will depend on the specific configuration of the scenario. However, what is important is that enabling a high MCS such as MCS 20 and allowing adaptation within that wider range ultimately result in the best performance. As such, it is paramount that within any study assuming variable packet sizes, common to almost all

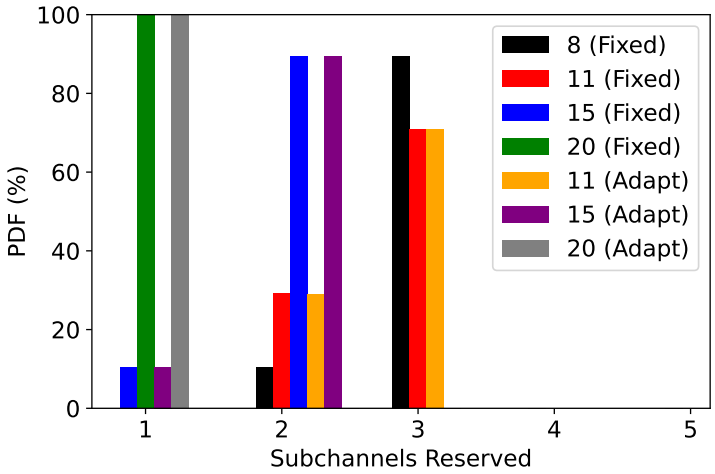

(a) Subchannels reserved.

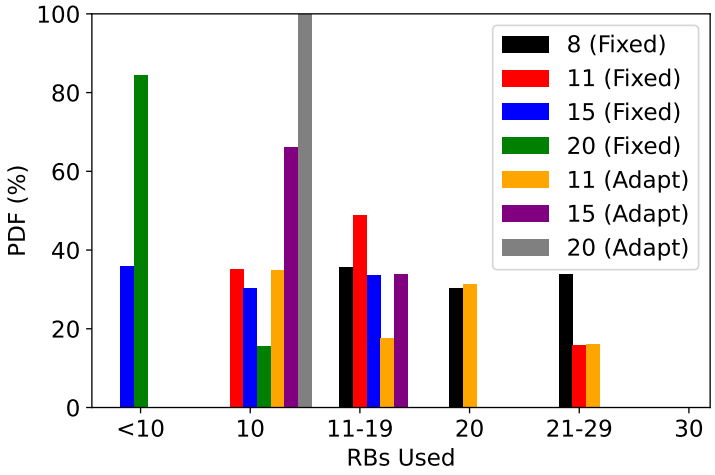

(b) RBs used.

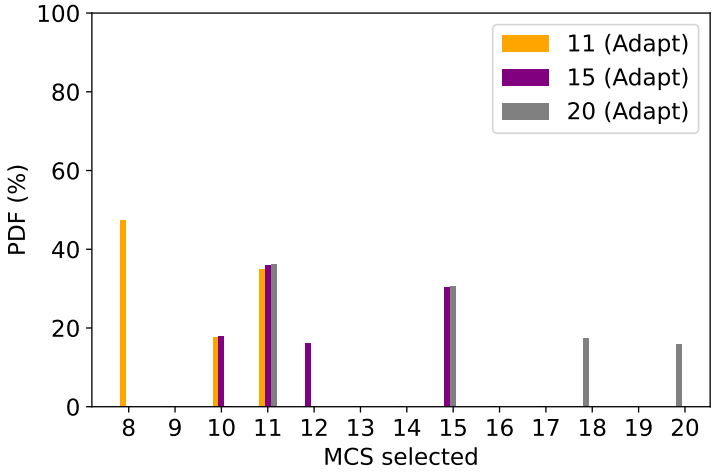

(c) MCS used.

Figure 5.18: Radio resource usage considering the Volkswagen packet size distribution.

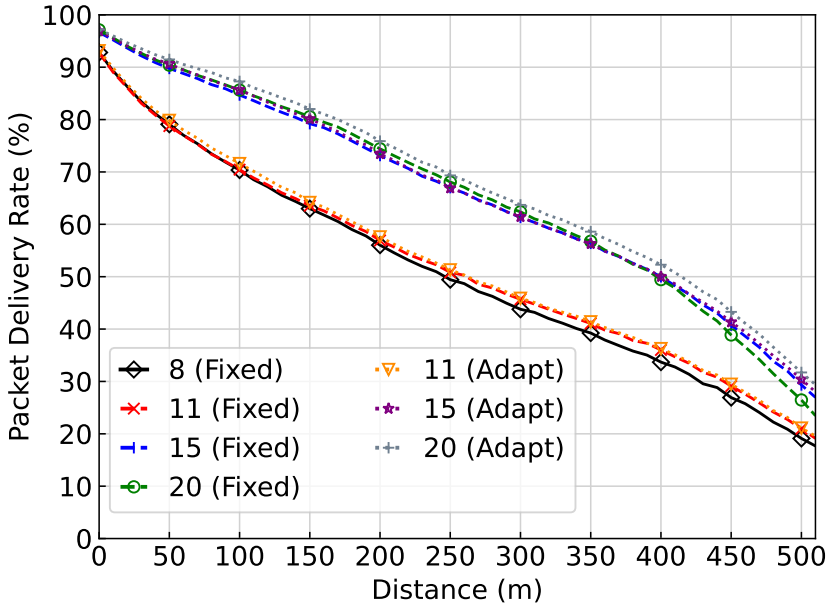

(a) PDR.

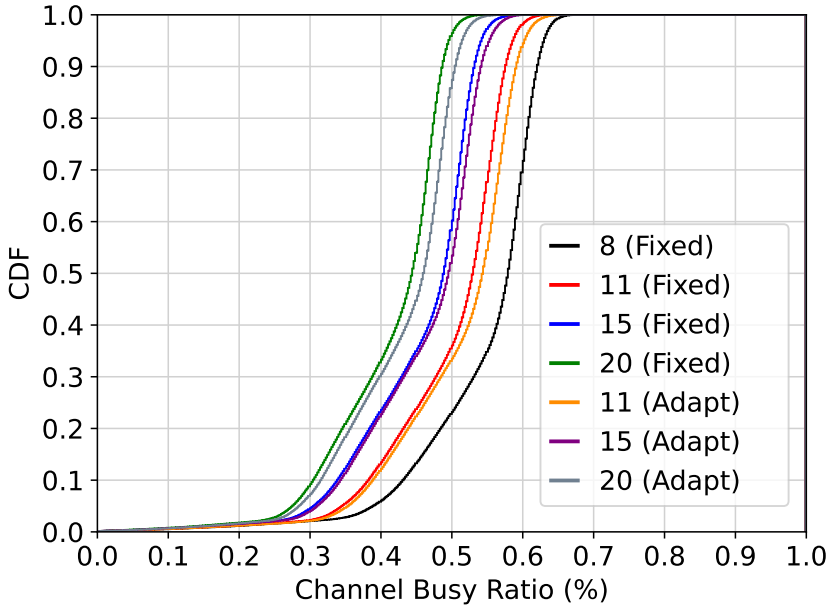

(b) CBR CDF.

Figure 5.19: Scheduling performance of various MCS selection schemes when considering the Renault packet size distribution.

realistic vehicular applications, the simulation model should implement MCS adaptation as per the C-V2X and NR-V2X standards.

5.5.3 Grant optimisation

As stated previously, the default adaptive MCS selection process in C-V2X and NR-V2X chooses an MCS to occupy as few subchannels as possible while ensuring the lowest possible MCS for robustness. The challenge is that packet sizes are variable and therefore can change over the course of an SB-SPS grant. If a packet is too large to transmit, the grant is broken and the grant selection process commences again.

This section seeks to investigate different MCS selection approaches to investigate the trade-off between the current default approach where the grant is determined based on the packet at generation time vs an approach that could seek to optimise MCS selection based on possible distribution of packet sizes incurred by a given application.

With this aim, the following approaches are evaluated:

- **SPS grant generation using the Default C-V2X/NR-V2X MCS Selection Scheme:** This is the approach adopted so far in the investigation and described in Section 2.3.6.
- **SPS grant generation using an MCS Selection Scheme to minimise grant breaking:** By examining typical distribution of packet sizes, an MCS will be selected based on the largest possible packet. In the case of Renault, this is an 800 byte packet. This approach will eliminate the need for grant reselections at the cost of increased subchannel occupation. This is labelled *(Max)* in Fig 5.20.
- **SPS grant generation using an MCS Selection Scheme based on mean packet size:** This selection scheme is a tradeoff between the first two where the MCS will be selected such that it fits most of the packet sizes that occur, with the exception of rare larger packets. This will balance the subchannels occupied while reducing the number of reselections. This is labelled *(Mean)* in Fig 5.20.

It is clear from Fig. 5.20 that adopting an approach other than the default MCS selection scheme is not beneficial. On the contrary, it can have a serious negative impact on performance even if it reduces the number of SB-SBS grant

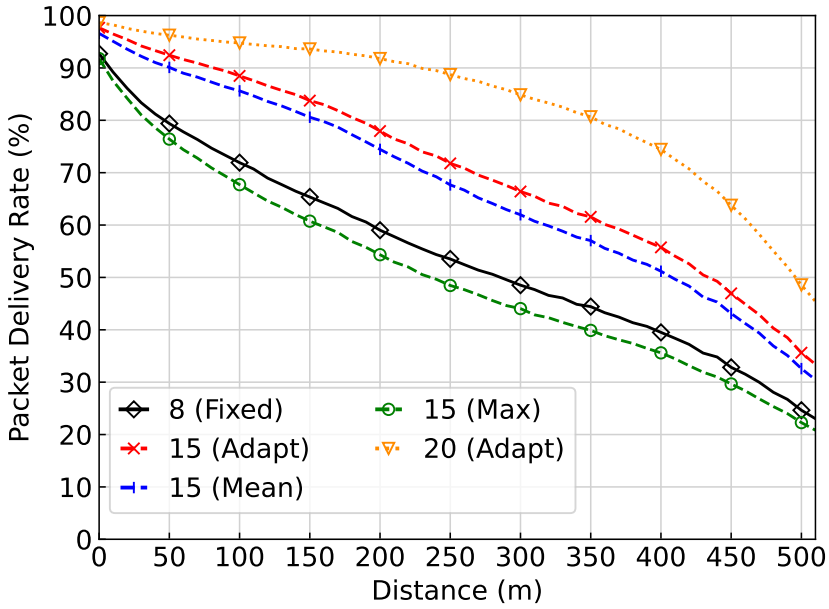

(a) Volkswagen.

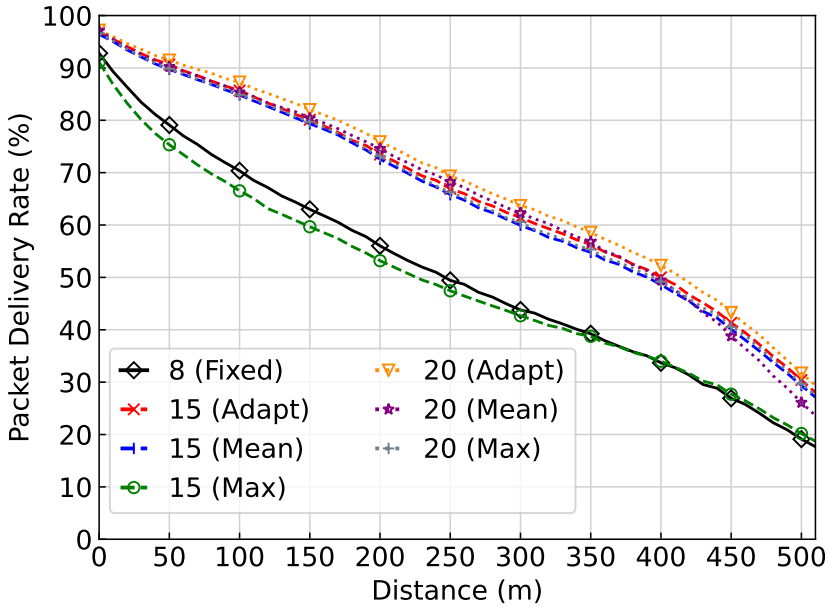

(b) Renault.

Figure 5.20: Performance based on alternative MCS Selection Strategies linked to packet size distributions i.e. deviating from the standard SB-SPS grant selection mechanism.

reselections. This is highlighted by the performance of the MCS 15 (Max) configuration for both Volkswagen and Renault. The negative impact of MCS 20 is less pronounced but ultimately still does not offer any benefit.

Table 5.8: Grant Breaking due to packet size

MCS	MCS adapt	Volkswagen	Renault
8	Fixed	2544	2715
15	Adapt	2519	1496
15	Mean	0	475
15	Max	0	0
20	Adapt	0	1669
20	Mean	-	1694
20	Max	-	0

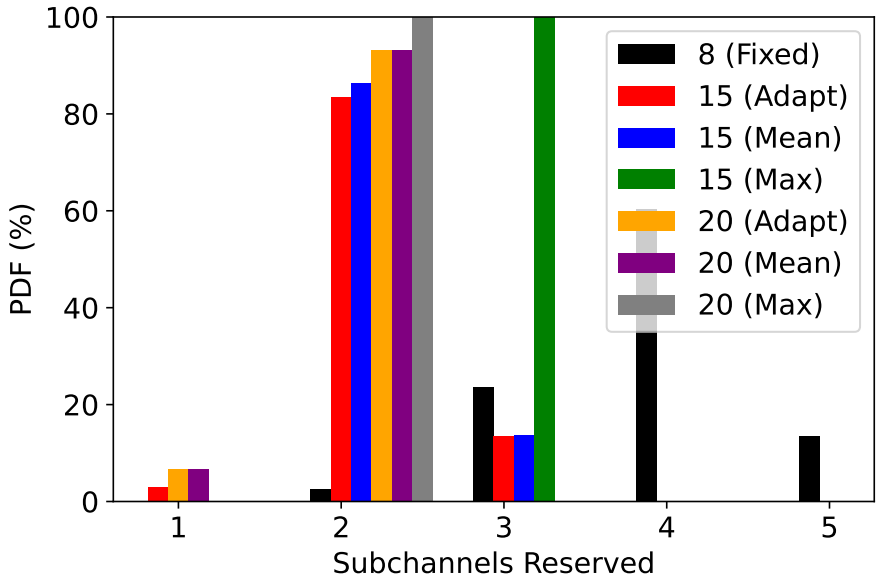

Figure 5.21: Subchannel Occupation (Renault)

This is somewhat unexpected, as it could be assumed that by reducing grant reselections, an improvement in PDR would be observed.

Table 5.8 shows that adopting the Max MCS selection approach eliminates all grant breaks due to size, eliminating the need to reschedule grants. Mean also reduces the number of rescheduling events, which generally results in improved performance. However, the explanation for this is shown in Fig. 5.21, where the occupation of the subchannel is compared for each scheme. In the case of MCS 15, adopting the Max approach results in 3 subchannels being used. With over half of the subchannels being reserved for all transmissions, this results in increased interference as there will always be an overlap in

subchannels if 2 or more vehicles transmit simultaneously. Clearly, while grant breaks can be reduced, the trade-off when increasing the subchannel occupation results in a degradation in performance greater than the benefit from increased grant duration.

5.6 Conclusion

This chapter included a full analysis of how the C-V2X and NR-V2X scheduling mechanism SB-SPS manages the current services. With particular focus placed on the CA and DEN services, which will form the core to any deployments. The highlights of this analysis are as follows:

- The initial element of this chapter provided an in-depth analysis of SB-SPS performance with aperiodic traffic patterns. The main outcome from this section was regardless of configurations, and parameterisation SB-SPS does not schedule these traffic patterns effectively.
- This study further focused on different aperiodic scheduling mechanisms and there are a few key findings in this area.
  - The dynamic grant mechanism proposed in NR-V2X is shown to not be adequate at managing fully aperiodic traffic.
  - Alternative scheduling approaches, are shown to perform similarly or better but with tradeoffs such as increased channel usage. As such, no solution is truly effective in managing high levels of aperiodic traffic.
  - Dynamic grant and other mechanisms can be effective in hybrid periodic and aperiodic traffic scenarios in managing low-levels of aperiodic traffic.
- This result motivated the adoption of a CAM IPT prediction mechanism to ensure accurate sensing history, which enables dynamic grant to manage aperiodic traffic more effectively.
- As shown, the proposed GRU based trajectory prediction mechanism can be used to accurately predict the CAM IPT and maintain the RRI mechanism more effectively.
- The final element of this chapter deals with variable packet size management and analysed two empirical models. The key conclusions

from this work are as follows:

- First, that adopting a higher order MCS which reduces subchannel occupation results in the best performance for SB-SPS due to reduced congestion even with reduced transmission robustness.
- Second, and often overlooked in other literature is that maintaining the adaptive mechanism is key to ensure the best scheduling performance. This avoids the pitfall of both using too low an MCS and incurring increased subchannel occupation or using too high an MCS which reduces transmission robustness unnecessarily.
- Lastly, with regard to over-provisioning mechanisms, which avoid grant breaking due to packet size changes. These were ultimately shown to be ineffective because the increased congestion from the use of multiple subchannels does not overcome scheduling degradation as a result of a grant break.

Chapter 6

Quantitative Performance Analysis: Decentralised Congestion Control (DCC) in C-V2X/NR-V2X, Proposed DCC Solutions and the role of MCS

6.1 Introduction

The previous chapter focused on how SB-SPS in both NR-V2X and C-V2X is able to operate with the current generation of V2X services, both CA and DEN. As highlighted, SB-SPS has challenges both when packets arrive in an aperiodic fashion and when packet sizes change dynamically. In congested scenarios, similar issues occur i.e. MTs and GBs, and as such, new approaches need to be explored that can overcome these challenges; specifically, in the case of TRC mechanisms, these will ultimately exhibit properties similar to those aperiodic arrival rates. The challenges faced by congestion control mechanisms are explored in more detail in Section 4.2.3.1. This chapter forms the quantitative analysis behind **contribution 4** (**C4**). This chapter will:

- Provide an in-depth quantitative investigation of the significant challenges that the C-V2X and NR-V2X scheduler faces when standardised congestion control mechanisms are employed. This is described in Section 6.3.
- Investigate the role of MCS congestion control as a technique in cellular vehicular networks. This is described in Section 6.4.

- Propose congestion control mechanisms that overcome these shortcomings while remaining compliant with the C-V2X and NR-V2X scheduling standards. These approaches will be evaluated for performance at the network and application layers but also with respect to fairness and stability. This is described in Section 6.5.

The investigation of existing standardised congestion control approaches with C-V2X and NR-V2X will focus on 3GPP based approaches, ie, CBR-to-CR-Limit tables used for packet dropping or TRC approaches [52]. How the table mechanism operates and the possible choices for congestion control technique are described in Section 2.3.7.1 of Chapter 2. An investigation of the TRC approaches of the ETSI standards will also be carried out, i.e., *DCC Adaptive* and *DCC Reactive* [107]. Ultimately, it will become clear that current standardised congestion control approaches cannot adequately provide improved performance, as they are incompatible with the C-V2X and NR-V2X scheduler.

Next, this chapter considers the role of MCS congestion control as an alternative technique to TRC approaches. There are four categories of mechanisms described in the standard, the first being TRC, then TPC, elimination of retransmissions, and finally adaptation of MCS [28]. The premise of adapting the MCS is to increase the MCS when congestion occurs, thereby reducing subchannel occupation and thus reducing congestion by freeing up radio resources. We find that this approach is not effective due to the limited availability of MCS in current standards, subchannel configuration, and the natural performance gains from always using the highest possible MCS.

Finally, to combat the incompatibility of existing approaches with the SB-SPS scheduler, new congestion control mechanisms are introduced that manage congestion while remaining compliant with C-V2X and NR-V2X. This contribution **[C4]** is a TRC mechanism based on the adaptation of the scheduler RRI parameter, depending on the number of transmissions that need to be skipped. Three variants of this approach are implemented:

- $\text{RRI}_{\text{CR\_Limit}}$: An approach compliant with the 3GPP table-based packet dropping mechanism.
- $\text{RRI}_{\text{Reactive}}$: An approach compliant with the ETSI DCC reactive table-based mechanism.

- $RRI_{Adaptive}$: An approach compliant with the ETSI DCC adaptive algorithmic approach which is a variant of the LIMERIC algorithm [25].

These three mechanisms will be shown to significantly improve on the standardised TRC based congestion control techniques.

6.2 Simulation scenario and models

The primary simulation parameters considered for this study are the same as defined in Table 3.1 of Chapter 3. However, to show the limitations of existing congestion control techniques and highlight the advantage of the proposed RRI techniques, we change certain parameters, shown in Table 6.1. In this case results are generated as highlighted in Section A, with the exception that simulations in this case are run for considerably longer 60s; this is to allow the DCC averaging mechanism used in *DCC Reactive* as well as $RRI_{Reactive}$ and RRI_{CR_Limit} to stabilise. Further, it should be noted in all cases the subsequent analysis will focus on non grant breaking (No-GB) as such *sl-reselectAfter* is set to greater than the RC meaning only MTs can occur when a packet is unavailable. This is as a result of what has been discussed thoroughly in the previous Chapter 5 namely that GB configurations result in degraded SB-SPS scheduling performance.

Table 6.1: Scenario Parameters for Sections 6.3 & 6.5

Parameter	Value
Vehicular scenario	
Vehicular density	0.46 veh/m
Road length	600 m
Number of lanes	3 in each direction (6 in total)
Vehicle Speed	50km/h
Vehicle Mobility	SUMO (step-length = 1ms)
Channel settings	
No. subchannels & size	3 & 16 RBs
MAC & PHY layer	
MCS	6 (QPSK 0.5)

For the analysis of MCS congestion control the selection is based on the ETSI standard for C-V2X [2] rather than the original 3GPP simulation scenario, as it provides more scope under which adaptation can occur, that is, 5 subchannels as opposed to 3. We consider the parameters shown in 6.1.

Table 6.2: Scenario Parameters for Section 6.4

Parameter	Value
Vehicular scenario	
Vehicular density	0.06, 0.09 and 0.20
Channel settings	
Number of subchannels	5
Subchannel size	10 RBs
MAC & PHY layer	
MCS	7, 11, adaptive
Modulation and Coding Rate	QPSK/0.7, 16QAM/0.5, adaptive
Transmission power	27.32 dBm, 25.97 dBm, adaptive

Three different vehicular densities are considered along with two fixed MCS configurations (that is, MCS 7 and MCS 11) and an adaptive configuration, where the MCS alternates between MCS 7 and MCS 11 depending on the congestion in the channel. Furthermore, to maintain the PSD constant and within the limit of 23 dBm/MHz, the transmission power was also adapted according to the selected MCS configuration.

6.3 Standardised DCC : Performance limitations in C-V2X & NR-V2X

This section conducts an in-depth evaluation of the performance of the current congestion control standards when directly applied to the C-V2X and NR-V2X networks and confirms the theorised limitations discussed in detail in Section 4.2.1 of Chapter 4. For clarity, the primary challenge is briefly reiterated: Existing standardised congestion control techniques are mainly based on TRC approaches. These TRC mechanisms will delay or drop packets to reduce congestion, which will induce a significant number of missed transmissions in SB-SPS, the MAC scheduler. Such missed transmissions can result in:

- Grant breaks which cause significant degradation in scheduling performance due to increased packet collisions, or
- In the case no grant breaking, increased collisions as a result of missing SCIs causing reserved resources to appear free.

Neither scenario represents ideal performance of SB-SPS, and as such this negatively impacts, and in some cases totally negates, the performance of the

congestion control mechanisms.

6.3.1 ETSI standardised packet delay (TRC) mechanisms

This section will seek to investigate the effectiveness of ETSI mechanisms that were initially envisaged to be applied to ITS-G5, but may be applicable to the C-V2X and NR-V2X standards. The first of these mechanisms is *DCC reactive*. This is another table-based approach where a delay is placed on message transmission based on CBR threshold values. This is shown in Table B.3.

The efficacy of applying the ETSI *DCC Reactive* mechanism, originally designed for the wireless ITS-G5 standard, directly to C-V2X is shown in Fig. 6.1 with the performance of NR-V2X shown in Fig. 6.2. This is compared to SB-SPS performance without congestion control, labelled *C-V2X No DCC* and *NR-V2X No DCC* respectively. This section will focus on them at a high level before discussing them in more detail in the context of the proposed RRI DCC mechanisms.

It can be observed from Fig. 6.1a that *DCC Reactive (No-GB)* has a positive impact in terms of PDR for C-V2X especially at distances greater than 200 m. Although in the ranges of close to 100 m there is a minimal improvement in PDR performance, this is in contrast to the reduction in terms of CBR, as can be seen in Fig. 6.1b. Overall there is a reduction in CBR of at least 20% but typically greater than 40% considering an average CBR of 28%. This ultimately shows that while *DCC Reactive* can improve scheduling performance, the improvement is not commensurate with the data rate loss due to much increased packet delay. It does not operate effectively with the SB-SPS scheduling mechanism, which results in a significant number of MTs occurring and interrupting the grant mechanism, as outlined in Section 4.2.1 of Chapter 4.

The performance in the NR-V2X standard is shown in Fig. 6.2. The overall trends are much the same as C-V2X with the exception of near distances where in ranges 0 - 100 m the performance of *DCC Reactive (No-GB)* is actually worse than the case of NR_{NoDCC} as shown in Fig. 6.2a. Again, the performance of CBR is similar as shown in Fig. 6.2b, in that CBR is some 40% lower than that of no congestion control being applied. This further highlights that the performance of *DCC Reactive* when applied to NR-V2X does not operate effectively with the scheduling mechanism and the number of MTs occurring is actually causing

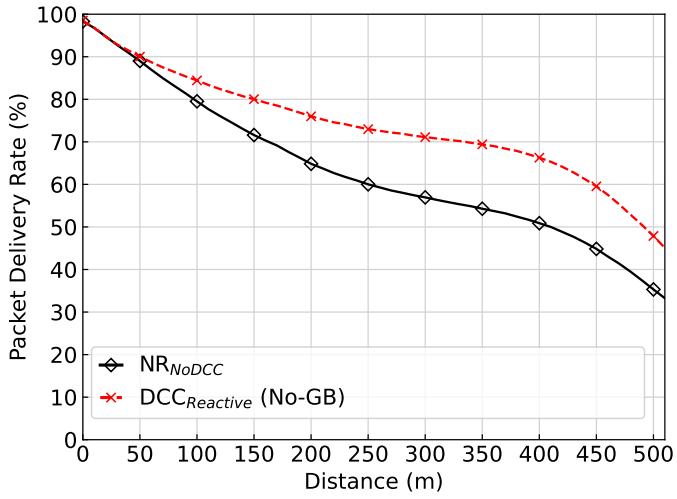

(a) PDR.

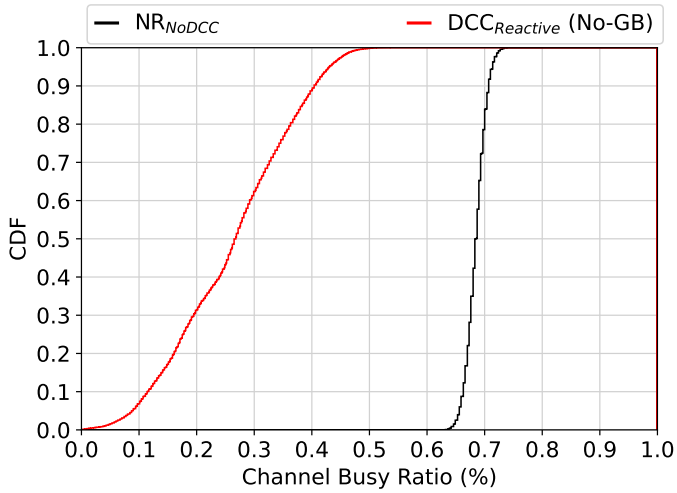

(b) CBR.

Figure 6.1: Performance of ETSI DCC Reactive for C-V2X (Rel. 14).

degraded performance at near distances.

These results are provided to be indicative of the performance of *DCC Reactive* when applied to the SB-SPS mechanism. Ultimately, the introduction of MTs as a result of packet delays being applied results in scheduling performance being degraded. There are a large number of MTs in this case due to the way in which the *DCC Reactive* mechanism operates, namely in adding a delay to packet transmission. On average the delay applied of 252 ms, results in a considerable number of MTs referring back to the 2n-2 (n=RRI) analysis provided in Section 4.2.2. In this configured scenario the RRI is set to 100 ms

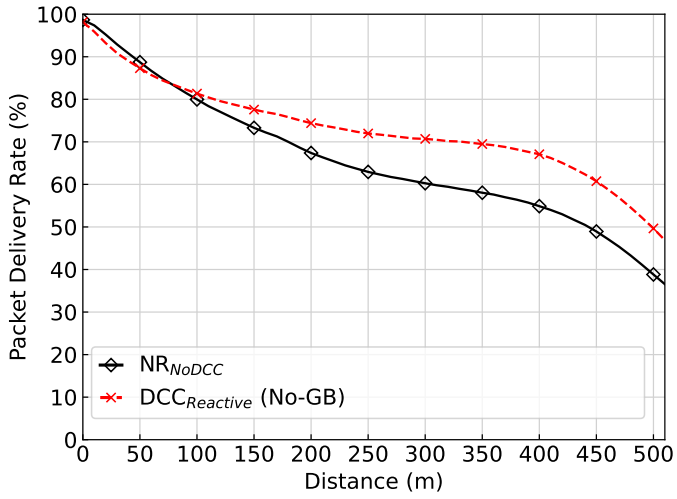

(a) PDR.

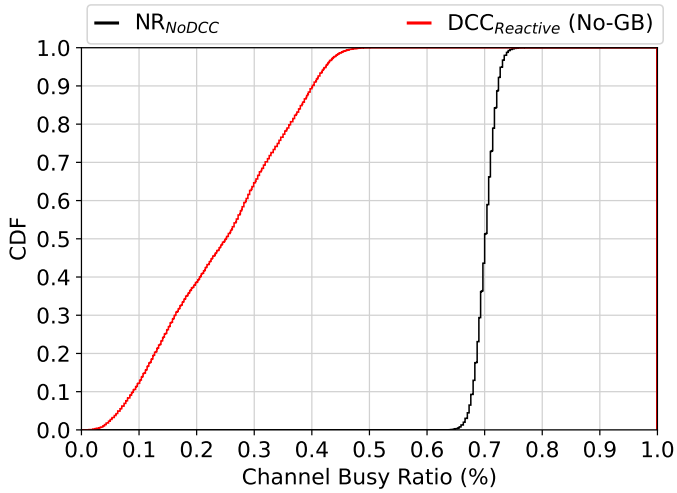

(b) CBR.

Figure 6.2: Performance of ETSI DCC Reactive for NR-V2X (Rel. 16).

as such the minimum IPT must be 198 ms to avoid grant breaking which is exceeded by *DCC Reactive*, resulting in a higher number of MTs. In both technologies this exhibits as either the same or degraded performance at near distances, while the reduced CBR level of 40% results in the performance at distance beyond 150 m improving by a significant amount of 10%. It is also noteworthy that for both technologies *DCC Reactive* exhibits significant CBR instability with the CDF for both ranging from 5% CBR up to 48%.

For subsequent sections given the low average CBR of *DCC Reactive* to provide a more fair analysis of each mechanism and also to show the impact of high

levels of packet dropping or delay a CBR target of 20% has been configured for *DCC Adaptive* mechanisms and for packet dropping approaches based on 3GPP CBR to CR-Limit tables have been designed to target a 20% CBR as well. This enables a more direct comparison between all the applied approaches given the similar average CBR level.

The performance of the other ETSI-standardised approach, namely *DCC Adaptive* is now discussed. Based on LIMERIC [25], this adjusts the packet transmission rate to converge to a target CBR. Default *DCC Adaptive* behaviour is assumed, where the CBR converges to a target of 68%. A CBR target of 20% is also evaluated to allow for direct comparison to *DCC Reactive*. The impact on PDR is shown in Fig. 6.3a. It can be observed that *DCC Adaptive* (68%) has the same performance as no congestion control (*NR-V2X No DCC*) with a slight reduction in the CBR experienced by vehicles, as shown in Fig. 6.3b. For a comparable CBR of 20%, *DCC Reactive* shows marginally better performance than *DCC Adaptive*. This is as a result of the higher packet inter-arrival times for the more stable *DCC Adaptive* which maintains an average of 300ms inter-arrival rate throughout. *DCC Reactive* has a more variable inter-arrival time, characterised by a mean of 252ms, but ranging from 100ms to 400ms. Up to 50% of traffic arrives within 200ms. In accordance with *2n-2* (where n=RRI) as observed earlier, this allows the SB-SPS grant to be maintained more frequently, thus reducing the requirement for rescheduling and allowing for more accurate historical sensing. However, the *DCC Adaptive* approach has advantages, as it avoids the drawbacks of table-based mechanisms, such as CBR instability around fixed thresholds [26] and the need to fit the table to variable network demand.

So far, we have investigated the effectiveness of the current generation of ETSI standard congestion control techniques being applied to the C-V2X and NR-V2X standards. Our overall observation is that although performance improvements can be made considering an approach such as *DCC Adaptive*, the performance improvement is not commensurate with the CBR reductions imposed. TRC based approaches need to be adapted to be able to conform with the SB-SPS mechanism and avoid issues with grant breaking or missed transmissions, which result in poor overall performance.

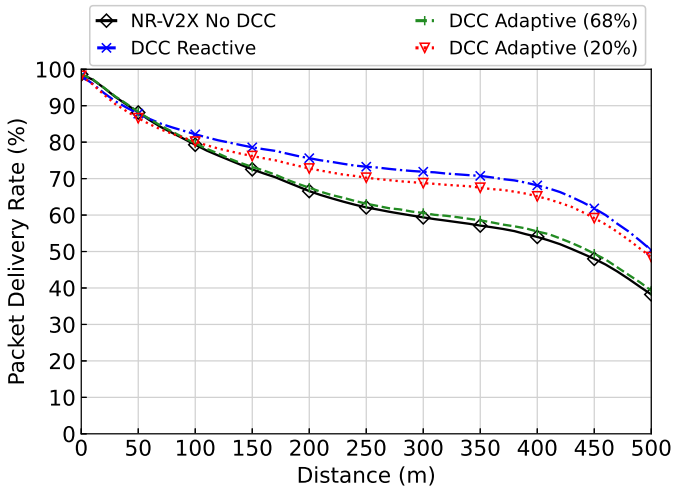

(a) ETSI DCC PDR (Reactive & Adaptive) Mode 2.

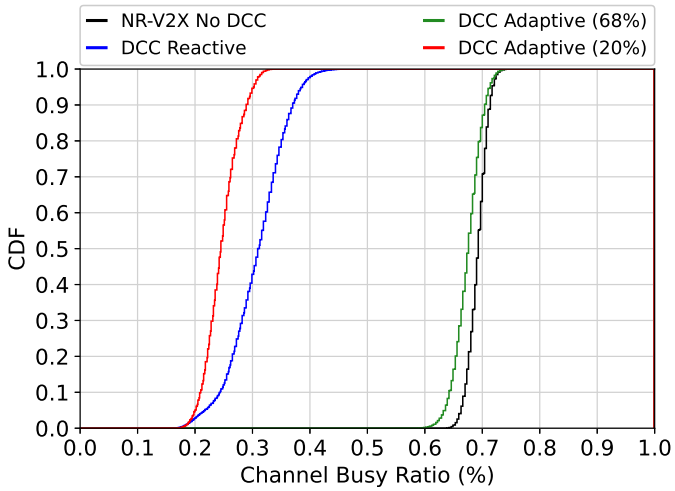

(b) ETSI DCC CBR (Reactive & Adaptive) Mode 2.

Figure 6.3: Performance of ETSI DCC (Reactive & Adaptive) for NR-V2X (Rel. 16) with GB.

6.3.2 ETSI & 3GPP packet dropping (TRC) mechanisms

As described previously in Section 2.3.7, 3GPP have outlined several approaches to congestion control, all of which are based on CBR to CR-Limit tables. These tables work on a lookup basis; depending on the inferred channel congestion, a limit is set on the channel usage for each vehicle. One of the approaches to maintain this limit is packet dropping, where packets are simply dropped before transmission to maintain the limit. Three different CR limit lookup tables are considered in this study, two of which are proposed by

the standards bodies:

- The ETSI CR limit table is shown in Table B.2. A traffic priority of 6-8 is assumed, as other priorities had no impact on channel conditions. This is labelled *Packet Dropping (ETSI)*.
- As part of a 3GPP working group, Qualcomm has also proposed a CR-Limit table shown in Table B.1 [1]. This is labelled *Packet Dropping (3GPP)*.
- We proposed an adapted 3GPP CR limit table, with lower CBR thresholds. We specified this to more aggressively manage congestion and to allow us to directly compare against approaches such as the *DCC Reactive* scheme, that produces a mean CBR closer to 20%, this is shown in Table 6.3. This is labelled *Packet Dropping (Aggressive)*.

Table 6.3: CBR to CR-Limit table achieving a 20% CBR target.

CBR	**CR limit**
CBR $\leq$ 0.14	0.8e-3
0.14 $<$ CBR $\leq$ 0.22	0.7e-3
0.22 $<$ CBR $\leq$ 0.3	0.6e-3
0.3 $<$ CBR $\leq$ 0.38	0.5e-3
0.38 $<$ CBR $\leq$ 0.46	0.4e-3
0.46 $<$ CBR $\leq$ 0.6	0.3e-3
0.6 $<$ CBR $\leq$ 0.7	0.2e-3
0.7 $<$ CBR $\leq$ 0.8	0.1e-3
0.8 $<$ CBR $\leq$ 0.9	0.1e-3
0.1 $<$ CBR	0.1e-3

The PDR performance for a packet dropping congestion control mechanism according to these CR limit tables can be seen in Fig. 6.4a. It can be observed that *Packet Dropping (ETSI)* performs identically to a scenario where no congestion is employed, labelled *No DCC*, both in terms of PDR and CBR as shown in Fig. 6.4b. This is because the ETSI CR limits associated with the CBR thresholds are high and hence not sufficiently responsive. *Packet Dropping (3GPP)* performs better due to reduced CBR, as shown in Fig. 6.4b, although this only results in an improvement in PDR of up to 9% at distances exceeding 150 m. *Packet Dropping (Aggressive)* adopts a more aggressive approach to reducing CBR which improves PDR by up to 26% at distances beyond 100m. However, the improvement in PDR does not correlate with a significantly lower CBR. This is caused by the incompatibility between the packet

generation rate and the SB-SPS scheduling algorithm; this will be explored in greater detail in subsequent sections comparing both these approaches and the proposed mechanisms.

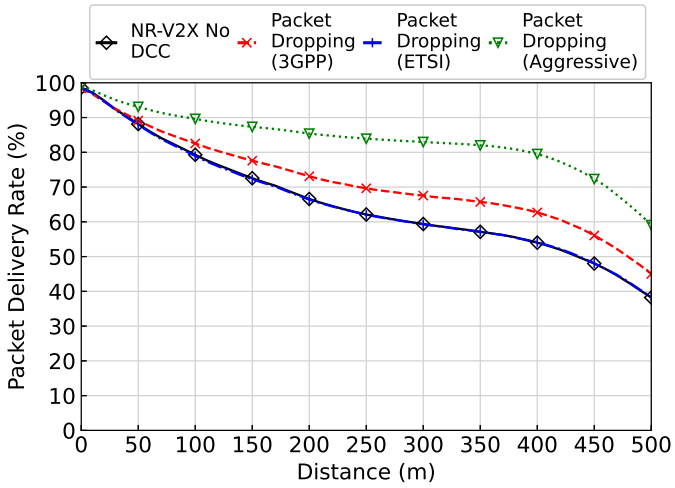

(a) Packet Dropping PDR.

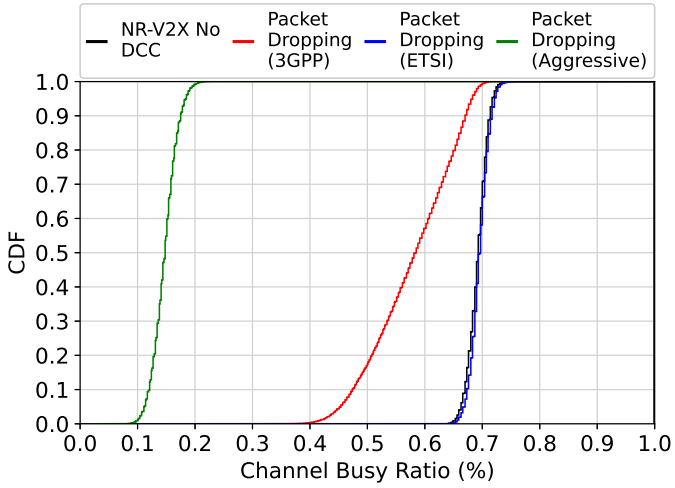

(b) Packet Dropping CBR.

Figure 6.4: Performance of ETSI & 3GPP Packet Dropping for NR-V2X (Rel. 16).

6.3.3 Summary of findings based on evaluation of standardised DCC

The primary findings are as follows.

1. Current congestion control mechanisms based on transmission rate control will result in missed transmissions and/or grant breaks, which

degrades SB-SPS performance.

2. While some performance improvements can still occur, the improvement in performance is not proportional to the decrease it incurs in channel usage, as shown by the differential in the CBR in Fig. 6.3b.

Before proposing a solution to this problem, the next section considers whether the use of MCS congestion control might be preferable to TRC mechanisms as an approach for congestion control.

6.4 MCS congestion control techniques

MCS congestion control has been proposed as a mechanism to address the distributed control congestion by 3GPP and ETSI [41] and in the literature [28, 42]. We now model and evaluate how MCS congestion control would perform if used for this purpose. This research was a collaboration with researchers in UNAM. This forms an elements of **contribution 4** focusing on congestion control in C-V2X and NR-V2X.

Using CBR and CR-limit tables, e.g. Table B.2, each vehicle can implement different mechanisms to reduce congestion i.e. packet dropping, increasing the MCS, power control etc. As described in Section 2.3.6, increasing the MCS can reduce the number of required subchannels for a given transmission and therefore reduce the CR. Moreover, even if the number of subchannels remains the same due to the C-V2X channelisation scheme, increasing the MCS reduces the number of required RBs, which can reduce the average RSSI of the subchannel and ultimately the CBR.

We evaluate the strategy of using MCS congestion control as a DCC mechanism, where each vehicle adapts the MCS according to the measured CR and CBR. To this end, two MCS configurations are implemented. MCS 7 is set by default in all vehicles as specified in [82] when congestion is not detected in the channel. MCS 11, which is the maximum MCS specified by ETSI in [10], is implemented if congestion conditions are met. With this mechanism, each vehicle increases the MCS from 7 to 11 when the CR exceeds the limit for its corresponding CBR and returns to MCS 7 otherwise.

Figure 6.5 shows the PDR versus distance for all MCS configurations and vehicular densities. As expected, the PDR decreases for both MCS 7 and MCS 11 as the distance and vehicular density increase, due to the reduction in SINR

and the increase in interference, respectively. It can also be noted that MCS 7 always outperforms MCS 11 with the difference in PDR increasing at higher vehicular densities. This can be attributed to the more robust modulation and coding rate of MCS 7 (QPSK/0.7), which outperforms MCS 11 (16QAM/0.5) at higher interference levels.

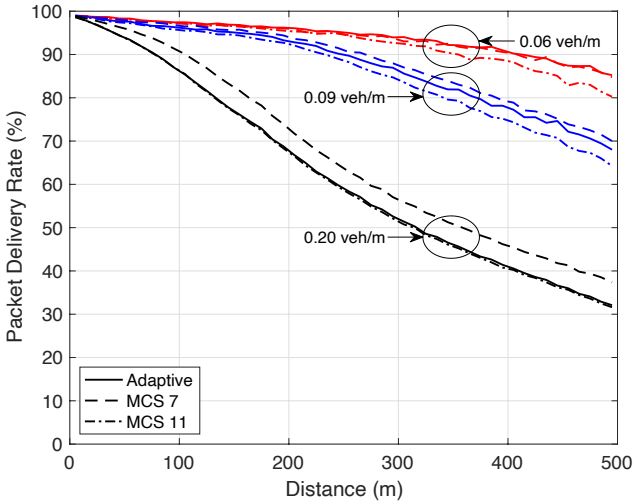

Figure 6.5: Packet Delivery Rate MCS congestion control.

Figure 6.5 also shows how the performance of MCS congestion control varies depending on the density of the vehicle. For example, at 0.06 veh/m, the adaptation of MCS achieves the same PDR as MCS 7 because the vehicular density is not high enough to trigger the congestion control mechanism. However, as the density increases to 0.09 veh/m, some vehicles in the scenario trigger MCS congestion control changing the MCS from 7 to 11, slightly reducing the overall PDR. This is more evident at 0.20 veh/m, where MCS congestion control has the same performance as MCS 11 due to all vehicles increasing their MCS due to the higher vehicular density. Ultimately, it can be seen that adapting the MCS does not improve PDR in response to congestion. This is attributable to the constraints imposed by the standard regarding the subchannel and MCS configurations, which limit the reduction in subchannel occupation despite the reduction in the number of required RBs.

Figure 6.6 shows the percentage of transmission errors versus the cause of error for all MCS configurations. The results are shown for a density of 0.09 veh/m and correspond only to errors within a range of 500 m. It can be seen that the percentage of half-duplex errors is the same for all MCS configurations. This is expected, as the transmission rate for MCS 7 and MCS 11 is the same, resulting in the same number of lost transmissions due to half

duplex limitations. In terms of sensing errors, which occur when the received power is below the sensing threshold, MCS 7 performs slightly better than MCS 11. This occurs because MCS 7 can implement a higher transmission power than MCS 11 while maintaining the PSD within the 23 dBm/MHz limit. Specifically, the transmission power can be set at 27.32 dBm for MCS 7 and 25.97 dBm for MCS 11 because MCS 7 occupies a larger bandwidth (15 RBs) than MCS 11 (11 RBs). In the case of MCS congestion control, the percentage of sensing errors is within the range of MCS 7 and MCS 11 as some vehicles increase the MCS when the mechanism is triggered at the configured density.

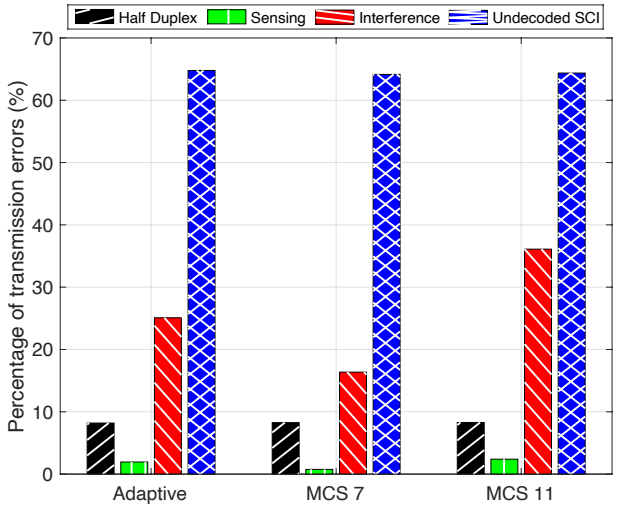

Figure 6.6: Percentage of transmission errors versus cause of error for all MCS configurations at 0.09 veh/m.

Figure 6.6 also shows the percentage of transmission errors due to interference, which occur when the packet cannot be decoded due to low levels of SINR. In this case, MCS 7 (QPSK/0.7) outperforms MCS 11 (16QAM/0.5), as the former is more robust when the same levels of interference are present in the channel. In the case of MCS congestion control, the errors due to interference are within the range of MCS 7 and MCS 11 as the MCS increases in some of the vehicles in the scenario.

Finally, Figure 6.6 shows the percentage of transmission errors due to undecoded SCIs. This type of error occurs when the SCI cannot be decoded at the receiver, resulting in the corresponding TB being immediately discarded. This is attributable to the fact that errors are analysed within a 500 m range, where interference from SCI collisions is more likely to cause undecoded SCIs and impact TB transmissions. Figure 6.6 also shows that the selection of different MCS configurations has no impact on SCI collisions, as the MCS only

modifies the number of RBs for the TB and not the SCI.

To validate the operation of MCS congestion control, Figure 6.7 shows the CBR for all vehicular densities and MCS configurations, while Figure 6.8 shows the MCS usage for MCS congestion control for all vehicular densities.

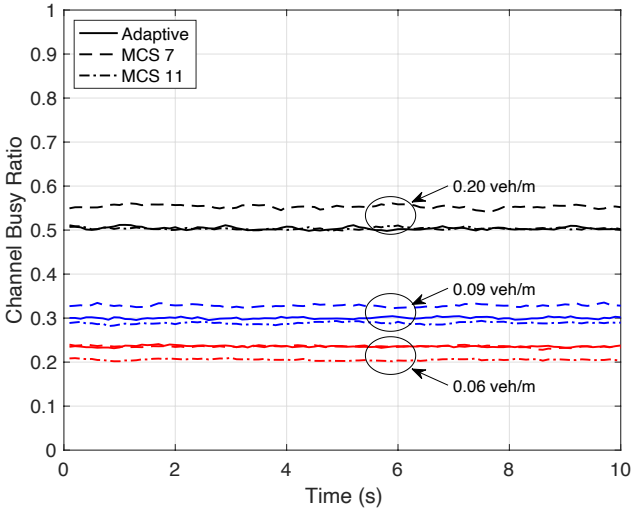

Figure 6.7: Channel Busy Ratio MCS congestion control.

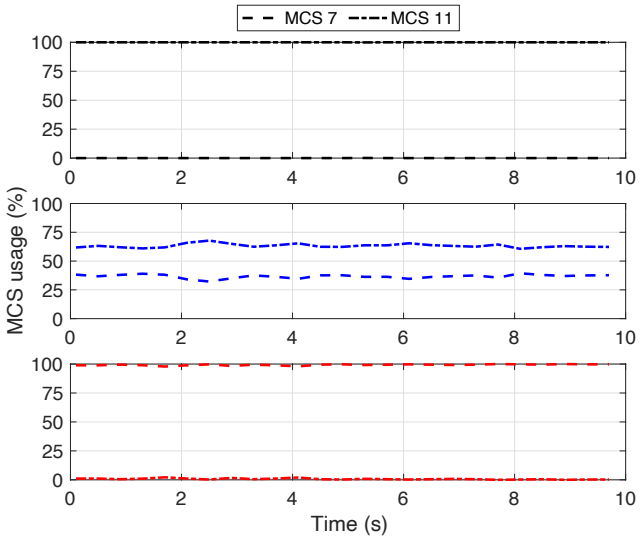

Figure 6.8: MCS usage for MCS congestion control at 0.06 veh/m (lower figure), 0.09 veh/m (middle figure) and 0.20 veh/m (upper figure).

In Figure 6.7, the CBR levels at 0.06 veh/m are lower than 0.3, which is the minimum CBR required to trigger MCS congestion control according to Table B.2. Therefore, at 0.06 veh/m and with MCS congestion control enabled, almost all vehicles in the scenario implement MCS 7, as shown in Figure 6.8. It can also be seen that for 0.09 veh/m, the CBR increases around the minimum

CBR limit of 0.3 which is enough to trigger MCS congestion control. This is validated in Figure 6.8 where it is shown that around 60% of vehicles increase their MCS to 11, while the remaining 40% implement MCS 7. At 0.20 veh/m, Figure 6.7 also shows that the CBR reaches 0.5 as its highest value, which is higher than the minimum limit of the CBR. Consequently, MCS congestion control is activated in all vehicles of the scenario, increasing their MCS to 11 as shown in Figure 6.8.

Importantly, the results of Figure 6.7 indicate that CBR levels are generally lower for MCS 11 across all vehicular densities. This is the result of the reduction in RSSI of the subchannel that occurs as a result of the lower RB occupation of the higher MCS. Moreover, this result highlights the main advantage of increasing the MCS in congested scenarios, which is an overall reduction in channel congestion.

Finally, in the case of MCS congestion control, Figure 6.7 shows how the CBR levels vary depending on the vehicular density. For instance, at 0.06 veh/m the CBR of MCS 7 and MCS congestion control are practically the same due to the lower density in the scenario, which is below the CBR limit required to trigger MCS congestion control. As the density increases to 0.09 veh/m, MCS congestion control is able to reduce the overall CBR, as 60% of the vehicles in the scenario switch to MCS 11 to cope with the increased congestion in the channel. This trend is even clearer at 0.20 veh/m, where all vehicles switch to MCS 11 due to high congestion, achieving a higher reduction on CBR.

This reduction in CBR can be validated through Figure 6.9 where the average RSSI of the subchannel is shown for all MCS configurations and vehicular densities. It can be seen that the average subchannel RSSI increases with the vehicular density due to the higher levels of interference in the scenario. More importantly, it shows that the average subchannel RSSI of MCS 7 is higher than that of MCS 11 for all vehicular densities, indicating that the higher RB occupation of MCS 7 with respect to MCS 11 increases the subchannel RSSI and ultimately the CBR, as shown in Figure 6.7.

Figure 6.9 also shows how the average subchannel RSSI of MCS congestion control depends on the configured vehicular density. For example, at 0.06 veh/m the number of vehicles in the scenario is not enough to trigger MCS congestion control, and therefore subchannel RSSI is the same as MCS 7. However, as the density increases to 0.09 veh/m and 0.20 veh/m the average subchannel RSSI is lower for MCS congestion control in comparison with MCS

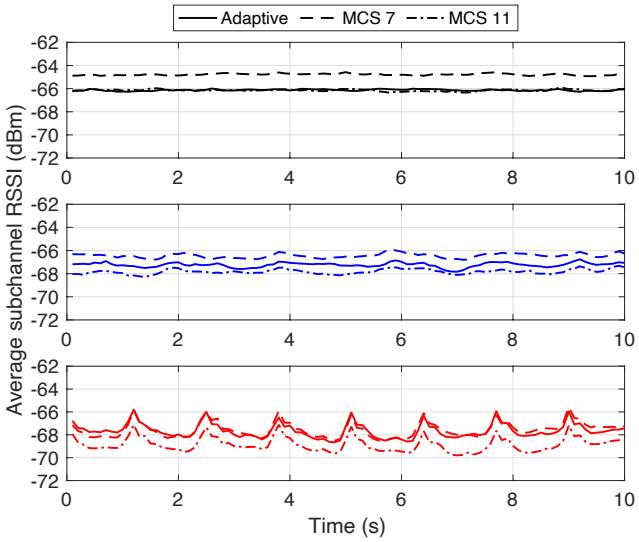

Figure 6.9: Average subchannel RSSI at 0.06 veh/m (lower figure), 0.09 veh/m (middle figure) and 0.20 veh/m (upper figure).

7 because the vehicles in the scenario increase their MCS due to the increased congestion in the channel.

The results indicate that under high vehicular density, MCS congestion control is capable of reducing subchannel interference and channel congestion. However, the results also highlight the limitations due to the channelization constraints of C-V2X and demonstrate that the technique does not improve the overall performance of C-V2X under the analysed congestion conditions.

6.4.1 Summary of findings based on evaluation of MCS congestion control

A thorough analysis of MCS congestion control shows that the approach has limited impact when applied to C-V2X and NR-V2X due to subchannel configuration, limited standardised availability of MCS and finally due to packet size variability.

6.5 Resource Reservation Interval Decentralised Congestion Control (RRI DCC)

Taking into account our findings from Section 6.3 and 6.4, it can be concluded that current mechanisms are not suitable for use in C-V2X and NR-V2X and a dedicated mechanism must be implemented to significantly improve

performance.

In this section, we evaluate the performance of the proposed RRI DCC mechanisms described in Section 4.3.2, which is a major contribution of this thesis **(C4)**. To briefly reiterate the proposed mechanisms: All are designed to adapt the RRI parameter to ensure that skipped transmission resources are correctly identified by SB-SPS. By ensuring that these resources are correctly skipped over, the SB-SPS algorithm can maintain a correct sensing history, improving scheduling performance and reducing the overall number of transmissions and network congestion.

Three variants are suggested:

- RRI_{CR_limit}: Approach based on CBR to CR-Limit tables, which is the 3GPP standard.
- $RRI_{Reactive}$: Approach based on the DCC reactive table shown in Table B.3.
- $RRI_{Adaptive}$: Approach based on the DCC adaptive algorithm which is derived from LIMERIC.

Specifically, the performance across 3 areas is considered; traditional access (network) layer performance, application layer performance, and finally a thorough analysis with respect to fairness and stability. The performance of all three schemes are evaluated for the network and application layer evaluations and it is ultimately shown $RRI_{Adaptive}$ exhibits the best performance. As such, $RRI_{Adaptive}$ will be the primary focus when discussing fairness and stability.

6.5.1 RRI DCC - Network Performance

The proposed *RRI DCC* congestion control schemes are compared against *NR-V2X No DCC*, *DCC Reactive*, *DCC Adaptive* (CBR target of 20%) and *Packet Dropping (Aggressive)*. The first network configuration considers a target load of 20% CBR in the network. Although this is excessively restrictive, it is used to highlight the impact of significant packet dropping on the performance of the SB-SPS scheduler, as well as to enable a fairer comparison with the *DCC Reactive* mechanism. A similar phenomenon might occur in a network simultaneously hosting multiple services such as CAMs, CPMs, DENMs and Local Dynamic Map packets. In such cases, while the load in the network will be significantly higher than what we consider as a benchmark (a single service without congestion control), it will have the same effect; it allows for

comparison on the performance of the congestion control techniques in a scenario where there is a significant level of packet dropping. This is to demonstrate how SB-SPS performance can be degraded when there is no effective means of maintaining the SB-SPS RRI.

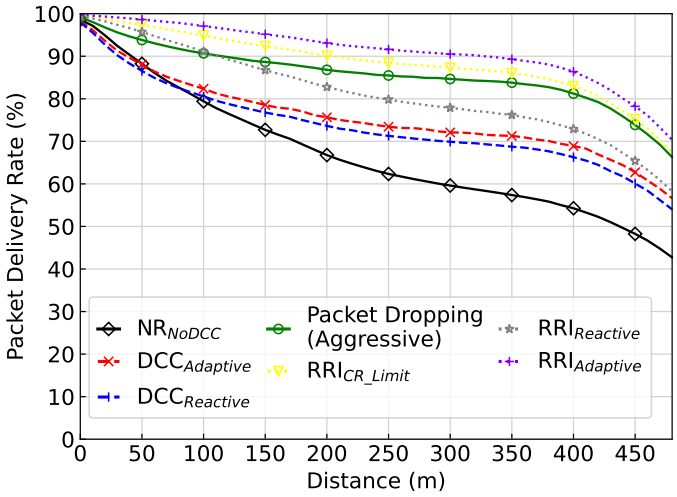

(a) Packet Delivery Rate.

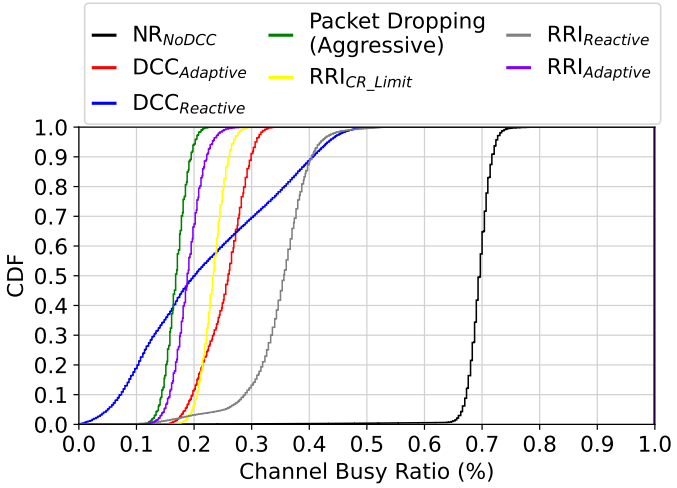

(b) Channel Busy Ratio.

Figure 6.10: RRI DCC versus standardised Congestion Control at 20% CBR (no grant breaking).

Fig. 6.10a shows the PDR performance of $RRI_{Reactive}$ outperforms *DCC Adaptive* by up to 9% PDR at near distances, but Fig. 6.10b demonstrates less stability with respect to CBR. This is because it inherits the instability of the *DCC Reactive* table lookup mechanism. It exhibits considerably less mean colliding grants compared to standardised DCC approaches, as shown in Table 6.4. The

mean colliding grants are a significant metric as they represent the resources utilised by two or more vehicles simultaneously, due to poor CSR selection in SB-SPS. Colliding grants can occur for four reasons:

1. A missed transmission, causing a vehicle to not send an SCI. Neighbouring vehicles will mistakenly believe a resource to be free in future reserved slots, when they may be utilised. This is indicated as γ_{MT} in Table 6.4.
2. In the case of *RRI DCC* based mechanisms, collisions can occur due to the RRI changing over time. This happens when there is a change in congestion. Resources may be used by multiple vehicles on alternating RRIs (collision-free), and when congestion decreases (causing the RRI to return to default 100ms interval), this can result in additional collisions. This is represented as γ_{RRI} in Table 6.4.
3. Under congested radio conditions, when resources are not determined to be free, vehicles will select resources with low RSRP measurements. This is indicated as γ_{NF} in Table 6.4.
4. Failure to decode the SCI of a neighbouring vehicle can result in a resource being seen as free. This can be as a result of propagation or interference on the resource, and it is still possible to be selected in congested scenarios similar to γ_{NF}. This is indicated as γ_{NSCI} in Table 6.4.
5. When neighbouring vehicles reserve the same resources within a single RRI as a result of similar selection windows. As neighbouring vehicles may have similar RSRP measurements, they are likely to have similar CSR pools, which can result in selecting the same resource(s). This is indicated as γ_{TSim} in Table 6.4.

More notably, RRI_{CR_limit} and $RRI_{Adaptive}$ show significantly improved PDR e.g. an increase of up to 18% on *DCC Adaptive*, when compared to other standardised schemes, even those with comparable CBR. They also exhibit a low number of mean colliding grants, while maintaining highly stable CBR. The primary advantage of $RRI_{Adaptive}$ is that it avoids the drawbacks of table-based mechanisms, which we have outlined previously. It also exhibits a more precise maintenance of CBR around the target of 20%, this is the primary explanation for its performance improvement over RRI_{CR_limit} which averages a CBR slightly above the target of 25%.

Table 6.4: Absolute number of colliding grants & causation for each congestion control mechanism.

CBR	Congestion Control Mechanism	Mean Colliding Grants (γ)	γ_{MT}	γ_{RRI}	γ_{NF}	γ_{TSim}	γ_{NSCI}
70%	*NR-V2X No DCC*	7444	-	-	4531	1276	1637
20%	*DCC Adaptive*	7247	2127	-	0	747	4373
	DCC Reactive	8866	2198	-	33	1074	5561
	Packet Drop. (Aggr.)	2060	302	-	0	227	1531
	$RRI_{Reactive}$	2327	-	50	1718	184	375
	$RRI_{CRlimit}$	1408	-	180	787	68	373
	$RRI_{Adaptive}$	2133	-	39	1552	286	256
60%	*DCC Adaptive*	9117	1366	-	3553	1175	3023
	$RRI_{CRlimit}$	5242	-	1216	2648	557	821
	$RRI_{Adaptive}$	6043	-	1133	3397	827	686

DCC mechanisms are very susceptible to colliding grants due to γ_{MT}, γ_{TSim}, and γ_{NSCI} due to the incompatibility of their schemes with the SB-SPS RRI mechanism. As reserved resources are maintained for 5-15 transmissions, as determined by the RC, recurring colliding grants represent multiple avoidable half-duplex errors in the channel as well as increasing interference. This is evident in Table 6.4. On the contrary, the *RRI DCC* approaches account for this and therefore incur significantly fewer colliding grants by eliminating γ_{MT} and reducing γ_{TSim} and γ_{NSCI}. Their additional source of error also does not contribute as significantly to degradations in performance as shown by relatively low collisions γ_{RRI}. The variance in $RRI_{Reactive}$ CBR results in a higher number of colliding grants than the more stable RRI_{CR_limit} and $RRI_{Adaptive}$ approaches.

The second network configuration considered has a CBR target of 60%, considered the default in many academic papers [27], for both congestion control mechanisms. This represents a more likely network target load. In this case, we will focus only on the mechanisms that easily conform to this target, i.e. RRI_{CR_limit} with a new table and $RRI_{Adaptive}$ and *DCC Adaptive* which can be configured to target it.

In this case, there is a minimal performance improvement of $RRI_{Adaptive}$ and RRI_{CR_limit} over *DCC Adaptive* as shown in Fig. 6.11a and all mechanisms are effective in meeting the CBR target of 60% as shown in Fig. 6.11b. This is due to the much smaller reduction in CBR compared to NR_{NoDCC}. The less congested scenario means that packet dropping/delay is not as prevalent as in the

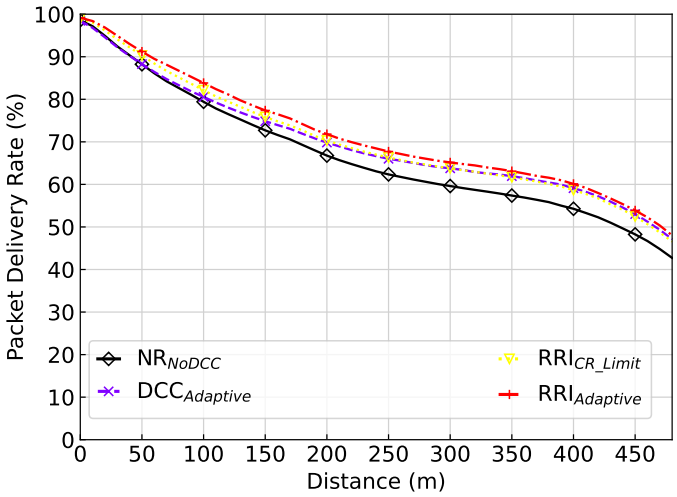

(a) Packet Delivery Rate.

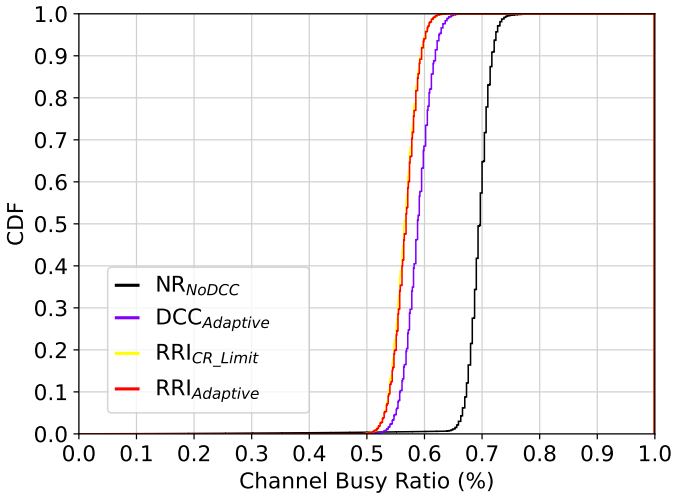

(b) Channel Busy Ratio.

Figure 6.11: RRI DCC versus standardised Congestion Control at 60% CBR (no grant breaking).

previous scenario. As shown in Table 6.4, $RRI_{Adaptive}$ reduces the mean number of colliding grants by 20% compared to $NR\text{-}V2X_{NoDCC}$ and the trend across the other categories is similar to the 20% CBR scenario, with RRI_{CR_limit} reducing them even more effectively. The main observation is the disadvantage in applying RRI_{CR_limit} i.e. that a new table needs to be designed for each scenario rather than simply enabling a new configuration of the LIMERIC algorithm. This makes it significantly more inflexible than an adaptive approach.

6.5.1.1 Impact of scenario density

This section focuses on the impact that various densities have on the performance of the proposed solutions. In this case 3 different techniques are considered to demonstrate this those being *DCC Reactive*, *DCC Adaptive* and $RRI_{Adaptive}$. All are compared to the default NR_{NoDCC} and the parameters are the same as described in Table 6.1. The key change being the density of traffic increasing across the scenarios with β = 0.12v/m and 0.2v/m and the road networking being the 3GPP standard 2km 6 lane highway scenario. In both cases *DCC Adaptive* and $RRI_{Adaptive}$ are configured with a target of 60% CBR.

The first set of results focuses on the low density of β 0.12v/m as shown in Fig. 6.12. In this case it is clear that the performance of $RRI_{Adaptive}$ and *DCC Adaptive* are identical to that of NR_{NoDCC}, this is due to the fact that this low density scenario exhibits low CBR as shown in Fig. 6.12b. Only for *DCC Reactive* is congestion control applied due to its more restrictive table mechanism and results in improved performance in terms of PDR shown in Fig. 6.12a. This again shows similar trends as previously explored with those PDR gains only being at distances greater than 200m with a reduced CBR of 18%. This ultimately represents a significant under-utilisation of channel resources with no significant benefit in terms of PDR particularly at near distances.

The next density considered is 0.2v/m as shown in Fig. 6.13. This scenario represents a significant increase in congestion with NR_{NoDCC} seeing congestion increase to an average of 70% CBR as seen in Fig. 6.13b.*DCC Reactive* again exhibits similar performance to the 0.12v/m density and again under utilises channel resources with a low level of CBR at 32% shown in Fig. 6.13b while showing similar trends in PDR at distances beyond 200m as shown in Fig. 6.13a. In this case *DCC Adaptive* and $RRI_{Adaptive}$ both exhibit similar performance with marginal PDR differences between them and both performing 10% better in terms of PDR than NR_{NoDCC} as shown in Fig. 6.13a. This being as a result of the reduced congestion in the channel as shown in Fig. 6.13b, in this case $RRI_{Adaptive}$ exhibits a higher CBR than the *DCC Adaptive* case, while maintaining the same PDR performance.

As shown the adaptive mechanisms maintain better channel utility across both densities as well as those already explored in previous sections on the shorter road network. This flexibility is a useful aspect of the mechanisms as it avoids the drawbacks of table based approaches which may require tuning to new scenarios to better utilise channel resources without incurring too greater a

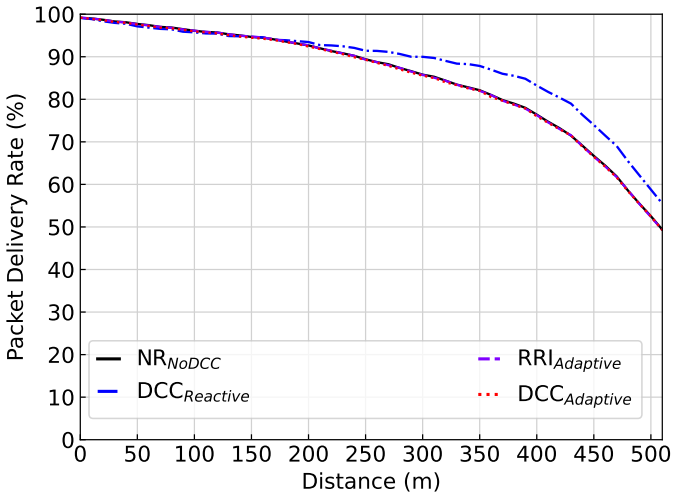

(a) Packet Delivery Rate.

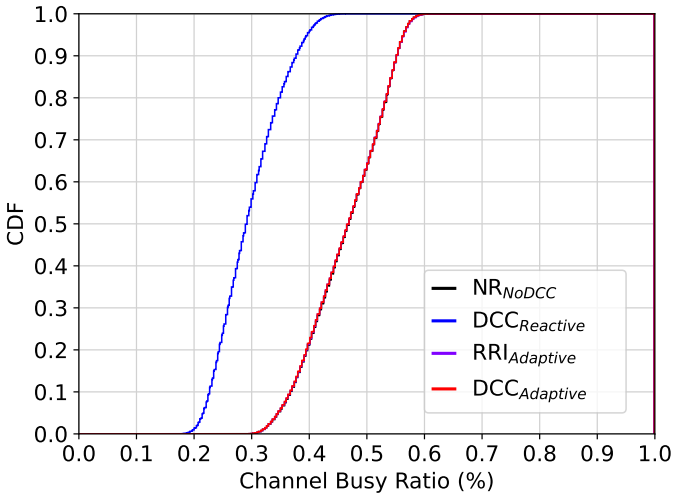

(b) Channel Busy Ratio.

Figure 6.12: RRI Adaptive versus standardised Congestion Control at β = 0.12v/m.

load such that transmission performance degrades significantly. It is also is of note that $RRI_{Adaptive}$ maintains the same high performance as shown in previous scenarios with similar high channel utilisation when compared to *DCC Reactive* and *DCC adaptive*. The mechanism is successful in adapting to various vehicular densities.

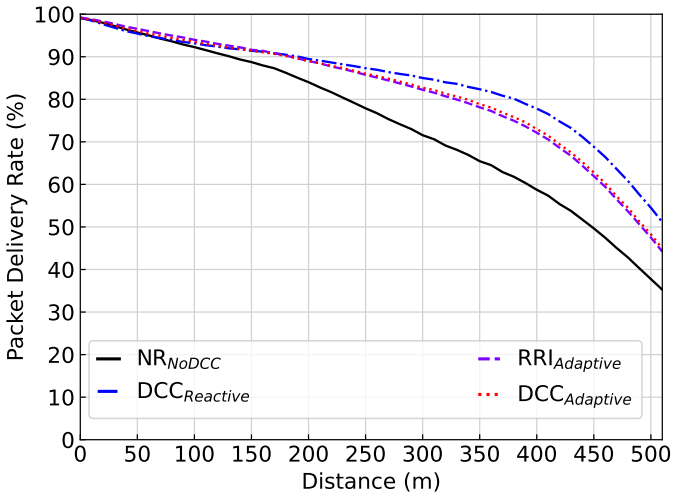

(a) Packet Delivery Rate.

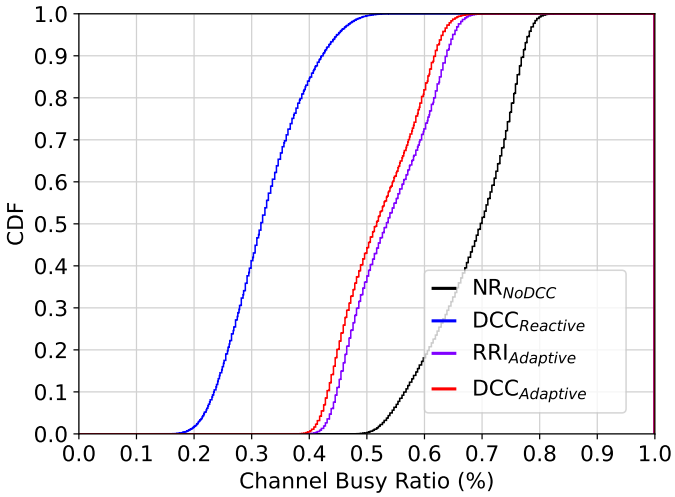

(b) Channel Busy Ratio.

Figure 6.13: RRI Adaptive versus standardised Congestion Control at β = 0.2v/m.

6.5.2 RRI DCC - Application Performance

Next, we investigate the performance of the proposed congestion control mechanisms relative to application layer KPIs. The primary KPIs considered in this study are the Inter-Packet Gap (IPG) and the mean awareness of neighbours in the 200-300m range. IPG represents the average inter-packet arrival rate on a single vehicle and is compared over distance between the sender and receiver. This metric indicates the regularity of updates that a vehicle receives from its neighbouring vehicles in a given range. It is impacted

by channel conditions and applied congestion control mechanisms. Mean neighbour awareness represents the proportion of vehicles in the 200m-300m range of which a vehicle is aware, i.e. it has received a CAM from those vehicles and knows their position, heading, speed, etc.

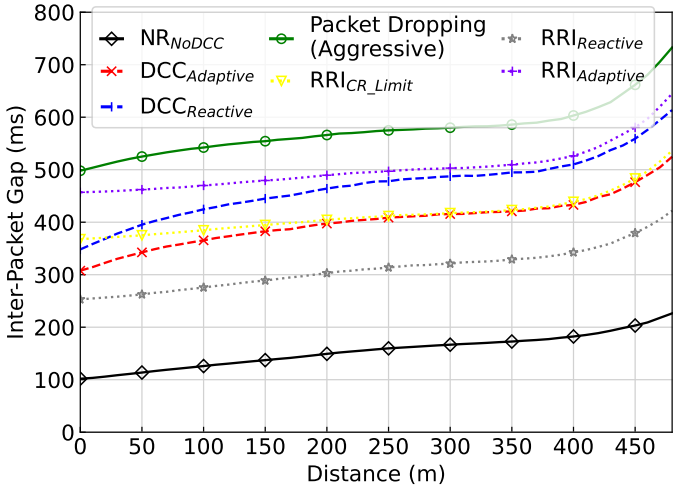

(a) IPG at 20% CBR.

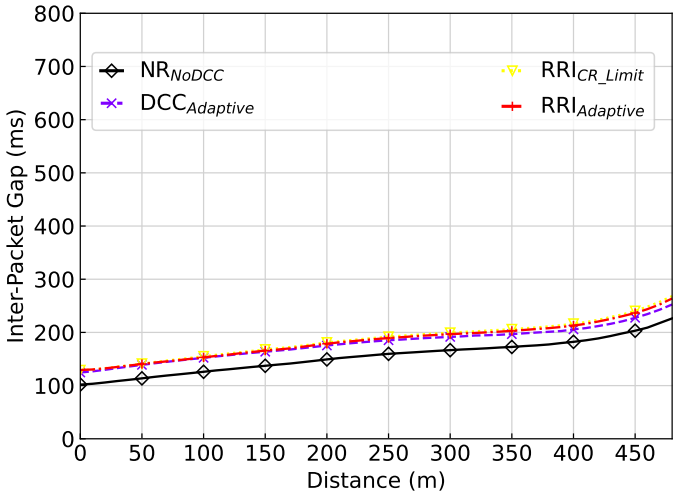

(b) IPG at 60% CBR.

Figure 6.14: Inter-Packet Gap (IPG) performance of congestion control mechanisms.

In Fig. 6.14a, $RRI_{Reactive}$ demonstrates the lowest IPG while exhibiting better PDR (Fig. 6.10a), than standardised approaches except for the two other RRI schemes and *Packet Dropping (Aggressive)* which exhibits significantly higher IPG, followed by $RRI_{Adaptive}$, due to their much more aggressive approaches maintaining 20% CBR. Of the remaining mechanisms at the 20% CBR target,

all exhibit an IPG of 360 to 450 ms on average with RRI approaches exhibiting slightly higher IPG than *DCC* variants. Neighbour vehicle awareness is shown in Table 6.5 and is comparable with other approaches. The best performing schemes are RRI_{CR_limit} and *DCC Reactive* however RRI_{CR_limit} achieves better PDR performance and lower congestion than *DCC Reactive*.

Table 6.5: Neighbour Vehicle Awareness (200m-300m)

CBR Target	Congestion Control Mechanism	Awareness %	Std. Dev. %
70%	*NR-V2X No DCC*	91.5	1.5
20%	*DCC Adaptive*	91.5	1.7
	DCC Reactive	96.4	1.2
	Packet Dropping (Aggressive)	94.9	1.3
	$RRI_{Reactive}$	92.4	1.7
	RRI_{CR_limit}	96.4	1.2
	$RRI_{Adaptive}$	94.8	0.6
60%	*DCC Adaptive*	97.2	0.5
	RRI_{CR_limit}	97.4	1.0
	$RRI_{Adaptive}$	97.8	0.4

The 60% CBR target scenario provides a more fair comparison between the three methods, and here they exhibit almost identical IPG in Fig. 6.14b but with *RRI DCC* mechanisms that exhibit marginally lower CBR. It is clear in all cases that applying any of these congestion control mechanisms will result in increased IPG as illustrated by the increase over the $NR_{No\ DCC}$ but this represents the trade-off between maintaining a high PDR at the expense of latency. However, they all exhibit improvements in terms of mean neighbour awareness of up to 6% so the trade-off is not such that it impacts overall awareness. It is noteworthy that no mechanism clearly prevails in this comparison with nearly identical awareness between all three.

6.5.3 RRI DCC - Fairness and Stability

In this section, the focus will be on the fairness amongst vehicles and the stability of channel conditions for the proposed RRI DCC mechanisms. This is a topic which has been overlooked by the academic community in the context of C-V2X and NR-V2X congestion control.

The performance of $RRI_{Adaptive}$ is now evaluated against that of *DCC Adaptive*, which was shown to have the best performance of the currently standardised congestion control approaches. RRI_{CR_limit} has been eliminated as it requires

table tuning, and as such the fairness of the mechanism is almost entirely dependent on the configured table. It may be possible to design a highly stable and fair table, but the requirement to re-configure this for each scenario is simply prohibitive. This is a vital aspect of any congestion control mechanism and in the context of both C-V2X and NR-V2X has been widely overlooked by the academic community; as such, we believe it is vital to consider it carefully when suggesting a possible technique to adopt. Ultimately, a technique cannot be considered to be effective if while maintaining a stable CBR it is forcing some vehicles to be starved of resources. This could result in increased collision risks considering the CA service. In this study, three network configurations are considered, all based on the same road network as described in Table 6.1. The first two are identical except that in both cases the target network load for *DCC Adaptive* and $RRI_{Adaptive}$ are configured to 20% CBR and 60% CBR respectively. The final scenario considers a long-form simulation that represents a dynamic scenario in which the channel load changes over time. This is designed to evaluate the stability of congestion control mechanisms with varying densities of network traffic. This is achieved stepwise, with only 25% of the vehicles initially transmitting and subsequently increasing by 25% over 20 s time intervals. The channel load is then reduced in a similar way, with 25% fewer vehicles transmitting per 20 s interval.

As we explicitly consider the fairness of $RRI_{Adaptive}$ when compared to *DCC Adaptive*, we use the Jain's fairness index [108]. This well-known index operates based on the desired transmission rate for each vehicle compared to their allocated rate. The goal is to achieve a score of 1 which is considered perfect fairness with low scores indicating unfairness where vehicles are starved of radio resources. We also investigated the convergence of congestion control techniques in terms of CBR and the stability of CR and message transmission rate for each node throughout our simulation study.

Table 6.6: Jain's Fairness evaluation

CBR Target	Congestion Control Mechanism	Fairness %
20%	*DCC Adaptive*	100
	$RRI_{Adaptive}$	99.75
60%	*DCC Adaptive*	99.41
	$RRI_{Adaptive}$	99.98
60%	*DCC Adaptive (longform)*	95.66
	$RRI_{Adaptive}$ *(longform)*	99.07

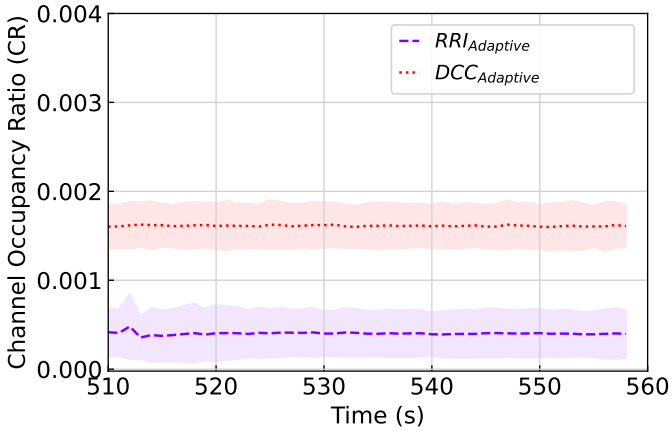

(a) Channel Occupancy Ratio.

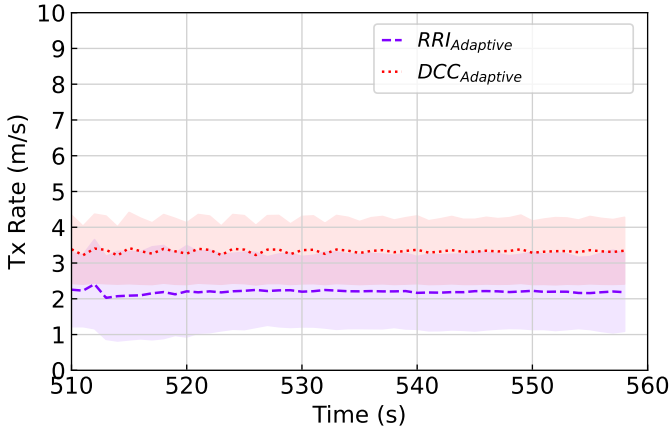

(b) Average Message Rate.

Figure 6.15: Comparative stability statistics at 20% CBR.

In Table 6.6, when considering the two network configurations already evaluated in the previous sections, i.e. 20% and 60%, it can be observed that both congestion control mechanisms demonstrate comparable fairness. However $RRI_{Adaptive}$ demonstrates improved fairness of 3.4% when considering the long-form simulation that shows fluctuations in congestion levels.

Both congestion control mechanisms applied are based on the LIMERIC algorithm, and this is well known to operate in an effective and fair manner. The explanation for the difference in performance is integration with the underlying scheduling approach. The greatest challenge is in managing stability effectively, which is where *DCC Adaptive* exhibits drawbacks due to its incompatibility with the scheduling approach.

In Figures 6.15a and 6.15b, we consider the CR and the average transmission rates to examine stability performance. In all cases, the CR cannot exceed 0.003 and the transmission rate cannot exceed 10Hz. In particular, *DCC Adaptive* exhibits a higher CR and therefore transmission rate, especially for a target of 20% CBR, which explains the higher measured CBR shown earlier in Fig. 6.10b. For both mechanisms, there are no significant changes in the confidence intervals of both metrics, which are shown in shaded areas. This aligns with the fairness data shown earlier in Table 6.6, with both showing high stability and fairness.

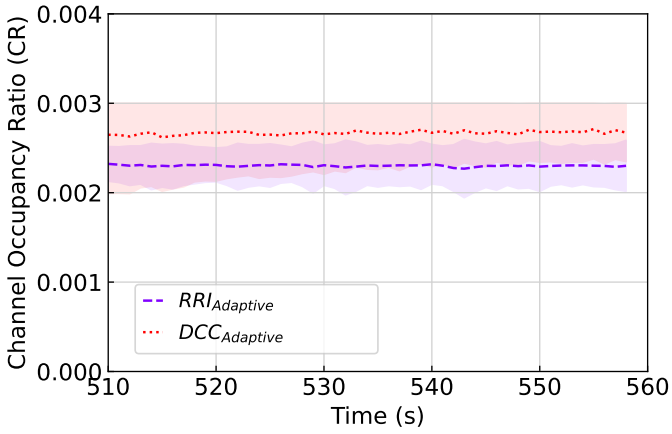

(a) Channel Occupancy Ratio.

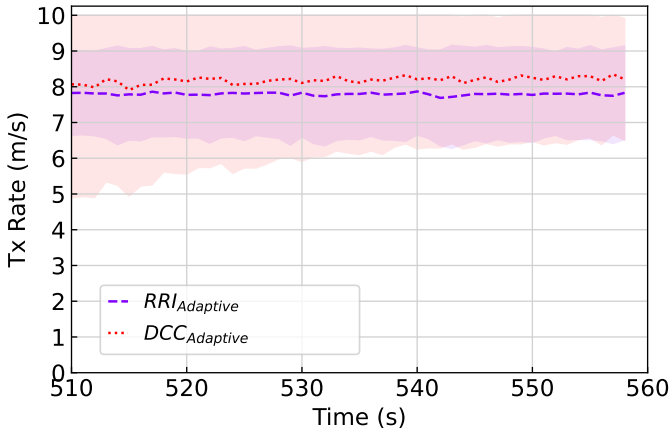

(b) Average Message Rate.

Figure 6.16: Comparative stability statistics at 60% CBR.

For the 60% target CBR, $RRI_{Adaptive}$ exhibits a more stable performance. While both congestion control mechanisms exhibit comparable fairness, they diverge in terms of CR and message rate stability. In Figure 6.16, it is clear that *DCC*

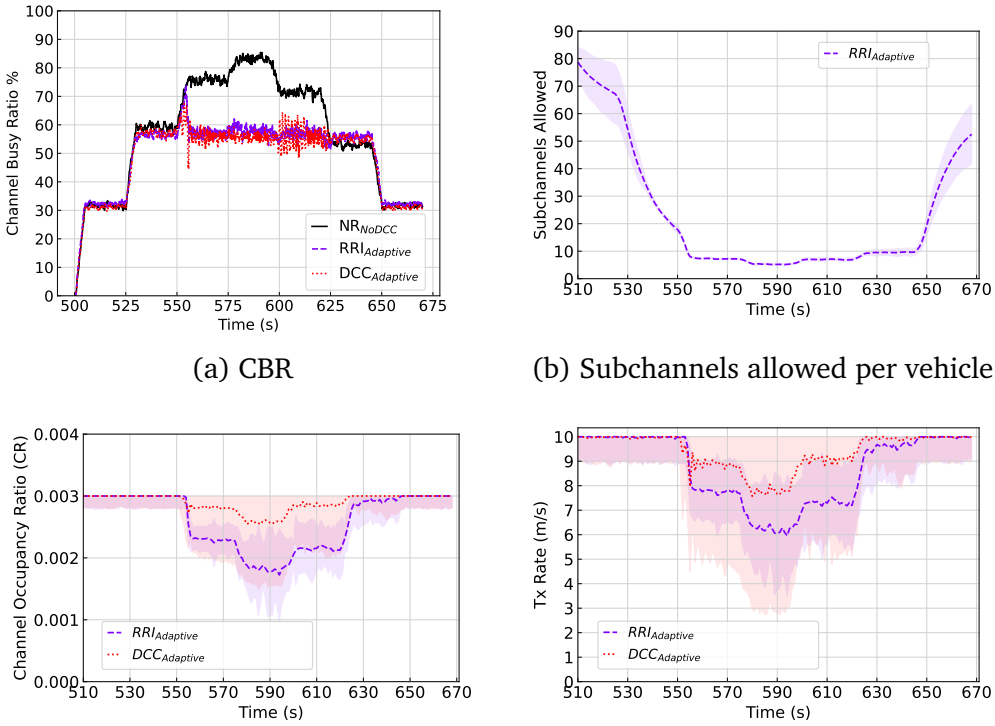

(a) CBR

(b) Subchannels allowed per vehicle

(c) CR

(d) Transmission Rate (Tx)

Figure 6.17: Long-form variable CBR simulation. Comparative stability statistics for a long form simulation exhibiting fluctuating congestion in the channel.

Adaptive exhibits higher CR and message rate. It is clear from the 95% confidence intervals (shaded areas) that *DCC Adaptive* exhibits reduced stability, particularly when examining transmission rates that can fluctuate by up to 5Hz early in the simulation and always exceed the rates maintained by $RRI_{Adaptive}$.

The long-form simulation results are shown in Fig. 6.17a. NR_{NoDCC} (black line) represents the baseline network performance without applying congestion control techniques. When $RRI_{Adaptive}$ and *DCC Adaptive* are applied, both reduce congestion i.e. CBR by over 20% in most densely congested time period between 550s-625s. However, it is clear that $RRI_{Adaptive}$ is more effective at maintaining a stable CBR while converging to the target CBR, especially when congestion is reduced between 600 and 625s.

It is clear that $RRI_{Adaptive}$ is more stable, and this is verified by examining the CR and transmitter rate performance of both approaches, as shown in Fig. 6.17c and Fig. 6.17d. Fig. 6.17c clearly shows how CR is stepped down by each approach and how *DCC Adaptive* exhibits much more variables CR

confidence intervals. This translates into a more extreme divergence in transmission rate as shown in Fig. 6.17d, especially in the most dense time period 590s where the transmission rate can vary from 3Hz to 10Hz. While, $RRI_{Adaptive}$ does show some instability ranging from 4Hz to 8Hz, it still exhibits a significant improvement.

The reason why $RRI_{Adaptive}$ exhibits better stability can be visualised in 6.17b, which shows how the underlying parameter δ is translated into the effective number of subchannels allowable for each vehicle per second. It is clear that initially vehicles can use a significant number of resources, that is, up to 80 subchannels the entire channel. The application running on each vehicle only requires 10 subchannels, demanding 1 subchannel per transmission and 10 transmissions per second. As such, congestion control mechanisms are only applied when δ drops below 0.003 i.e. the maximum CR needed for each vehicle, or fewer than 10 subchannels per second. At this stage, stability is high because the LIMERIC algorithm is strict on the number of resources available for each vehicle. This is demonstrated by the confidence interval 95% (shaded area) tightly bound to the average in this area (dashed line). As such, it is operating effectively and results in better stability performance in terms of CR and average transmission rate, as previously shown. The challenge for *DCC Adaptive* is due to its incompatibility with the SB-SPS scheduling mechanism. This results in inconsistent performance for the channel as a result of increased collisions which can result in increased CBR. Additionally, the introduction of missed transmissions and the consequent additional collisions can impact channel congestion. This can result in oscillations in CBR, where a collision occurs, reducing measured congestion, thereby resulting in rates increasing and thus CBR increasing again, causing a back off in transmission rate. These oscillations can be clearly seen in Fig. 6.17d and are much more severe than any exhibited by $RRI_{Adaptive}$.

The ultimate conclusion with respect to both fairness and stability is that $RRI_{Adaptive}$ can maintain a stable target CBR while achieving high levels of fairness amongst vehicles. Importantly, it can deal with high fluctuations in congestion in the network.

6.6 Conclusion

This chapter has focused on the fourth contribution of this thesis (**C4**) which relates to publications (*1, 3 and 6*), namely congestion control in the context of C-V2X and NR-V2X. The key conclusions from this chapter are as follows.

- A full analysis of the performance of the current generation of congestion control mechanism when applied to C-V2X and NR-V2X, this includes:
  - ETSI DCC approaches both *DCC Reactive* which is the table based mechanism and *DCC Adaptive* the algorthmic approach based on the LIMERIC algorithm [25].
  - The 3GPP CBR to CR-Limit table based mechanism, specifically employing packet dropping as a TRC mechanism.
  - Further, applying the CBR to CR-Limit table based mechanism but to MCS congestion control.
- In this context this chapter has two key findings:
  1. That the current generation of standards introduce a significant number of MTs when applied to either C-V2X or NR-V2X that results in degraded SB-SPS scheduling. As such the reduction in CBR does not enable an improvement in PDR commensurate with the level of packet dropping or delay being applied.
  2. The application of MCS congestion control is ultimately not an effective mechanism as a result of the fact that the use of a high MCS and reducing subchannel occupation is always preferable regardless of congestion. This follows from the analysis in the previous chapter which highlights that reduced subchannel occupation is preferable and this applies regardless of congestion. As such, there are limited scenarios underwhich it would be logical to use a lower order MCS that occupies more subchannels. This conclusion was reached through collaboration with UNAM in a study on which the author of this thesis was the second author on the assocaited paper.
- The primary contribution of this chapter is the proposal of three congestion control techniques based on RRI Adaptation which are referred to as RRI DCC approaches, namely.

1. $RRI_{Reactive}$ a mechanism based on ETSI table based solution *DCC Reactive*.
2. RRI_{CR_Limit} a mechanism based on 3GPP table based CBR to CR-limit approach.
3. $RRI_{Adaptive}$ a mechanism based on ETSI algorithmic solution *DCC Adaptive*.

- The chapter shows that the proposed approaches both improve scheduling performance significantly from the current standard approaches, but also do not significantly degrade on application QoS.
- The $RRI_{Adaptive}$ mechanism represents the first application of an algorithmic approach to SB-SPS and has significant advantages when compared to the current generation of solutons, these include.
  - Avoiding the main challenges of table based solutions, namely the requirement of needing to tune the table to optimise performance.
  - Improved stability and fairness of the applied approach, which is often a challenge with table based solutions.
- This study to the authors knowledge is the first to provide a comprehensive analysis of the fairness and stability of the proposed congestion control technique which is a key area when discussing any such mechanism.

Chapter 7

Conclusion and future outlook

7.1 Introduction

This work has focused on the scheduling approach adopted in both NR-V2X and C-V2X specifically the SB-SPS mechanism. The main topic throughout this thesis was in how this scheduler performs with the current generation of applications, as well as in dense vehicular environments. It is clear from the analysis provided throughout the previous chapters that SB-SPS has considerable challenges both in the management of the current application standards but also in congestion management. There are two key challenges, firstly the impact of IPT changes either based on applications being event-triggered or congestion control techniques intervening in packet generation. Secondly, in the management of variable packet sizes which again ties into the application traffic patterns. These challenges will continue into further iterations of vehicular applications as it is likely to be the case that future applications will also exhibit these event-triggered patterns and variable size. In addition the introduction of a plethora of new applications is likely to result in congestion control becoming ever more important considering the limited C-ITS allocated spectrum.

As such, this chapter will focus on both summing up the contributions in these areas. Then changing the focus to future avenues of investigation in both the near-term and the longer-term.

7.2 Contributions and key findings

The following is a compiled list of the contributions of this work with the key conclusions associated with each of the contributions provided.

- **Contribution 1**: (*Pub 4, Pub 8*)
  Implementation of an open source full-stack standards compliant C-V2X and NR-V2X model, that has been fully validated and is in use by the wider V2X academic community. This model ultimately became the main vehicle by which further investigation was carried out in this thesis.
- **Contribution 2**: (*Pub 2 and poses the underlying research challenge addressed in Pub 1 & Pubs 5 - 8*)
  Thorough investigation of the performance issues faced by SB-SPS when handling commonly occuring vehicular service characteristics (aperiodic packet inter-arrival rates and variable packet sizes) and network management techniques (congestion control). This includes:
  - An in-depth theorised analysis of the challenges posed to the C-V2X and NR-V2X SB-SPS mechanism by MTs and GBs. These impact on the scheduling performance significantly by introducing inconsistencies in the sensing history in the case of MT and increasing contention when GBs are introduced due to the increased number of rescheduling events which occur.
  - A detailed quantitative analysis of the impact of both MT and GB on SB-SPS scheduling performance with a brief quantitative overview provided and a detailed analysis of the consequences of a) aperiodic traffic patterns, b) variable packet sizes in and c) congestion control mechanisms. Across all scenarios, the introduction of MT or GB results in degraded scheduling performance, especially GB which in the case of variable packet sizes is unavoidable. With regards to aperiodic traffic patterns, several alternative mechanisms were investigated, including the NR-V2X dynamic grant mechanism. These were shown to be ineffective or have trade-offs of increased signalling overhead, though they were shown to be effective in mixed aperiodic periodic traffic scenarios with high levels of periodic traffic.
- **Contribution 3**: *(Manuscript under review)*

A machine learning based approach to improved scheduling of aperiodic packet inter-arrival rates (CAMs) within SB-SPS to overcome MT and GB. The proposed model was able to improve the performance of SB-SPS scheduling with CAMs. Considering that CAMs form the core service for V2X applications it then shouldbe possible to employ the dyanmic grant mechanism more effectively for application traffic that exhibits more unpredictable behaviour, such as many DEN use-cases would exhibit.

- **Contribution 4**: *(Pub 1, Pub 3, Pub 6)*
  Proposal of congestion control mechanisms that are compliant with SB-SPS including a quantitative evaluation of their effectiveness. This comprises:
  - A RRI DCC mechanism that reduces congestion while maintaining effective scheduling. Ultimately, the proposed $RRI_{Adaptive}$ mechanism was shown to manage congestion effectively, without incurring an undue loss of QoS, and to maintain stable channel conditions and high fairness between vehicles. The ability to do so with an algorithmic approach rather than being bound to a table-based mechanism represents a significant improvement on the current state of the art. Importantly this represents the first study that investigated fairness and stability in the context of C-V2X and NR-V2X.
  - An analysis of the effectiveness of MCS adaptation as a means of congestion control, which was found to be ineffective.
- **Contribution 5**: *(Pub 5)*
  This minor contribution analyses the role of MCS adaptation in assisting SB-SPS to effectively schedule variable sized packets. This includes:
  - A quantitative analysis of the performance implications of MCS adaptation versus fixed MCS selection (widely assumed in academic literature). Importantly MCS adaptation within the grant ensures the best SB-SPS performance as it avoids unnecessary reductions in transmission robustness or over-provisioning of subchannels. As such, further academic studies in this space should avoid the assumption of a fixed MCS as often the choice of MCS will not be able to manage each packet size effectively.
  - A quantitative evaluation highlighting the importance to consider

MCS ranges beyond those defined in the existing standards. The outcome of this investigation should inform the standardisation efforts of C-V2X and NR-V2X as current limits result in reduced overall scheduling performance. By extending the MCS range to beyond 11 it enables SB-SPS to manage a more dynamic range of packet sizes more effectively.

- An alternative MCS selection strategy to minimise SB-SPS grant breaks. The current approach adopted of reducing subchannel occupation when possible is preferable, and the other over-provisioning techniques are ineffective. This is due to the fact that while reducing or eliminating GBs is possible the increased subchannel occupation as a result of the over-provisioning required to do so results in a greater number of collisions. This then leads to the reduced performance of SB-SPS overall in these cases, as such, scheduling based on the current packet size, while applying MCS adaptation and ensuring a wide range of possible MCS is the best configuration of SB-SPS.

These have formed the contributions of this work and now the focus will shift to the future work which may have been completed had the project continued, as well as longer-term research outlooks which may form the basis of a future project.

7.3 Future Work

There are a number of short-term implementation extensions and longer-term open research challenges that follow from this work.

Firstly, addressing the shorter-term implementation extensions; it can be envisaged that OpenCV2X can be extended in a number of ways to advance and maintain its relevance to the vehicular communications community:

- Integration of further NR-V2X features such as re-evaluation, preemption, unicast and group cast communication paradigms and finally variable sub-carrier spacing mechanisms.
  - The first mechanism allows for NR-V2X SB-SPS grants to be cancelled before transmission of the first packet when it is judged that another vehicle has already selected the resource.

- Preemption enables vehicles to relinquish a selected resource when it determines that a higher priority message requires the resource. This feature is relevant in the context of applications requiring ultra reliable low latency communication.
- Unicast and Groupcast communication patterns should be introduced into the model to replicate the feature set introduced in the NR-V2X release. This would require the introduction of the 2 stage SCI message which enables these communication paradigms. The 2 stage SCI message allows for the indication of the group or singular receiver in the first stage and the second stage then acts in a similar manner to the current SCI message with specific message information such as the MCS used for transmission. In addition such a change would require the introduction of HARQ retransmissions for these group or unicast message patters. HARQ is already part of the SimuLTE standard but any changes made to the process for the sidelink would need to be introduced.
- Variable sub-carrier spacing enables SB-SPS to adapt the transmission time at the cost of resource occupancy, this enables low latency applications. Modelling this accurately and enabling it to be configurable with a range of MCS is a significant undertaking.
- The final element to introduce to the model is actually in updating the integration, since the initial implementation of OpenCV2X a new version of SimuLTE has been introduced Simu5G [109]. Porting the model to this newer model will introduce a number of improvements on the LTE side which should enable users to integrate interesting LTE based elements to their simulations such as edge enabled by the cellular network. Updating to Simu5G will ensure that the latest LTE standards are used when running such simulations.

Next we discuss five open/longer-term research opportunities that leverage or build upon the research in this thesis.

7.3.1 Multi service management at the facilities layer

The facilities layer is a construct in C-ITS that refers to an intermediate layer between the applications and the radio access technology [110]. It is

envisaged as a midddleware management layer aiding in the various functions of a C-ITS node, from application, communication and information support management. It has a broad remit for aiding the C-ITS station to operate more effectively.

The next generation in vehicle V2X services will have diverse QoS requirements. Some packets must be delivered with high reliability, while others have low latency requirements, with their usefulness degrading over time. As investigated in this thesis, these services may also differ in how packets are generated (event-driven or periodic) and variability in packet size, which has impacts for the scheduler. Using static QoS classes can starve lower priority services, potentially diminishing their usefulness when actually transmitted.

Research to intelligently characterise and prioritise based on the service characteristics and condition of the network by utilising contextual factors like usefulness and urgency while balancing fairness across services is needed.

This proposal aligns well with recent standards [111] by ETSI for Release of 2 C-ITS standards. This proposes that applications be prioritised at the facilities layer by negotiation or optimisation of the application requirements (called FCPs) and knowledge of the lower layers (called FCLs). Essentially, this proposes an intelligent cross-layer design to fairly and intelligently scheduler packets.

7.3.2 DCC at the facilities layer

Current C-ITS Rel. 1 systems manage congestion based on access layer DCC mechanisms. These are totally incompatible with the C-V2X and NR-V2X scheduler. In this thesis we proposed a standard compliant way of solving these scheduling challenges. However, such schemes are based on crude packet drop/delay (TRC) mechanisms with no knowledge of the potentially negative impact this has on a given service or whether the packet that is eventually transmitted is still useful. As such, there is a need to intelligently manage scarcely available radio resources amongst apps, enable an app to intelligently reduce generated traffic or be able to dynamically react to congestion by offloading traffic to a different channel. This should be investigated in a cellular context.

7.3.3 Multi-RAT load balancing

It can be expected that future C-ITS systems (OBU + RSUs) may have multiple radios. Allowing such systems to co-exist presents a challenge that is as yet not clearly solved. One of the primary challenges that can be envisaged is the fact that when managing multiple different technologies, there is interference between technologies, that is, C-V2X interferes with ITS-G5 transmissions and vice versa. If technologies have no means of interpreting the other's transmissions and do not take them into account, the system will encounter increased interference that will likely severely impact performance. There is the option of dedicating specific channels for each technology and avoiding this challenge, but the C2CCC has conducted some investigation on co-channel co-existence methods [112]. There is significant scope for investigating both these co-channel co-existence challenges building on the work of C2CCC but further investigating the means of load balancing across technologies and ultimately proposing an improved overall system performance through this multi-RAT approach.

With respect to load balancing, the primary challenge is to identify for each application type which radio technology is best suited to its requirements. This could be based both on the application characteristics e.g. generation rate and packet sizes, but also on channel conditions such as congestion. Finally, any mechanism would need to be able to dynamically manage overtime with ever changing channel conditions.

7.3.4 Impact of retransmissions and aperiodic application traffic.

This is a key area of interest when it comes to focusing on aperiodic application patterns and the dynamic grant mechanism introduced in NR-V2X. In this study there was no focus on retransmission but the standard allows for it and dynamic grant employs retransmissions to effectively change its operation. In the case where retransmissions are enabled dynamic grant operates like a normal SB-SPS grant, where the retransmissions make up the subsequent reservations in the grant. As such, increasing the number of retransmissions improves both the redundancy of each transmission ensuring that each packet will be received by more neighbouring vehicles, but also improves the scheduling performance of dynamic grant. This should eliminate

highlighted issue where dynamic grant due to it being grantless and having no retransmissions would operate in a random manner as there was no sensing history to use to inform the selection of subchannels in scenarios with only aperiodic traffic.

The key issue to investigate though are the trade-offs in introducing retransmissions. While the individual packet will be sent in a more robust manner due to the retransmission there is a significant increase in channel usage as a result of those retransmissions which can degrade overall network performance. In dense scenarios this increase in channel usage may become prohibitive and result in degraded performance for all traffic and this may be exasperated when there are numerous services running simultaneously. This aspect ties in nicely with the work focusing on the facilities layer management, as it may be possible to intelligently use retransmissions in low density scenarios for the benefits in terms of robustness and improve scheduling performance with dynamic grant. Then in high density scenarios reducing or eliminating retransmissions can occur to ensure the best channel conditions across a diversity of different scenarios and maintain the best performance for the various applications. This would be a significant area of investigation to balance all these competing elements and to implement a system that successfully manages retransmissions across the various and dynamic scenarios that are possible in vehicular networks.

7.3.5 A focus on Cooperative Perception

CP is a key service for the next generation of vehicular safety applications, it will be key in improving the safety of VRUs such as cyclists and pedestrians. It operates by vehicle using onboard sensors to sense the environment around them and then communicate that information to other neighbouring vehicles to ensure that they have a more complete understanding of their environment. There are numerous challenges with such an application from a communications perspective which are similar to those explored in the CA service this study has focused on. The key issues are the same in that these packets will exhibit variable size as well as an aperiodic arrival rate. This is due to the fact that the packets are based on sensed objects in the environment, which makes it highly unpredictable and depending on the environment there may be significant fluctuations in the number of sensed objects.

This poses some very interesting challenges which are difficult to address to

take them in turn firstly with aperiodic arrival rate. In this study a machine learning approach was employed to improve the performance of the CA service by predicting the short term trajectory of the vehicle and using this to inform the upcoming RRI. This is much more challenging with the CP service, due to the fact that it is based on sensed objects in the vicinity of the vehicle as such it requires the prediction of when objects may enter sensing range. In addition it may require predicting how those sensed object might move and when this will result in a trigger condition and require transmission. This is very challenging as it would be dependent on the environment and is inherently unpredictable unless through understanding trends in the individual areas. It may be possible to use reinforcement learning in areas where a vehicle commonly travels to understand the typical trends such as usual number of cyclists in an area at particular periods in the day and have this inform the RRI. Still this would be a significant challenge and there is considerable scope for investigation on this topic.

The second significant challenge is that of packet size; for the CA service it has been shown that there can be significant fluctuations in packet size depending on the recorded model. Still the sizes of CA packets are somewhat bounded based on the configurable fields in the message and a typical size model can be generated as shown in other research [59]. This would be more challenging for the CP service as again it is based on the sensed object in the environment, which can vary significantly. This can result in cases with a dense number of sensed object in packets becoming prohibitively large and causing significant increases in resource usage and interrupt the SB-SPS grant mechanism as we have shown for the CA service. There is significant scope for investigation in this area as well both in terms of intelligent MCS and subchannel configuration but also in determining the correct sensed objects to include in the message to reduce packet size by eliminating redundant information without decreasing the safety of those VRUs.

Appendix A

Overview of results generation and analysis

It is of importance to fully explain and explore the means by which results are generated and the ultimate statistical confidence on which all subsequent results are based. This is vital for reproducability of results, an underlying tenet of the OpenCV2X contribution. Throughout the work, all experiments have been completed across five simulation runs. For each graphed configuration, every recorded result representing a batch mean across those simulations. Additionally, it is the case that each simulation is of at least 12 second length with the first 2 seconds disregarded in the generated results. Taking the packet delivery rate as the primary example and the main statistic of interest throughout this work, this represents the total number of decoded packets by each vehicle over the total received. These values are then categorised based on distance into 10 metre buckets and graphed with this granularity; the batch means approach results in a 95% confidence interval on PDR values with a difference of 0.05% . These confidence intervals are not represented in the graphs for the sake of readability.

Appendix B

Standardised and recommended congestion control mechanisms.

Table B.1: 3GPP C-V2X Congestion Control - Maximum CR limit per CBR range and packet priority [1].

CBR	**CR limit**
CBR $\leq$ 0.65	no limit
0.65 $<$ CBR $\leq$ 0.675	1.6e-3
0.675 $<$ CBR $\leq$ 0.7	1.5e-3
0.7 $<$ CBR $\leq$ 0.725	1.4e-3
0.725 $<$ CBR $\leq$ 0.75	1.3e-3
0.75 $<$ CBR $\leq$ 0.775	1.2e-3
0.8 $<$ CBR $\leq$ 0.825	1.1e-3
0.825 $<$ CBR $\leq$ 0.85	1.1e-3
0.85 $<$ CBR $\leq$ 0.875	1.0e-3
0.875 $<$ CBR	0.8e-3

Table B.2: ETSI C-V2X Congestion Control - Maximum CR limit per CBR range and packet priority [2].

CBR	**Priority 1-2**	**Priority 3-5**	**Priority 6-8**
0 $\leq$ CBR $\leq$ 0.3	no limit	no limit	no limit
0.3 $<$ CBR $\leq$ 0.65	no limit	0.03	0.02
0.65 $<$ CBR $\leq$ 0.8	0.02	0.006	0.004
0.8 $<$ CBR $\leq$ 1	0.002	0.003	0.002

$$t_{\text{go}} = t_{\text{pg}} + min(max(\frac{T_{\text{on}_{\text{pp}}}}{\delta}, 25ms), 1s) \tag{B.1}$$

Table B.3: ETSI Access Reactive DCC as specified in [3].

CBR	**State**	**Packet Rate**	$\mathbf{T_{off}}$
CBR $<$ 0.3	Relaxed	10Hz	100ms
0.3 $\leq$ CBR $\leq$ 0.40	Active 1	5Hz	200 ms
0.40 $\leq$ CBR $\leq$ 0.50	Active 2	2.5Hz	400 ms
0.50 $\leq$ CBR $\leq$ 0.60	Active 3	2Hz	500 ms
CBR $>$ 0.60	Restrictive	1Hz	1000ms

Algorithm B.1 Algorithm for δ calculation

$CBR_{\text{ITS-S}} \leftarrow 0.5 \times CBR_{\text{ITS-S}} + 0.5 \times ((last\_CBR + 2nd\_Last\_CBR)/2)$
if $CBR_{\text{target}} - CBR_{\text{ITS-S}}) > 0$ **then**
    $\delta_{\text{offset}} \leftarrow min(\beta \times (CBR_{\text{target}} - CBR_{\text{ITS-S}}), G_{+\text{max}})$
else
    $\delta_{\text{offset}} \leftarrow max(\beta \times (CBR_{\text{target}} - CBR_{\text{ITS-S}}), G_{\text{-max}})$
end if
$\delta \leftarrow (1 - \alpha) \times \delta + \delta_{\text{offset}}$
if $\delta > \delta_{\text{max}}$ **then**
    $\delta \leftarrow \delta_{\text{max}}$
end if
if $\delta < \delta_{\text{min}}$ **then**
    $\delta \leftarrow \delta_{\text{min}}$
end if

$$t_{\text{go}} = t_{\text{pg}} + min(max(\frac{T_{\text{on}_{\text{pp}}}}{\delta} \times \frac{t_{\text{go}} - t}{t_{\text{go}} - t_{\text{pg}}}, 25ms), 1s) \tag{B.2}$$

Appendix C

Additional implemented ML models for CAM IPT prediction.

Throughout the process of implementing ML approach for CAM IPT prediction there were a number of iterations of different models with various approaches adopted that incrementally improved prediction accuracy coming to the final proposal in this thesis. This appendix catalogues a number of the models attempted and highlights the pitfalls of each and provides some indicative results as to their performance and how this informed the next iteration of models. What is not highlighted here is that for each implemented model a process of model searching and refinement was completed through a baysian search approach, this being the selection of elements such as LSTM or GRU based models, number of layers, neurons per layer, learning rates and additional dense layers. In each of the models provided below this approach was adopted and the best version from this process is ultimately discussed. In additional some information with regards to data preparation will be highlighted as this is most relevant when transitioning from the first proposed model to subsequent iterations.

A key element to note with all ML approaches is that ultimately the proposed solutions act as a blackbox, it is difficult to precisely diagnose why the model might be performing poorly. There are common issues that can occur such as limited data availability which might not reflect the real world scenarios resulting in the model having poor generalisation. Alternatively over training or models with too many layers and neurons can become so tied to their datasets and again will generalise poorly as predictions will be entirely based

on training data.

In the subsequent discussed models these issues are addressed through model selection with baysian search though the earliest models would have had smaller datasets, further into the process additional datasets were introduced to bring a broader range of scenarios and greater diversity. This would result in improved model performance for subsequent iterations of CAM IPT models.

A final topic before discussing models is the datasets involved. The final datasets are shown in Table 4.4, initial iterations of the model would not have included the full dataset and would have generated the dataset in different forms. For example the initial model which directly predicts the RRI included information on when CAMs were triggered and the conditions underwhich they were triggered. The final model on the other hand would have included more datasets from different scenarios to include a more diverse set of scenarios. In addition the dataset generated included additional information about the lead vehicle such as its heading, speed, acceleration and other elements which could be used to infer the future trajectory of the vehicle. As such, overtime the data will change and this will be highlighted in the analysis of each model and form an element of the challenge for each model.

The following simulations use the same set of parameters as described in Table 3.1 with the Cork real road network shown in Fig. 5.11c, which gives the best indication of overall performance of the model considering diverse changes in heading and speed. In all cases for comparison the periodic result is included as a baseline towards which the result should trend

C.1 Direct RRI Prediction

The first approach adopted is the most obvious choice of directly predicting the CAM IPT. This was achieved by generating a dataset in which the trajectory information tied to the CAM trigger rules such as speed change, heading change, position change and elapsed time were recorded. In addition, we included information on why a CAM was triggered such as due to heading, speed, or position changes.

Indicative results of the performance of this model are provided in Fig. C.1. In this case there is a significant difference in terms of performance of the model with significantly degraded performance in terms of PDR as shown in Fig. C.1a

while exhibiting lower overall CBR shown in Fig. C.1b.

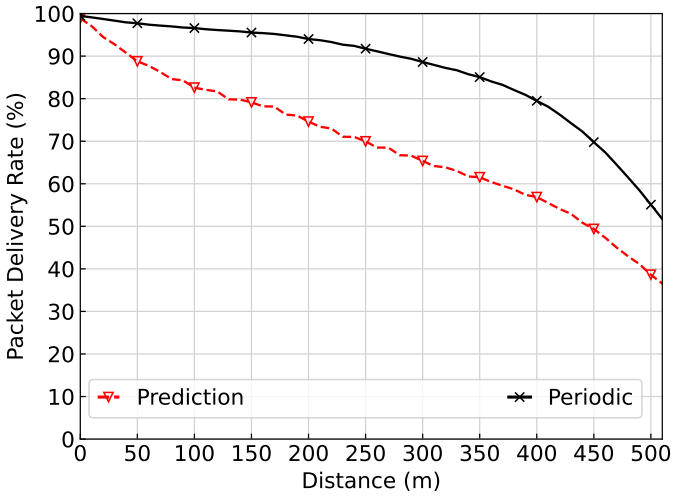

(a) Packet Delivery Rate.

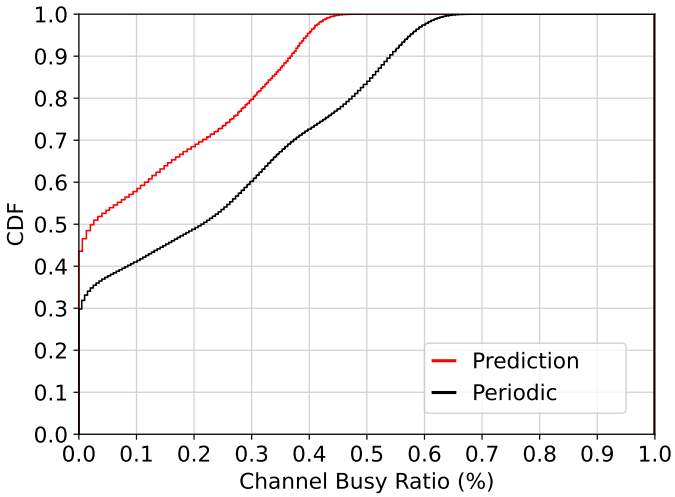

(b) Channel Busy Ratio.

Figure C.1: Performance of direct RRI prediction model.

It is evident that the model is not overall successful in determining the RRI accurately and as such mimic the periodic performance of SB-SPS. From investigation there were a few factors that influenced this poor performance. The first being based on the dataset, by only including data surrounding the trigger conditions of a CAM this significantly reduced the amount of data that could be used in training. In addition this required more time to generate a comprehensive dataset as it required either additional simulation overhead to track CAM trigger conditions or post processing to determine the trigger conditions overtime. The second element was down to the accuracy of the

model in predicting heading which became the most challenging aspect to predict for all the models.

There were a number of outcomes from the analysis of this initial model. The first being the decision to focus entirely on trajectory prediction and use this in the subsequent prediction in determining the possible trigger time and condition for the upcoming CAM transmission. This overcame one of the main challenges with the RRI model specifically in that it allowed for a more extensive dataset to be developed by reducing the need to consider CAM triggers throughout and focus purely on the trajectory information for each vehicle. Additional datasets were introduced for the subsequent models extending the number of scenarios to the list compiled in Table 4.4, this was with the aim of introducing a more diverse data set. This would ensure both more data which helps in all machine learning approaches but additionally to try to overcome challenges with predicting heading by introducing scenarios with greater diversity of travel conditions.

C.2 First trajectory prediction approach

The second model explored was the initial model which introduced trajectory based prediction, in this case the model predicted future heading and speed for the subsequent second. This then was also translated to positional change and those elements when compared with the CAM trigger rules were used in determining the RRI for the current transmitted CAM. The indicative results in this case are highlighted in Fig. C.2, with PDR improving significantly from the previous scenario but still with significant under performance when compared with periodic SB-SPS as shown in Fig. C.2a. It is noteworthy that again CBR performance is lower than that of the periodic scenario as shown in Fig. C.2b.

In this case the improvements but pitfalls are the same as in the previous scenario, primarily being in heading prediction. Importantly performance was improved by the introduction of new scenarios but significant error was still occurring with heading prediction. The key issue identified in this case was in the way in which heading information was being handled and the unit circle itself. It is difficult for ML models to determine relationships where the values are not clearly linked in the case of the unit circle and heading 359° is only a single degree from 0° rather than 359° away. This ultimately led to inaccuracies in the prediction of future heading as the model had difficulty in

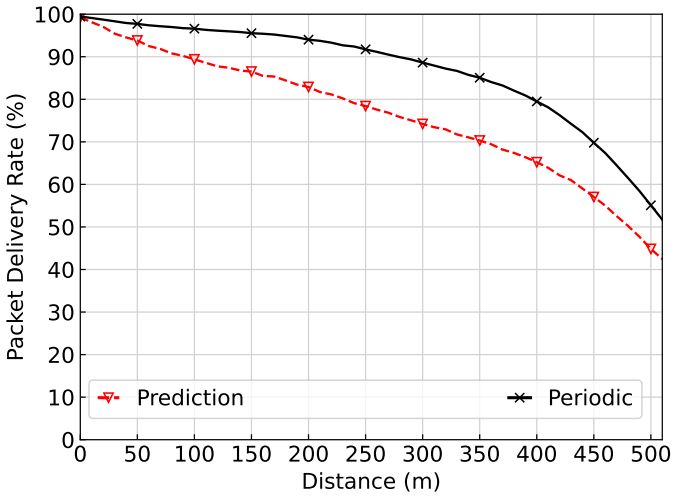

(a) Packet Delivery Rate.

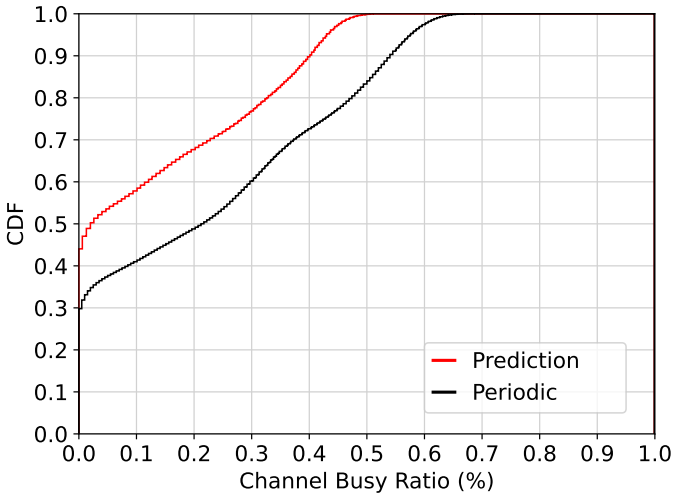

(b) Channel Busy Ratio.

Figure C.2: Performance of initial trajectory based model.

accurately modelling elements of this relationship.

This lead to the changing of the model in subsequent iterations the specific change being in the handling of heading information in prediction. To overcome this specific challenge all heading information was translated into 2 elements the sine and cosine of the heading value. This could then be translated back into the actual heading using the atan2 function but the relationship between sine and cosine of heading is more easily managed in a machine learning model as the values are continuous between -1 and 1 rather than reverting between 0 and 360.

C.3 Accounting for heading changes

As outlined this version of the model adopted the same trajectory prediction approach for determining the RRI but accurately modelled the heading relationship by splitting its components into sine and cosine such that the model can more easily track the subtle changes in heading.

This model proved decidedly more accurate as shown in Fig. C.3. In this case the PDR performance is significantly better than the previous iterations and closely tracks with the periodic performance as seen in Fig. C.3a. Still there is an issue with the channel usage with CBR still tending to be lower than the periodic scenario highlighted in Fig. C.3b, this trend has remained throughout the different iterations of the model. This ultimately means that performance is still somewhat degraded from the periodic scenario and that the ML model is introducing some latency to the transmission of packets.

The main outcomes from this analysis is that the model proposed was highly effective in maintaining good performance in this scenario with some minor performance drawbacks when compared to the periodic scenario. There were significantly fewer inaccuracies in prediction and these were introduced again solely through heading errors. In addition the ML model introduces some additional latency to the transmission of packets most likely due to some RRIs being set slightly longer than the periodic scenario due to inaccurate predictions i.e. predicting an RRI of 200ms which was ultimately 100ms. As this model did not give ideal performance a final iteration was investigated.

C.4 Lead vehicle information incorporated in prediction

The final model introduced saught to fully overcome the challenges with heading prediciton and to do so incorporated information about lead vehicles. In scenarios where a vehicle had a vehicle which was travelling in the same direction as the ego vehicle and was clearly ahead of it, that lead vehicles information was incorporated in the prediction in an attempt to discern information about upcoming road and traffic conditions. This is exclusively achieved through exploiting the information available in transmitted CAMs from that lead vehicle which contain all the necessary information such as lead vehicle speed, heading and position.

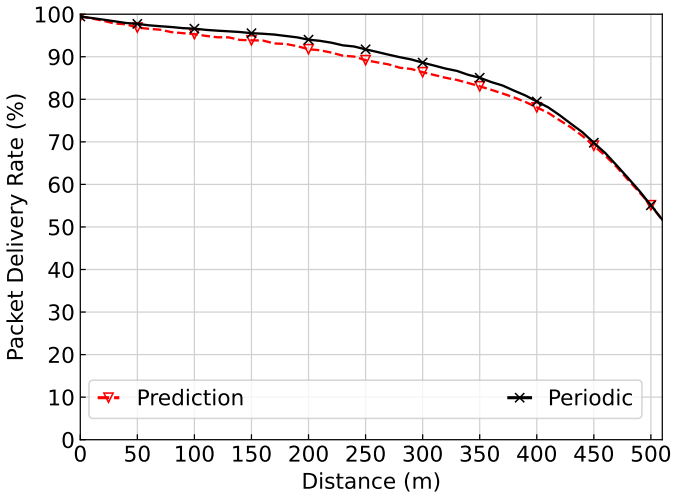

(a) Packet Delivery Rate.

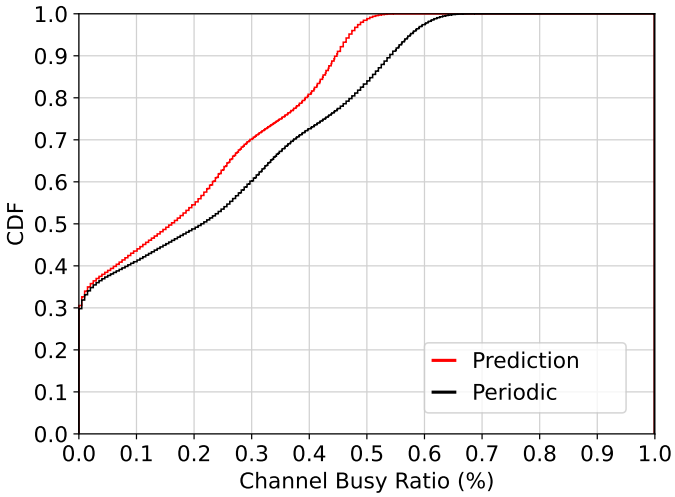

(b) Channel Busy Ratio.

Figure C.3: Performance of trajectory based model accounting for heading.

As shown in Fig. C.4, the performance in this scenario both does not improve the CBR differential as seen in Fig. C.4b it also degrades PDR performance in this case highlighted in Fig. C.4a. Evidently this approach did not achieve the goal it set out to and elimate the final issues with heading error.

The ultimate cause of the issues in performance in this model is actually exasperating the issue that it sought to overcome, when using the leader based model heading prediction degraded significantly from the previous iteration. Still performance overall is better than the initial trajectory based approach indicating improvements from that version but the introduction of the lead

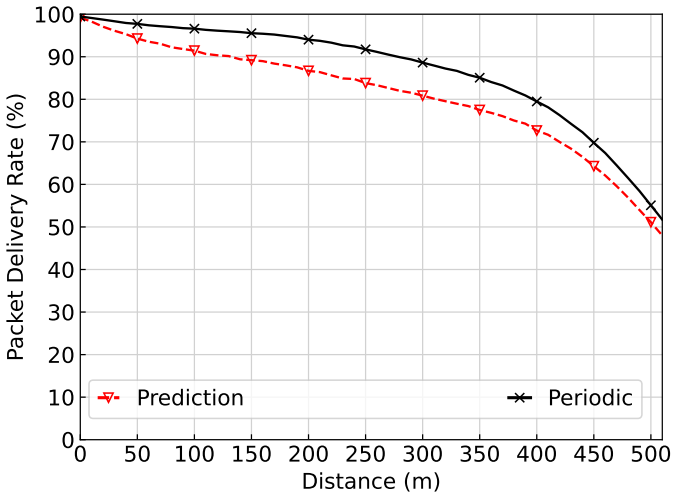

(a) Packet Delivery Rate.

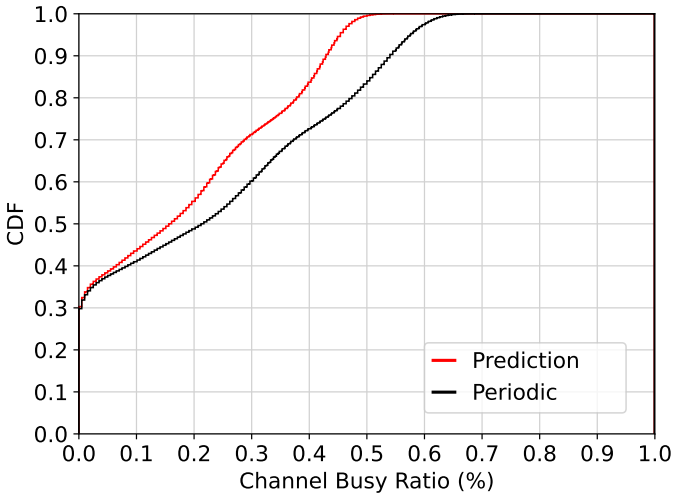

(b) Channel Busy Ratio.

Figure C.4: Performance of ML model incorporating lead vehicle information.

vehicle information ultimately lead to more error in heading. The diagnosis of this is more difficult to refine, logically understanding the dynamics of the vehicle directly in front of the ego vehicle would be highly relevant to the ego vehicles future position and dynamics. As such, the only possible reason behind the degraded performance is in the data used in training. It is likely that the datasets did not include a good number of scenarios with a lead vehicle but also may have introduced noise to the training data. A key example being in a city scenario where the lead vehicle may change frequently i.e. turning off into a side street while the ego vehicle continues straight. This would lead to inconsistencies in the training data where highly dynamic

scenarios for the leader result in minimal changes for the ego vehicle. Ultimately such an approach should be able to improve performance especially considering scenarios such as highways where the lead vehicle might change more infrequently but due to limitations with time to further refine this model it was decided to revert to a model which did not include the lead vehicle data in prediction.

References

[1] Qualcomm, “R1-1611594 Congestion control for V2V,” tech. rep., 2016. 3GPP TSG-RAN WG1 No.87.

[2] ETSI TS 103 574 V0.3.1, “Intelligent Transport Systems (ITS); Access layer part; Congestion Control for the Cellular: V2X PC5 interface,” vol. 1, pp. 1–12, 2018.

[3] ETSI, “ETSI TS 102 687 - V1.2.1 - Intelligent Transport Systems (ITS); Decentralized Congestion Control Mechanisms for Intelligent Transport Systems operating in the 5 GHz range ;,” *ETSI standard*, vol. 2, pp. 1–14, 2018.

[4] European Climate, Infrastructure and Environment Executive Agency, “Eu road safety: Towards "vision zero",” Nov 2022. Available at https://cinea.ec.europa.eu/publications/eu-road-safety-towards-vision-zero_en.

[5] Irish Road Safety Authority, “Provisional review of fatal collisions, 1 january to 31 december 2021,” 2022. Available at https://www.rsa.ie/docs/default-source/road-safety/r2---statistics/provisional-reviews/provisional-review-of-fatal-collisions-2021.pdf?Status=Master&sfvrsn=f3461509_3.

[6] European Commission, “A european green deal,” Dec 2019. Available at https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/european-green-deal_en.

[7] The government of Ireland; Department of Transport, “Transport vision in climate action plan will transform how we travel over the coming 7 years,” Dec 2022. Available at https://www.gov.ie/en/press-release/d3341-transport-vision-in-climate-action-plan-will-transform-how-we-travel-over-the-coming-7-years/.

[8] The government of Ireland, "Project ireland 2040; national planning framework," Sep 2017. Available at https://www.npf.ie/wp-content/uploads/Project-Ireland-2040-NPF.pdf.

[9] D. A. G. Alina Pinkelnig, "Annual pilot overview report 2021," 8 2022. C-Roads EU.

[10] ETSI, *Intelligent Transport Systems (ITS); Access layer specification for Intelligent Transport Systems using LTE Vehicle to everything communication in the 5,9 GHz frequency band (V1.1.1)*, Nov. 2018.

[11] S. o. A. E. SAE International, "On-board system requirements for v2v safety communications. Standard Doc J2945/1," 03 2016.

[12] 3GPP, *Study on LTE-based V2X services (v14.0.0, Release 14)*, July 2016. 3GPP TR 36.885.

[13] 3GPP, "Release 16 description; summary of rel-16 work items (v1.0.0 , release 16))," Dec. 2020. TR 21.916.

[14] 3GPP, *Proximity-based services (ProSe); Stage 2, Vx14.1.0*, Mar 2017. 3GPP TS 23 303.

[15] ETSI, "En 302 637-2 intelligent transport systems (its); vehicular communications; basic set of applications; part 2: Specification of cooperative awareness basic service (v1.4.1)," Jan. 2019.

[16] ETSI, "Etsi en 302 637-3; intelligent transport systems (its); vehicular communications; basic set of applications; part 3: Specifications of decentralized environmental notification basic service v1.3.1," April. 2019.

[17] 3GPP, "Tr 36.213 physical layer procedures (v14.2.0, release 14)," *3GPP, Tech. Spec.*, Dec. 2017.

[18] M. H. C. Garcia, A. Molina-Galan, M. Boban, J. Gozalvez, B. Coll-Perales, T. Sahin, and A. Kousaridas, "A Tutorial on 5G NR V2X Communications," *IEEE Communications Surveys & Tutorials*, p. 1, 2021.

[19] 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. PP, pp. 1–1, 2021.

[20] Ericsson, *Resource allocation procedures for Mode 2*, March 2019. 3GPP TSG RAN WG1 No.96 R1-1903166.

[21] LG-Electronics, *Discussion on resource allocation mechanism for NR V2X*, March 2019. 3GPP TSG RAN WG1 No.96 R1-1901933.

[22] ETSI, *LTE; Evolved Universal Terrestrial Radio Access (E-UTRA) Physical layer procedures (V14.17.0 Release 14)*, June. 2021.

[23] 3GPP, *User Equipment (UE) radio transmission and reception (v14.0.0, Release 14)*, Mar. 2017. TR 38.786.

[24] NXP, *Analysis of CAMP Report on the C-V2X Performance Assessment Project*, 2019.

[25] G. Bansal, J. B. Kenney, and C. E. Rohrs, "Limeric: A linear adaptive message rate algorithm for dsrc congestion control," *IEEE Transactions on Vehicular Technology*, vol. 62, no. 9, pp. 4182–4197, 2013.

[26] A. Rostami, S. Member, B. Cheng, S. Member, G. Bansal, K. Sjoberg, M. Gruteser, and J. B. Kenney, "Channel Congestion Control Approaches," vol. 17, no. 10, pp. 1–14, 2016.

[27] J. B. Kenney and G. Bansal, "Controlling Congestion in Safety-Message Transmissions," no. December, pp. 1–8, 2013.

[28] R. Molina-Masegosa and J. Gozalvez, "Lte-v for sidelink 5g v2x vehicular communications: A new 5g technology for short-range vehicle-to-everything communications," *IEEE Vehicular Technology Magazine*, vol. 12, no. 4, pp. 30–39, 2017.

[29] O. Amador, I. Soto, M. Uruena, and M. Calderon, "GoT: Decreasing DCC queuing for CAM messages," *IEEE Communications Letters*, vol. 24, no. 12, pp. 2974–2978, 2020.

[30] 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, 2019.

[31] R. Molina-Masegosa, J. Gozalvez, and M. Sepulcre, "Configuration of the c-v2x mode 4 sidelink pc5 interface for vehicular communication," in *2018 14th International Conference on Mobile Ad-Hoc and Sensor Networks (MSN)*, pp. 43–48, 2018.

[32] R. Molina-Masegosa and J. Gozalvez, "System level evaluation of lte-v2v mode 4 communications and its distributed scheduling," in *2017 IEEE 85th Vehicular Technology Conference (VTC Spring)*, pp. 1–5, 2017.

[33] 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.

[34] A. Bazzi, A. O. Berthet, C. Campolo, B. M. Masini, A. Molinaro, and A. Zanella, "On the design of sidelink for cellular v2x: A literature review and outlook for future," *IEEE Access*, vol. 9, pp. 97953–97980, 2021.

[35] A. Bazzi, "Congestion control mechanisms in ieee 802.11p and sidelink c-v2x," in *2019 53rd Asilomar Conference on Signals, Systems, and Computers*, pp. 1125–1130, 2019.

[36] 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.

[37] S. Bartoletti, B. M. Masini, V. Martinez, I. Sarris, and A. Bazzi, "Impact of the Generation Interval on the Performance of Sidelink C-V2X Autonomous Mode," *IEEE Access*, vol. 9, pp. 35121–35135, 2021.

[38] F. Romeo, C. Campolo, A. Molinaro, and A. O. Berthet, "Asynchronous traffic on the sidelink of 5g v2x," in *2020 IEEE International Conference on Communications Workshops (ICC Workshops)*, pp. 1–6, 2020.

[39] F. Romeo, C. Campolo, A. Molinaro, and A. O. Berthet, "Denm repetitions to enhance reliability of the autonomous mode in nr v2x sidelink," in *2020 IEEE 91st Vehicular Technology Conference (VTC2020-Spring)*, pp. 1–5, 2020.

[40] 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.

[41] ETSI, *Congestion Control Mechanisms for the C-V2X PC5 interface; Access layer part (V1.1.1)*, Nov. 2018.

[42] A. Mansouri, V. Martinez, and J. Härri, "A first investigation of congestion control for LTE-V2X mode 4," in *2019 15th Annual Conference on Wireless On-demand Network Systems and Services*, pp. 56–63, 2019.

[43] A. Bazzi, "Congestion control mechanisms in IEEE 802.11p and sidelink C-V2X," in *2019 53rd Asilomar Conference on Signals, Systems, and Computers*, pp. 1125–1130, 2019.

[44] Qualcomm OnQ Blog, "C-v2x performance under aperiodic messages of variable size," Feb 2021. Available at `https://www.qualcomm.com/news/onq/2021/02/c-v2x-performance-under-aperiodic-messages-variable-size`.

[45] 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?," in *2022 IEEE 23rd International Conference on High Performance Switching and Routing (HPSR)*, pp. 183–188, 2022.

[46] J. Yan and J. Härri, "Mcs analysis for 5g-nr v2x sidelink broadcast communication," in *2022 IEEE Intelligent Vehicles Symposium (IV)*, pp. 1347–1352, 2022.

[47] A. Reyhanoglu, E. Kar, F. E. Kumec, Y. S. Can Kara, S. Karaagac, B. Turan, and S. Coleri, "On the reliability analysis of c-v2x mode 4 for next generation connected vehicle applications," in *2022 IEEE 96th Vehicular Technology Conference (VTC2022-Fall)*, pp. 1–5, 2022.

[48] Y. Yao, X. Zhou, and K. Zhang, "Density-aware rate adaptation for vehicle safety communications in the highway environment," *IEEE Communications Letters*, vol. 18, no. 7, pp. 1167–1170, 2014.

[49] Y. Yao, X. Chen, L. Rao, X. Liu, and X. Zhou, "LORA: Loss differentiation rate adaptation scheme for vehicle-to-vehicle safety communications," *IEEE Transactions on Vehicular Technology*, vol. 66, no. 3, pp. 2499–2512, 2017.

[50] J. Camp and E. Knightly, "Modulation rate adaptation in urban and vehicular environments: Cross-layer implementation and experimental evaluation," *IEEE/ACM Transactions on Networking*, vol. 18, no. 6, pp. 1949–1962, 2010.

[51] 3GPP, *Release 16 Description (v16.0.0)*, Jun. 2021. TR 21.916.

[52] ETSI, *Intelligent Transport Systems (ITS); LTE-V2X Access layer specification for Intelligent Transport Systems operating in the 5 GHz frequency band*, Jan 2020. ETSI EN 303 613 V1.1.1.

[53] ETSI, *Decentralized Congestion Control Mechanisms for Intelligent Transport Systems Operating in the 5 GHz Range: Access Layer Part, ETSI TS 102 687 v1.1.1*, Jul. 2011.

[54] C. B. Math, H. Li, S. H. de Groot, and I. Niemegeers, "Fair decentralized data-rate congestion control for v2v communications," in *2017 24th International Conference on Telecommunications (ICT)*, pp. 1–7, 2017.

[55] A. Rostami, B. Cheng, G. Bansal, K. Sjöberg, M. Gruteser, and J. B. Kenney, "Stability challenges and enhancements for vehicular channel congestion control approaches," *IEEE Transactions on Intelligent Transportation Systems*, vol. 17, no. 10, pp. 2935–2948, 2016.

[56] G. Bansal and J. B. Kenney, "Controlling congestion in safety-message transmissions: A philosophy for vehicular dsrc systems," *IEEE Vehicular Technology Magazine*, vol. 8, no. 4, pp. 20–26, 2013.

[57] A. Mansouri, V. Martinez, and J. Harri, "A First Investigation of Congestion Control for LTE-V2X Mode 4," *2019 15th Annual Conference on Wireless On-demand Network Systems and Services, WONS 2019 - Proceedings*, pp. 56–63, 2019.

[58] J. Hu, L. Zhao, Y. Feng, Y. Liu, and M. Gao, "Research, implementation and practice of congestion control mechanism in lte-v2x," in *2022 IEEE 96th Vehicular Technology Conference (VTC2022-Fall)*, pp. 1–5, 2022.

[59] 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, 2020.

[60] A. Burbano-Abril, B. McCarthy, M. Lopez-Guerrero, V. Rangel, and A. O'Driscoll, "Mcs adaptation within the cellular v2x sidelink," in *2021 IEEE Conference on Standards for Communications and Networking (CSCN)*, pp. 111–117, 2021.

[61] P. Wendland, G. Schaefer, and R. S. Thoma, "LTE-V2X Mode 4: Increasing robustness and DCC compatibility with reservation splitting," *2019 8th IEEE International Conference on Connected Vehicles and Expo, ICCVE 2019 - Proceedings*, pp. 0–5, 2019.

[62] Y. P. Fallah, N. Nasiriani, and H. Krishnan, "Stable and fair power control in vehicle safety networks," *IEEE Transactions on Vehicular Technology*, vol. 65, no. 3, pp. 1662–1675, 2016.

[63] B. Toghi, M. Saifuddin, Y. P. Fallah, and M. O. Mughal, "Analysis of distributed congestion control in cellular vehicle-to-everything networks," in *2019 IEEE 90th Vehicular Technology Conference (VTC2019-Fall)*, pp. 1–7, 2019.

[64] B. Toghi, M. Saifuddin, M. O. Mughal, and Y. P. Fallah, "Spatio-temporal dynamics of cellular v2x communication in dense vehicular networks," in *2019 IEEE 2nd Connected and Automated Vehicles Symposium (CAVS)*, pp. 1–5, 2019.

[65] Y. Yoon and H. Kim, "Balancing Power and Rate Control for Improved Congestion Control in Cellular V2X Communication Environments," *IEEE Access*, vol. 8, pp. 105071–105081, 2020.

[66] J. Choi and H. Kim, "A QoS-Aware Congestion Control Scheme for C-V2X Safety Communications," *IEEE Vehicular Networking Conference, VNC*, vol. 2020-Decem, pp. 23–26, 2020.

[67] S. Sabeeh and K. Wesołowski, "Congestion control in autonomous resource selection of cellular-v2x," *IEEE Access*, vol. 11, pp. 7450–7460, 2023.

[68] C. Wang, L. Ma, R. Li, T. S. Durrani, and H. Zhang, "Exploring trajectory prediction through machine learning methods," *IEEE Access*, vol. 7, pp. 101441–101452, 2019.

[69] W. Xu, Z. Chen, C. Zhang, X. Ji, Y. Wang, H. Su, and B. Liu, "Globalinsight: An lstm based model for multi-vehicle trajectory prediction," in *ICC 2020 - 2020 IEEE International Conference on Communications (ICC)*, pp. 1–7, 2020.

[70] A. Ip, L. Irio, and R. Oliveira, "Vehicle trajectory prediction based on lstm recurrent neural networks," in *2021 IEEE 93rd Vehicular Technology Conference (VTC2021-Spring)*, pp. 1–5, 2021.

[71] 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, 2022.

[72] H. Seon, H. Lee, and H. Kim, "Predicting cam generation times through machine learning for cellular v2x communication," *ICT Express*, 2022.

[73] B. McCarthy and A. O'Driscoll, "Opencv2x mode 4: A simulation extension for cellular vehicular communication networks," in *2019 IEEE 24th International Workshop on Computer Aided Modeling and Design of Communication Links and Networks (CAMAD)*, pp. 1–6, 2019.

[74] C. Sommer, R. German, and F. Dressler, "Bidirectionally Coupled Network and Road Traffic Simulation for Improved IVC Analysis," *IEEE Transactions on Mobile Computing (TMC)*, vol. 10, pp. 3–15, January 2011.

[75] R. Riebl, H. Gunther, C. Facchi, and L. C. Wolf, "Artery: Extending veins for VANET applications," in *2015 International Conference on Models and Technologies for Intelligent Transportation Systems (MT-ITS), Budapest, Hungary, June 3-5, 2015*, pp. 450–456, 2015.

[76] G. Cecchini, A. Bazzi, B. M. Masini, and A. Zanella, "Ltev2vsim: An lte-v2v simulator for the investigation of resource allocation for cooperative awareness," in *2017 5th IEEE International Conference on Models and Technologies for Intelligent Transportation Systems (MT-ITS)*, pp. 80–85, June 2017.

[77] 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.

[78] S. Bartoletti, B. M. Masini, V. Martinez, I. Sarris, and A. Bazzi, "Impact of the generation interval on the performance of sidelink c-v2x autonomous mode," *IEEE Access*, vol. 9, pp. 35121–35135, 2021.

[79] F. Eckermann, M. Kahlert, and C. Wietfeld, "Performance analysis of c-v2x mode 4 communication introducing an open-source c-v2x simulator," in *2019 IEEE 90th Vehicular Technology Conference (VTC2019-Fall)*, pp. 1–5, 2019.

[80] M. Malinverno, F. Raviglione, C. Casetti, C.-F. Chiasserini, J. Mangues-Bafalluy, and M. Requena-Esteso, "A multi-stack simulation framework for vehicular applications testing," in *Proceedings of the 10th ACM Symposium on Design and Analysis of Intelligent Vehicular Networks and Applications*, DIVANet '20, (New York, NY, USA), p. 17–24, Association for Computing Machinery, 2020.

[81] F. Raviglione, C. R. Carletti, M. Malinverno, C. Casetti, and C. Chiasserini, "ms-van3t: An integrated multi-stack framework for virtual validation of v2x communication and services," *Computer Communications*, vol. 217, pp. 70–86, 2024.

[82] 3GPP, *Study on LTE-based V2X services (v14.0.0)*, Jul. 2016. TR 38.885.

[83] P. A. Lopez, M. Behrisch, L. Bieker-Walz, J. Erdmann, Y.-P. Flötteröd, R. Hilbrich, L. Lücken, J. Rummel, P. Wagner, and E. Wießner, "Microscopic traffic simulation using sumo," in *The 21st IEEE International Conference on Intelligent Transportation Systems*, IEEE, 2018.

[84] S. Krauss, P. Wagner, and C. Gawron, "Metastable states in a microscopic model of traffic flow," *Phys. Rev. E*, vol. 55, pp. 5597–5602, May 1997.

[85] A. Virdis, G. Stea, and G. Nardini, "Simulating lte/lte-advanced networks with simulte," in *Simulation and Modeling Methodologies, Technologies and Applications* (M. S. Obaidat, T. Ören, J. Kacprzyk, and J. Filipe, eds.), (Cham), pp. 83–105, Springer International Publishing, Jan 2016.

[86] 3GPP, "Technical specification group radio access network; study on lte device to device proximity services; radio aspects (v12.0.1, release 12)," Mar. 2014. TR 36.843.

[87] ETSI, "Etsi en 302 636-4-1 intelligent transport systems (its); vehicular communications; geonetworking; part 4: Geographical addressing and

forwarding for point-to-point and point-to-multipoint communications; sub-part 1: Media-independent functionality (v1.3.1)," Sept. 2017.

[88] 3GPP, *Radio Resource Control (RRC); Protocol specification (v14.8.0)*, Sep. 2019. TS 36.331.

[89] 3GPP, *Physical layer; Measurements (v14.0.0)*, Dec. 2017. TS 36.214.

[90] 3GPP, "Tr 36.885 study on lte-based v2x services (v14.0.0, release 14)," *3GPP, Tech. Rep.*, Jul. 2016.

[91] 3GPP, "Tr 36.214 physical layer; measurements (v14.4.0, release 14)," *3GPP, Tech. Spec.*, Dec. 2017.

[92] 3GPP, "Technical specification group radio access network; evolved universal terrestrial radio access (e-utra); radio frequency (rf) requirements for lte pico node b (v16.0.0, release 16)," Jun. 2020. TR 36.931.

[93] R. R. Jian Wang, *BLER Performance Evaluation of LTE Device-to-Device Communications*, Nov 2016. Wireless Networks Division Communications Technology Laboratory (NIST).

[94] M. Nakagami, "The m-distribution—a general formula of intensity distribution of rapid fading," in *Statistical methods in radio wave propagation*, pp. 3–36, Elsevier, 1960.

[95] T.-H. Lee and C.-F. Lin, "Reducing collision probability in sensing-based sps algorithm for v2x sidelink communications," in *2020 IEEE REGION 10 CONFERENCE (TENCON)*, pp. 303–308, 2020.

[96] C. Sommer, O. K. Tonguz, and F. Dressler, "Adaptive beaconing for delay-sensitive and congestion-aware traffic information systems," in *2010 IEEE Vehicular Networking Conference*, pp. 1–8, 2010.

[97] C. Sommer, O. K. Tonguz, and F. Dressler, "Traffic information systems: efficient message dissemination via adaptive beaconing," *IEEE Communications Magazine*, vol. 49, no. 5, pp. 173–179, 2011.

[98] CAR 2 CAR Communication Consortium, *TR2052 Survey on CAM statistics (v1.0.1)*, Dec. 2018.

[99] D. Gyawali, "Comparative analysis of cpu and gpu profiling for deep learning models," 2023.

[100] 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.

[101] 3GPP, "Study on evaluation methodology of new vehicle-to-everything (v2x) use cases for lte and nr (v15.3.0, release 15)," Jul. 2019. TR 37.885.

[102] A. Mansouri, V. Martinez, and J. Härri, "A first investigation of congestion control for lte-v2x mode 4," in *2019 15th Annual Conference on Wireless On-demand Network Systems and Services (WONS)*, pp. 56–63, 2019.

[103] 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, p. 71685–71698, 2018.

[104] 3GPP, "Study on nr vehicle-to-everything (v2x) (v16.0.0, release 16)," Mar. 2019. TR 38.885.

[105] I. Corporation, *Sidelink Resource Allocation Mechanisms for NR V2X Communication*, August 2018. 3GPP TSG RAN WG1 Meeting no. 94 R1-1808696.

[106] B. McCarthy. and A. O'Driscoll., "Exploring nr-v2x dynamic grant limitations for aperiodic traffic," in *Proceedings of the 8th International Conference on Vehicle Technology and Intelligent Transport Systems - VEHITS,*, pp. 17–25, INSTICC, SciTePress, 2022.

[107] ETSI, "Ts 103 574 intelligent transport systems (its); congestion control mechanisms for the c-v2x pc5 interface; access layer part," Jan. 2018.

[108] R. Jain, D. M. Chiu, and H. WR, "A quantitative measure of fairness and discrimination for resource allocation in shared computer systems," *CoRR*, vol. cs.NI/9809099, 01 1998.

[109] G. Nardini, D. Sabella, G. Stea, P. Thakkar, and A. Virdis, "Simu5g–an omnet++ library for end-to-end performance evaluation of 5g networks," *IEEE Access*, vol. 8, pp. 181176–181191, 2020.

[110] ETSI, *Intelligent Transport Systems (ITS); Users and applications requirements; Part 1: Facility layer structure, functional requirements and specifications*, Aug 2013. ETSI TS 102 894-1 V1.1.1.

[111] ETSI, *Intelligent Transport Systems (ITS); Multi-Channel Operation study; Release 2*, Oct 2021. ETSI TR 103 439 V2.1.1.

[112] Car 2 Car Communication Consortium (C2CCC), *White paper on Additional Investigation of ITS-G5 and Sidelink LTE-V2X Co-Channel Coexistence Methods*, July. 2022.

Acronyms

3GPP Third Generation Partnership Project. i, ii, iv, v, vii, viii, xiii, 10, 13, 18, 21, 25, 26, 30, 33–35, 38, 42, 44, 47, 50, 56, 57, 59, 61, 65, 67, 68, 70, 72, 73, 77, 85–90, 92, 95, 99, 104, 106, 107, 117, 122–127, 134–136, 138–140, 147, 174, 180–184, 190, 205

5GAA 5G Automotive Association. 10

ADAS Advanced Driver Assistance Systems. 2, 10, 40

BLER Block Error Rate. 66, 67, 78

C-ITS Cooperative Intelligent Transport Systems. 3, 8, 207, 211–213

C-V2X Cellular Vehicle-To-Everything. i–iii, v, vi, viii, xiii, 3, 4, 6–18, 20, 22, 24–35, 38–47, 50, 51, 53, 54, 61–68, 70, 73, 77, 78, 80, 83, 84, 92, 96–99, 101, 105, 117, 118, 124, 126–134, 137, 140, 147, 163, 168, 171, 173–178, 180, 184, 189, 200, 205, 208, 210, 212, 213

C2CCC Car 2 Car Communication Consortium. 213

CA Cooperative Awareness. iii, 2, 3, 9, 12, 20, 25, 27, 41, 45, 51, 55, 56, 80, 84–86, 89, 90, 95, 101, 105, 108–112, 122, 148, 149, 171, 200, 214, 215

CAM Cooperative Awareness Message. iii, vi, 9, 10, 12, 29, 37, 52, 58, 85, 86, 92, 93, 96, 101, 108–110, 112, 118, 119, 122, 124, 145, 148–152, 154, 156, 160, 161, 165, 171, 198

CBR Channel Busy Ratio. ii, iii, v, vii, viii, 24, 26, 27, 33, 35, 56, 59–61, 72–74, 77, 90, 91, 100, 104–107, 117, 125, 134, 139–141, 143–145, 162–165, 167, 174, 177–184, 187, 188, 190–194, 198–206

CCAM Cooperative, Connected and Automated Mobility. 1

CP Cooperative Perception. iv, 2, 12, 85, 89, 109, 214, 215

CR Channel Occupancy Ratio. ii, viii, 24, 26, 33, 35, 59, 60, 77, 104–108, 174, 175, 180–182, 184, 190, 192–194, 199, 200, 202–206

CSR Candidate Single-Subframe Resources. 16, 17, 19, 20, 126, 128, 134, 136, 140, 192

D2D Device-to-Device. 10, 41, 50, 56

DCC Decentralised Congestion Control. ii–iv, vi–viii, 24–27, 34, 35, 51, 52, 56, 59–61, 90, 91, 95, 104–107, 174–184, 190–195, 197, 199–206, 212

DEN Decentralised Environmental Notification. 2, 3, 10, 12, 13, 41, 45, 51, 55, 56, 80, 85, 89, 108, 109, 118, 148, 171

DENM Decentralised Environmental Notification Message. 29, 30, 48, 141

ETSI European Telecommunications Standards Institute. i–viii, xiii, 3, 4, 10, 12, 20, 24, 26, 27, 29–35, 40–42, 44, 45, 47, 51–53, 57, 59, 75, 77, 80, 85–90, 92, 95–99, 104–107, 111, 112, 115, 117, 122–127, 129, 130, 133–136, 138–143, 145, 147, 149, 152, 153, 155, 159, 161, 163, 174, 175, 177–184, 205, 206, 212

GB Grant Break. 4, 5, 35, 83, 175, 177, 208

GLOSA Green Light Optimised Speed Advisory. 2

GRU Gated Recurrent Unit. iii, v, viii, 36, 110–113, 148, 171

HD Half-Duplex. 16, 17, 19, 72, 73, 75

IPG Inter-Packet Gap. vii, 197–199

IPT Inter-Packet Time. iii, vi, 9, 37, 58, 108–110, 112, 119, 122, 123, 126, 129, 130, 148–154, 156, 171, 179

ITS Intelligent Transport System. 11

LIMERIC **LI**near **ME**ssage **R**ate **I**ntegrated **C**ontrol. 27, 33, 34, 59, 60, 107, 175, 180, 190, 194, 201, 204, 205

LOS Line-of-sight. 1

LSTM Long-Short Term Memory. 36, 110, 112

MCS Modulation and Coding Scheme. i–viii, 3, 7, 14, 15, 18, 22–27, 30–36, 42, 46, 50, 57–59, 61–63, 65–67, 69, 70, 73, 75–77, 92, 94, 99, 115, 116, 118, 122, 159, 161–170, 172–176, 184–189, 205, 209

MT Missed Transmission. ii, 4, 5, 21, 35, 79–81, 84, 85, 88, 90–92, 95, 101–104, 107, 109, 110, 123, 125, 208

NED Network Description. 49, 51

NLOS Non-Line Of Sight. 2

NR-V2X New Radio Vehicle-To-Everything. i, iii, v–vii, xiii, 3, 4, 6–13, 16, 18–20, 22, 24–36, 40–42, 45, 51, 54, 64, 65, 78, 80, 83, 84, 96–99, 101, 105, 107–109, 117, 118, 121–123, 126, 128–134, 136, 137, 140–144, 147, 148, 163, 168, 171, 173, 174, 176, 177, 179–181, 183, 184, 189, 190, 193, 194, 199, 200, 205, 208, 210, 212

OBU On-Board Unit. 1, 213

PDR Packet Delivery Rate. v, viii, 33, 69–73, 75, 76, 100, 123, 124, 128, 134, 138, 141–145, 151, 153, 155, 162–165, 167, 170, 177–185, 191, 192, 198, 199, 205

PSCCH Physical Sidelink Control Channel. v, 14, 15, 66, 70, 138–140, 145, 147

PSD Power Spectral Density. 23, 24, 61, 66, 67, 176, 186

PSSCH Physical Sidelink Shared Channel. v, 14–17, 66, 70, 74

RB Resource Block. v, 13, 23, 24, 54, 55, 61, 62, 66, 67, 77, 92, 165, 188

RC Resource Reselection Counter. 17, 21, 22, 34, 58, 81, 84, 131, 175, 193

RE Resource Element. 61, 67

RNN Recurrent Neural Network. v, 36, 110, 111

RRI Resource Reservation Interval. ii, viii, 15, 17, 18, 20, 29, 34, 37, 57–60, 64, 65, 81, 82, 85–92, 102–112, 121, 125, 126, 138, 148, 149, 151–156, 171, 174, 175, 177, 178, 180, 190–193, 198, 199

RRI DCC Resource Reservation Interval Decentralised Congestion Control. iv, 6, 104, 189, 199

RSRP Reference Signal Received Power. 16, 17, 19, 28–30, 35, 61–68, 70, 72, 96, 98, 126, 134, 136, 147, 192

RSSI Reference Signal Strength Indicator. vi, vii, 16, 17, 19, 20, 24, 28–30, 61–67, 70, 72, 77, 83, 84, 96–98, 126, 128, 134–136, 147, 184, 188, 189

RSU Road-side Unit. 38, 52, 213

SB-SPS Sensing-Based Semi-Persistent Scheduling. ii, iii, v, vi, 4, 5, 10, 12, 13, 15–23, 25, 28–37, 42, 50, 53, 55, 57, 59, 61–64, 77, 80–86, 88, 91, 92, 94–99, 101–105, 107–110, 115, 116, 118, 121–124, 126–129, 133–138, 140–145, 147–149, 151–154, 156, 168, 169, 171, 172, 174–178, 180, 183, 184, 190–193, 204–206, 208, 209

SC-FDMA Single-Carrier Frequency Division Multiple Access. 13, 62, 67

SCI Sidelink Control Information. viii, 14–21, 28, 34, 37, 58, 63–66, 77, 81–83, 88, 90, 98, 102, 103, 105–108, 125, 126, 137, 148, 186, 187, 192

SFI Science Foundation Ireland. xv

SINR Signal-To-Noise Ratio. 35, 61, 62, 66, 67, 73, 184, 186

STR short-term reservations. 21

SUMO **S**imulation of **U**rban **MO**bility. 45, 47, 48, 50, 51, 53, 55, 70, 96, 113, 149, 175

TB Transport Block. 14, 15, 18, 21, 66, 96, 134, 137, 186, 187

TPC Transmission Power Control. 25, 26, 34, 35, 174

TRACI **Tra**ffic **C**ontrol **I**nterface. 48, 51, 53

TRC Transmission Rate Control. iv, 25, 26, 34, 35, 174–176, 180, 181, 184, 205, 212

UNAM Universidad Nacional Autónoma de México. 36, 61, 69, 184, 205

V2I Vehicle-To-Infrastructure. 1, 9, 109

V2N Vehicle-To-Network. 9

V2P Vehicle-To-Pedestrian. 9, 16

V2V Vehicle-To-Vehicle. 1, 9, 10

V2X Vehicle-To-Everything. v, 3, 9, 12, 47–51, 53, 55, 56, 80, 95, 106, 208, 212

VRU Vulnerable Road User. 1, 214, 215