MoReV2X – A New Radio Vehicular Communication Module for ns-3

This paper was downloaded from TechRxiv (<https://www.techrxiv.org>).

LICENSE

CC BY 4.0

SUBMISSION DATE / POSTED DATE

07-09-2021 / 25-10-2021

CITATION

Lusvarghi, Luca; Merani, Maria Luisa (2021): MoReV2X – A New Radio Vehicular Communication Module for ns-3. TechRxiv. Preprint. <https://doi.org/10.36227/techrxiv.16577621.v2>

DOI

[10.36227/techrxiv.16577621.v2](https://doi.org/10.36227/techrxiv.16577621.v2)

MoReV2X – A New Radio Vehicular Communication Module for ns-3

Luca Lusvarghi^{*†} and Maria Luisa Merani^{*†}

^{*}Dipartimento di Ingegneria "Enzo Ferrari"

Università degli Studi di Modena e Reggio Emilia, Italy

Email: luca.lusvarghi5@unimore.it, marialuisa.merani@unimore.it

[†] CNIT, Consorzio Nazionale Interuniversitario per le Telecomunicazioni, Parma, Italy

Abstract—The Third Generation Partnership Project (3GPP) recently released the 5G New Radio (NR) Vehicle-to-Everything (V2X) standard. NR-V2X is expected to support a large set of enhanced V2X applications, characterized by stringent requirements in terms of latency, reliability and data rate. The use of simulations plays a crucial role in the preliminary assessment of NR-V2X achievable performance, providing valuable insights on the strengths and limits of the new solution. This work introduces MoReV2X, an open-source ns-3 module for the simulation of NR-V2X communications. The development of MoReV2X adheres to 3GPP specifications, includes various traffic models, features an accurate implementation of the NR propagation channel between vehicles and allows to determine a thorough set of Key Performance Indicators (KPIs). A first set of exemplary results is reported to testify the proper functioning of the module. Moreover, the results reveal the impact that different subcarrier spacings have on system performance, when periodic and aperiodic traffic sources are considered.

I. INTRODUCTION

Over the last years, vehicular communications emerged as one of the key enablers for the digital transformation of transport systems. As a matter of fact, vehicle-to-everything data exchange can significantly improve the range and the quality of information available at each road user, fostering the broad deployment of sophisticated services that range from vehicle platooning to advanced and remote driving.

3GPP devoted significant efforts to the development of the cellular V2X standards and recently released NR-V2X, the 5G solution for V2X communications. NR-V2X can operate both at sub-6 GHz and mmWave frequencies, and relies on two operating modes for the assignment of radio resources to vehicles, termed Mode 1 and Mode 2. Both Modes allow Vehicular-User Equipments (V-UEs) to directly communicate on the so-called sidelink channel through the PC-5 interface; however, Mode 2 delegates V-UEs the task of identifying available radio resources in a distributed and autonomous manner.

In the evaluation of cellular V2X communications, the use of network simulators is instrumental to an accurate analysis, providing insights on expected system performance in different scenarios. As of today, numerical simulations have been leveraged by several works for the investigation of LTE-based vehicular communications. In [1], Bazzi *et al.* provided a detailed analysis on the impact of PHY and MAC layer parameters on LTE-V2X, the predecessor of NR-V2X. The

authors of [2] also assessed the performance of LTE-V2X using a custom simulator based on Veins. In [3], Eckermann *et al.* investigated the behavior of LTE-V2X Mode 4 in the urban setting. Recently, numerical simulations have been leveraged by the authors of [4], [5] to analyze the dissemination of event-triggered messages via LTE-V2X. In [6] and [7], network simulators have also been used to evaluate the impact of aperiodic traffic on LTE-V2X Mode 4 performance. Finally, the work in [8] presented an initial study on vehicular mmWave communications.

On the other hand, the investigation of NR-V2X is still in its infancy, and simulation tools for its analysis in the sub-6 GHz band are not available at the time of this writing. This work therefore presents MoReV2X¹, the first open-source ns-3 module for the simulation of NR-V2X communications. The implementation of MoReV2X focuses on Mode 2, that guarantees the seamless support of V2X services in out-of-coverage scenarios, and therefore represents the most promising approach for future V2X applications.

The development of MoReV2X adheres to the latest 3GPP specifications, including the most relevant PHY and MAC layer features standardized within 3GPP Release 16. Following the evaluation guidelines summarized in [9], MoReV2X includes an accurate characterization of the NR sidelink channels that captures large scale and small scale fading effects. Furthermore, MoReV2X implements a new set of realistic traffic models and evaluates a complete set of KPIs. MoReV2X can also be easily interfaced with SUMO [10] for the realistic simulation of vehicular mobility.

In this work, MoReV2X is employed for a preliminary assessment of Mode 2 achievable performance, when periodic and aperiodic traffic sources are considered, for two different subcarrier spacing values. For the examined scenario and vehicular density, the obtained results show that periodic traffic is more effectively served than aperiodic. Moreover, in the presence of periodic traffic, system performance improves when the subcarrier spacing increases. Conversely, aperiodic traffic performance is insensitive to what spacing is adopted.

The rest of this paper is organized as follows: Section II provides an overview of the PHY and MAC layer elements that characterize NR-V2X Mode 2. Section III presents the most

¹The code is available at <https://github.com/LLusvarghi/MoReV2X>

relevant points of the evaluation methodology suggested by 3GPP, i.e., traffic models, channel models and KPIs. Section IV illustrates the implementation of the aforementioned NR-V2X features within MoReV2X. Section V reports a set of exemplary results and Section VI draws the conclusions.

II. OVERVIEW OF NR-V2X TECHNOLOGY

Finalized within 3GPP Release 16, NR-V2X technology has been designed to support safety-critical applications and advanced driving services. To this aim, NR-V2X features significant modifications with respect to its predecessor, LTE-V2X.

At PHY layer, the NR-V2X design focuses on the frequency range going from 410 MHz to 7.125 GHz, and the current work concentrates on the specifications the standard provides for this band. However, mmWave frequencies are also mentioned in the standard documents. In NR-V2X, vehicles transmit using Cyclic Prefix (CP) - Orthogonal Frequency Division Multiplexing (OFDM) waveforms and unlike in LTE-V2X, different subcarrier spacings (SCSs) can be employed. The SCS values are $2^\mu \times 15$ kHz, where μ is the OFDM numerology index, $\mu = 0, 1, 2, 3$. For the considered frequency range, the SCS values of interest are 15, 30 and 60 kHz. The maximum system bandwidth is 50 MHz for SCS=15 kHz, 100 MHz for SCS=30 kHz and 200 MHz for SCS=60 kHz. In the frequency domain, the smallest reference unit is called Resource Block (RB) and it is made of 12 consecutive subcarriers: so, it can be 180, 360 or 720 kHz wide.

In the time domain, 10 ms long radio frames are considered, and each frame is further subdivided in 1 ms long subframes. Depending on the value of the numerology index, $1/2^\mu = 1, 2$ or 4 slots are present in each subframe, wherefrom a slot duration t_s is equal to 1, 0.5 and 0.125 ms, respectively. The slot is the smallest temporal unit on which transmissions can be organized.

Any NR-V2X-based system operates on a number of RBs that are frequency-contiguous and employs the same numerology over all RBs. To gain in flexibility, time slots reserved to

NR-V2X communications can be either contiguous or non-contiguous. Overall, the grid of frequency-time elements is called the resource pool. Within the same time slot, more consecutive RBs are grouped in a subchannel, that is the smallest unit for data transmission and reception. According to the standard, a subchannel can be made of 10, 12, 15, 20, 25, 75 or 100 RBs. Fig. 1 offers an example of the NR-V2X resource grid with contiguous time slots.

Whereas LTE-V2X only supports broadcast transmissions, NR-V2X also allows groupcast and unicast communications. The data packet is encapsulated within a Transport Block (TB) and protected using Low Density Parity Check (LDPC) codes. Each TB is paired with its Sidelink Control Information (SCI). Differently from LTE-V2X, the SCI is organized in two different stages. The first stage SCI contains relevant information for the correct decoding of the TB and of the second stage SCI. It indicates the subchannels the TB occupies, the Modulation and Coding Scheme (MCS) its transmission adopts, the TB priority, the resources that have been reserved for potential retransmissions of the TB. It is encoded employing polar codes and multiplexed in the frequency domain with the TB; V-UEs that are not the intended recipients of the TB decode the first stage SCI only, so as to understand which RBs are occupied by the TB transmission. The second stage SCI carries the PHY source and destination identifiers, that is, the address of the transmitting V-UE and of the destination V-UE(s). It therefore serves the purpose of supporting unicast and groupcast communications. Furthermore, it provides essential information for the support of Hybrid ARQ (HARQ) retransmissions, an important novelty with respect to LTE-V2X. The second stage SCI is transmitted on the same resources of the associated TB. The TB occupies a number of subchannels that depends on its size and on the adopted Modulation and Coding Scheme (MCS). The first stage SCI must be contained within one subchannel and can only employ QPSK modulation.

As regards the MAC sublayer, the resource allocation procedure of NR-V2X includes two modes, Mode 1 and Mode 2. In Mode 1, the gNodeB is responsible for scheduling the communications of the V-UE. In Mode 2, which represents the focus of this work, the resource allocation task is delegated to each V-UE. Therefore, the ego-vehicle autonomously identifies the set of subchannels able to accommodate the transmission of the TB and its associated SCI, i.e., the Single-Slot Resource (SSR). The process of (re)selecting a new SSR is referred to as resource reselection, and can be performed via two alternative strategies:

- **Dynamic:** the ego-vehicle selects a new SSR every time a TB plus its SCI have to be transmitted;
- **Semi-persistent:** the ego-vehicle selects and periodically reserves subchannels for a number of consecutive transmissions.

When a resource reselection is triggered, the ego-vehicle first identifies the set of candidate resources that lie within the so-called selection window (W). The selection window extends from the packet arrival time up to its latency deadline,

Fig. 1. NR-V2X Mode 2 time-frequency resource grid

defined by the Packet Delay Budget (*PDB*).

The ego-vehicle identifies unavailable resources in W exploiting the channel status information collected during the interval that precedes the packet arrival; this is the so-called sensing window, whose duration S can be set to either 1100 or 100 ms. Unavailable resources are those that have already been reserved by neighboring V-UEs, and whose Reference Signal Received Power (*RSRP*) in S is larger than a pre-configured threshold. The *RSRP* threshold is iteratively increased by 3 dB until the portion of candidate resources is larger than the percentage β of those initially included in W ; depending on the packet priority level, β can be set equal to 20, 35 or 50%. To take into account half-duplex constraints, the RBs where the ego-vehicle was previously transmitting, and could not sense reservations from other vehicles, are also removed from the pool of available resources in S .

Last, the ego-vehicle randomly selects an SSR from the remaining list of candidate resources. In the semi-persistent case, the selected SSR is reserved for a number of reselection counter, C_{resel} , periodic transmissions. The time-interval between two consecutive reservations is the Resource Reservation Interval (*RR*), and it is included in the first stage SCI to notify neighboring V-UEs about the current reservation. After every scheduled transmission, the reselection counter is decremented by one; when $C_{resel} = 0$, a resource reselection is triggered with probability $1 - P$, $P \in [0, 0.8]$.

With respect to LTE-V2X, NR-V2X additionally includes a new mechanism for the re-evaluation of selected resources and a pre-emption mechanism for traffic prioritization. A complete and detailed description of the PHY and MAC features introduced in NR-V2X Mode 2 is reported in [11].

III. EVALUATION METHODOLOGY

In [9], 3GPP stated a set of guidelines for the performance assessment of cellular V2X communications. Such set includes specific traffic models, it introduces relevant evaluation scenarios and channel models, and it indicates the KPIs to determine.

A. Traffic Models

1) *Periodic traffic*: This traffic model foresees a periodic generation of packets. The period T can be configured in three different ways, $T \in [10, 30, 100]$ ms, and the packet size can take on significantly large values, up to 60 kilobytes.

The simple assumption of periodic traffic perfectly suits the semi-persistent strategy of NR-V2X Mode 2, but it hinders a thorough evaluation of its scheduling capability. As a matter of fact, the periodic generation of messages does not reflect the nature of the traffic originating from the majority of driving applications. More realistically, the NR-V2X performance has to be evaluated considering different application-layer traffic profiles, in which neither the packet inter-arrival time nor the packet size is constant.

2) *eV2X use cases*: 3GPP defined a set of enhanced V2X (eV2X) use cases, such as vehicle platooning, remote and advanced driving and sensors information sharing. Accordingly, a traffic model was introduced to mimic eV2X data sources,

that are likely to exhibit a non-deterministic packet inter-arrival time, τ , and a large variability in packet size. In detail, τ is modeled as [9]

$$\tau = c + r \quad (1)$$

where c is a constant and r is an exponentially distributed random variable with mean $\bar{r} = c$. The packet size, X , is modeled as a discrete random variable, uniformly distributed in $[X_{min}, X_{max}]$.

B. Scenarios and Channel Models

A key aspect of the NR-V2X evaluation guidelines concerns the definition of new sidelink channel models. In this regard, 3GPP devoted significant efforts to characterize the large and small scale fading that affect vehicular communications in *Highway* and *Urban* scenarios.

For the Highway scenario, 3GPP has introduced two possible channel states, namely:

- **LOS**: the link between the transmitting and receiving V-UE is characterized by a direct Line-Of-Sight (LOS) path;
- **NLOSv**: the LOS path is blocked by the presence of other vehicles.

For every received packet, the channel state is determined according to the LOS and NLOSv state probabilities, $P(LOS)$ and $P(NLOSv) = 1 - P(LOS)$, defined in [9].

For the Urban scenario, three channel states are possible. When the transmitting and receiving V-UEs are in the same street, the sidelink channel can be either in the LOS or in the NLOSv state; for this scenario too, the LOS and NLOSv state probabilities are detailed in [9]. On the other hand, when the transmitting and receiving V-UEs are in different streets, a new Non-Line-Of-Sight (NLOS) channel state is introduced:

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

and $P(NLOS) = 1$.

In the Highway setting, the propagation loss plus shadowing contribution, PL , which characterizes both the LOS and the NLOSv states is:

$$PL = 32.4 + 20 \log_{10}(D) + 20 \log_{10}(f_c) + L \quad (2)$$

where f_c is the center frequency and L is a random term that accounts for shadow fading, modeled as a zero-mean lognormally distributed random variable with $\sigma_L = 3$ dB.

In the Urban setting, the PL for the LOS and NLOSv states is:

$$PL = 38.77 + 16.7 \log_{10}(D) + 18.2 \log_{10}(f_c) + L \quad (3)$$

Conversely, the PL for the Urban NLOS state is:

$$PL_{NLOS} = 36.85 + 30 \log_{10}(D) + 18.9 \log_{10}(f_c) + L \quad (4)$$

where σ_L increases to 4 dB. In both the Highway and Urban scenarios, an additional blockage loss term is added to PL when the transmitting and receiving V-UEs are in the NLOSv state [9].

Besides such propagation loss and shadowing characterization, the 3GPP document provides an equally accurate set of

fast fading models. They are based on Clustered Delay Lines (CDLs), whose parameters depend on the scenario and the channel state.

C. KPIs

Along with the definition of traffic and channel models, 3GPP introduced the performance metrics that shall be used for the analysis of NR-V2X communications [9].

For broadcast communications, the first metric is the Packet Reception Ratio (*PRR*) type 1. In the region where the vehicular distances from the transmitting V-UE are in the $[a, b]$ range, the *PRR* type 1 of every transmitted packet is evaluated as Y/Z , where Y is the number of V-UEs in the region that successfully receive the packet and Z is the total number of receiving V-UEs in the same region. Given N packets are generated during the simulation, the average *PRR* is evaluated as:

$$PRR = \frac{\sum_{j=1}^N Y_i}{\sum_{j=1}^N Z_i} \quad (5)$$

with $a = i \cdot 20$ m, $b = (i + 1) \cdot 20$ m, $i = 0, 1, \dots, 25$. The second merit figure is the Packet Inter-Reception (*PIR*) type 1. For a given transmitter-receiver pair, the *PIR* is defined as the time between two consecutive successful receptions of packets belonging to the same application flow, when the transmitter-receiver distance falls within the $(0, D_{max}]$ range, D_{max} fixed, at reception time. For unicast and groupcast communications, *PRR* and *PIR* type 2 are further introduced, and their definitions take into account only the intended set of receivers.

MoReV2X computes additional metrics for the assessment of the NR-V2X scheduling accuracy. They are the Propagation Losses Ratio (*PLR*) and the Collision Losses Ratio (*CLR*). For a given $[a, b]$ distance range, the *PLR* is determined as

$$PLR = \frac{N_{PL}}{N_{PL} + N_{CL} + N_{SR}} \quad (6)$$

and the *CLR* is defined as

$$CLR = \frac{N_{CL}}{N_{PL} + N_{CL} + N_{SR}} \quad (7)$$

where: (i) N_{PL} is the number of packets incorrectly received within the same range, because of poor propagation conditions, that is, N_{PL} is the number of packets that did not collide, but experienced an insufficient received Signal-to-Noise Ratio (*SNR*); (ii) N_{CL} is the number of packets in the $[a, b]$ distance range that were lost in a collision, and could not be successfully retrieved because of an insufficient Signal-to-Interference-plus-Noise Ratio (*SINR*); (iii) the N_{SR} is the number of correctly received packets within the same range.

Another metric MoReV2X determines is the Channel Busy Ratio (*CBR*): at a given slot n , the *CBR* is the fraction of busy subchannels in the $[n - 100 \cdot 2^\mu, n - 1]$ previous slots. A subchannel is labeled as busy if its Received Signal Strength Indicator (*RSSI*) is larger than a pre-configured threshold. The *CBR* is employed to monitor the channel load and to proactively enforce congestion control, if needed.

When the NR-V2X semi-persistent strategy is employed, some additional metrics turn out useful to fully understand the scheduling strategy behavior:

- **Unused Reservations Ratio (*URR*):** it is the ratio between the reserved SSRs that have not been used for the transmission of a packet over the total number of performed reservations. Unused reservations occur when the reservation period *RRI* is smaller than the packet inter-arrival time;
- **Unused Subchannels Ratio (*USR*):** this indicator is the ratio between the reserved subchannels that have not been used for accommodating a transmitted packet over the total number of reserved subchannels. Unused subchannels occur when the reservation size is larger than the size of the packet to be transmitted;
- **Additional Reselections Ratio (*ARR*):** it is the ratio between the number of latency and size reselections with respect to the total number of transmitted messages. Latency and size reselections are generated when the reserved resources cannot satisfy the packet delay budget or the number of subchannels required for the packet transmission, respectively;
- **Counter Reselections Ratio (*CRR*):** it is the ratio between the number of packets which triggered a resource reselection due to the reselection counter time-out and the total number of packets.

IV. NS-3 IMPLEMENTATION

This Section presents the main features of the MoReV2X module, developed starting from the simulator presented in [4], and adhering to the 3GPP specifications and evaluation guidelines summarized in Sections II and III.

A. Application Layer

In the MoReV2X module, every simulation can be configured using three different command line arguments, namely:

- **Aperiodic:** when this option is selected, V-UEs generate packets in accordance to the aperiodic profile described in Sec. III-A2. The constant term c in (1), the smallest and the largest packet size values can all be adjusted;
- **Periodic:** when this option is chosen, all V-UEs generate periodic traffic. As indicated in [9], a pattern of five arbitrary packet sizes is chosen, that endlessly repeats. For each V-UE, the starting point is randomly set;
- **Mixed:** this configuration allows for the coexistence of periodic and aperiodic traffic sources within the same simulation. The percentage of periodic V-UEs with respect to the total number of vehicles, termed Δ , is set using the `Percentage` command line argument. The default value is $\Delta = 50\%$.

B. MAC Sublayer

On purpose, MoReV2X concentrates on the NR-V2X Mode 2 semi-persistent access strategy. The sensing window width

S , the $RSRP$ threshold and the minimum amount of candidate resources β are set employing dedicated variables. The width of the selection window W is dynamically adjusted, depending on the latency deadline which characterizes the packet that triggered the resource (re)-selection. The semi-persistent strategy is the default choice. However, the dynamic strategy can be enabled via the ns-3 attribute system, setting the Boolean MAC layer `DynamicScheduling` variable to true. The current version of MoReV2X supports broadcast communications.

Every time new radio resources have to be selected for the transmission of a packet, a valid scheduling grant is generated. Grants contain relevant information for the scheduling of each V-UE transmission. The content of a grant includes the position of the reserved SSR on the time-frequency grid, the RRI configuration and the reselection counter value.

At the beginning of each slot, if there is a new packet to be transmitted and the V-UE has a valid scheduling grant (i.e., $C_{resel} \neq 0$), the transmission opportunity is notified to the PHY layer. Otherwise, if there is no packet waiting for transmission, the V-UE switches to reception mode and listens for incoming received packets forwarded from lower layers.

MoReV2X computes all the KPIs introduced in Subsection III-C. Their values are periodically saved in dedicated output files.

C. Physical (PHY) Layer

1) *Subcarrier Spacing*: MoReV2X can accommodate different SCS choices, namely, 15, 30 and 60 kHz, via the SCS command line argument.

2) *Sidelink Channel*: MoReV2X features the 3GPP-defined *Highway* and *Urban* channel models discussed in subsection III-B. To evaluate the large and small scale fading effects on the decoding process of a TB, MoReV2X runs the following steps:

- 1) the average $SINR$ of the TB carrying the data payload and of its first stage SCI is evaluated taking into account the propagation loss plus the shadowing contribution. Depending on the selected scenario and on the relative position of the transmitting and receiving V-UEs, the path loss PL is provided by eq. (2), eq. (3) or eq. (4). The $SINR$ is averaged over the occupied RBs, interpreting the overlapping transmissions coming from neighboring V-UEs as interference.
- 2) if the average $SINR$ is below the receiver sensitivity, the TB cannot be recovered. Otherwise, the $SINR$ levels of the TB and of the first stage SCI are separately mapped into a Packet Error Rate (PER) value, P_{TB} and P_{SCI} , respectively, on the basis of Look Up Tables (LUTs) modeling fast fading effects. At the time of this writing, the LUTs included within the MoReV2X module have been obtained from [12] and [13] for the TB and the first stage SCI, respectively.
- 3) Next, the first stage SCI is declared successfully decoded with probability $1 - P_{SCI}$. As the first stage SCI contains decoding-critical information, the TB can be successfully

coded only if the SCI has been correctly retrieved. In turn, the TB is declared successfully decoded with probability $1 - P_{TB}$.

V. NUMERICAL RESULTS

In this Section, a set of exemplary results obtained through the MoReV2X module is reported. The simulation scenario consists of a 5 km long highway segment with 3 lanes per driving direction. Vehicular mobility has been recreated using SUMO and adhering to the 3GPP option that models all V-UEs as 5 m long, 2 m wide and 1.6 m high passenger cars, traveling at 140 km/h [9]; the average vehicular density is 60 vehicles/km.

The system bandwidth is 10 MHz wide, centered at 5.9 GHz; the SCS value is 15 or 30 kHz, thus resulting in a total number of RBs equal to 52 and 24, respectively [14]. In the frequency domain, RBs are organized in subchannels made of 12 RBs each. The TB and the SCI are multiplexed in non-overlapping resources and transmitted during the same time slot, employing QPSK and a code rate 0.5 and 0.23, respectively. The number of symbols occupied by the SCI in each slot is equal to 2, and 4 DMRS symbols are employed for the decoding of the TB. The transmission power is set to 23 dBm for all V-UEs.

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

The packet delay budget is equal to the minimum packet inter-arrival time, i.e., 100 ms for the `Periodic` profile, 50 ms for the `Aperiodic` alternative. All simulations adopt the Mode 2 semi-persistent strategy: the number of reserved subchannels is set so as to accommodate the largest expected packet size, i.e., 300 bytes, and subchannels are reserved with a periodicity $RRI = PDB$, thus avoiding any size and latency reselection. New resources are selected every time the reselection counter expires, i.e., $P = 0$, and the re-evaluation mechanism is not enabled. The Mode 2 initial $RSRP$ threshold is configured to -128 dBm and the β threshold is set to 20%.

Fig. 2 reports the average PRR type 1 as a function of the distance D between the transmitting and the receiving V-UE. The semi-persistent strategy achieves higher PRR levels when periodic traffic is considered, as constant packet inter-arrival times perfectly fit the periodic, semi-persistent resource reservation approach. This holds true irrespective of the SCS value. The figure also reveals that for periodic traffic, the PRR performance improves when increasing the SCS from 15 to 30 kHz. This effect has to be mainly ascribed to a different number of slots in the selection window, which is twice as high in the 30 kHz case; in turn, this means a lower probability that

Fig. 2. PRR as a function of the Tx-Rx distance D

Fig. 4. PLR as a function of D

Fig. 3. PIR CDF, $D_{max} = 520$ m

two or more V-UEs select a subchannel in the same slot. When this event happens, V-UEs transmitting within the same slot are “blind” to each other and persistently collide at other receiving V-UEs, if they select the same subchannel. On the other hand, the PRR performance for aperiodic traffic does not improve when increasing the subcarrier spacing. Here, the main cause of PRR degradation is the mismatch between the aperiodic packet generation and the periodic allocation of resources, that forbids V-UEs to correctly announce their reserved resources, and in turn prohibits neighboring V-UEs to gain a correct view of the actual reserved resources.

Next, Fig. 3 reports the PIR type 1 Cumulative Distribution Function (CDF) for $D_{max} = 520$ m. For all simulations, the CDF step behavior reflects the periodic nature of the

reservations placed using the semi-persistent strategy, and indirectly validates MoReV2X accurate implementation. For periodic traffic, the probability of observing PIR values lower than 100 ms is close to 0.9, further demonstrating the NR-V2X Mode 2 ability to cope with this traffic type.

In Fig. 4, the PLR as a function of the distance D is reported, where it is recalled that the PLR measures the fraction of packets exclusively lost because of poor propagation conditions. As expected, this figure shows that the PLR does not depend on the traffic type, but it exclusively depends on the SCS. In detail, it reveals that propagation losses are negligible for D values lower than 500 m and that above this value the SCS=15 kHz choice guarantees lower PLR values than SCS=30 kHz. As a matter of fact, the noise power of a subchannel is twice as high when SCS=30 kHz, for a given noise power spectral density. Therefore, packets transmitted in the system adopting SCS=30 kHz exhibit an SNR value which is 3 dB lower than in the system where SCS=15 kHz, on average. Moreover, given the vehicles relative speed is not excessively high (its maximum value is 280 km/h), the 30 kHz choice does not improve system robustness against Doppler effects.

Then, Fig. 5 reports the Collision Losses Ratio (CLR) as a function of D . In the presence of periodic traffic, collisions are constantly lower than for the aperiodic traffic case. This figure corroborates the intuition that a smaller time slot duration, due to the adoption of a larger subcarrier spacing, reduces the number of collisions only in the case of periodic traffic.

Finally, Table I reports the average values of the remaining KPIs. The $RSSI$ threshold employed for the CBR evaluation is 0.5 dB greater than the receiver sensitivity, i.e., -89.9 dBm. Note that the CBR for the case of periodic traffic is the highest. As a matter of fact, the periodic and the aperiodic profiles pour the same amount of traffic on the channel, but the semi-persistent strategy is able to allocate

Fig. 5. CLR as a function of D

resources more effectively in the former circumstance. For the same channel load, a smaller number of collisions results in a more efficient occupation of the channel, hence leading to a larger CBR . In the case of periodic traffic, the RRI perfectly matches the packet generation period T , and no reserved subchannels are left idle. Accordingly, $URR = 0.0$ in the first row. Conversely, the aperiodic traffic profile generates packets with an average inter-arrival time equal to $\bar{\tau} = 100$ ms, but $RRI = 50$ ms; as a consequence, one reservation out of two is left unused, on average, which allows to predict $URR = 0.5$. And indeed, the URR obtained by simulation is 0.48. On the third column, observe that the ratio of unused subchannels over the total number of reserved subchannels is equal to 0.4 for both traffic profiles. It is easy to verify that the numerically determined USR coincides with the value we can forecast. In Periodic simulations, the packet size pattern is $[300, 100, 100, 100, 100]$, which translates in the pattern of occupied subchannels $[2, 1, 1, 1, 1]$; as the semi-persistent strategy has been configured to reserve the largest required number of subchannels, i.e., 2, the USR value we should expect is $1 - \frac{2+4 \cdot 1}{5 \cdot 2} = 0.4$. In Aperiodic simulations, 100 and 300 bytes are selected with probabilities 0.8 and 0.2, respectively. Therefore, the expected average number of unused subchannels is $1 - \frac{1 \cdot 0.8 + 2 \cdot 0.2}{2} = 0.4$.

The fourth column of Table I confirms that our configuration of the semi-persistent strategy triggers no additional reselections (neither size nor latency reselections), that is, $ARR = 0$. Only counter reselections are triggered, and the CRR values in the fifth column quantify to what extent different traffic profiles impact on their number. As a matter of fact, the CRR can be easily predicted for both traffic profiles: in the periodic case, when $RRI = 100$ ms, the average C_{resel} value is equal to 10, on average leading to a resource reselection once every ten packets, i.e., $CRR = 0.1$. For aperiodic traffic, where $RRI = 50$ ms, the average C_{resel} value is equal to 20. As, on

TABLE I
ADDITIONAL METRICS

	CBR	URR	USR	ARR	CRR
Periodic	0.38	0	0.4	0	0.1
Aperiodic	0.36	0.48	0.4	0	0.1

average, one reservation out of two is left idle, only one packet out of ten is triggering a resource reselection, i.e., $CRR = 0.1$.

The perfect match between the *a priori* estimated values of URR , USR , ARR and CRR and those numerically determined testifies the correct design of the MoReV2X module.

VI. CONCLUSIONS

This work has presented MoReV2X, an ns-3 module for 5G NR-V2X vehicular communications. Its development adheres to the 3GPP specifications and performance assessment guidelines: as such, MoReV2X replicates many significant elements of NR-V2X and evaluates a complete set of KPIs. The first obtained results validate the proper functioning of MoReV2X, and offer an interesting insight on the achievable performance of NR-V2X for two alternative SCS choices, when periodic and aperiodic traffic patterns are examined.

REFERENCES

- [1] 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.
- [2] R. Molina-Masegosa and J. Gozalvez, "System Level Evaluation of LTE-V2V Mode 4 Communications and Its Distributed Scheduling," in *2017 IEEE 85th Veh. Technol. Conf. (VTC Spring)*, 2017, pp. 1-5.
- [3] 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 Veh. Technol. Conf. (VTC2019-Fall)*, 2019, pp. 1-5.
- [4] L. Gibellini and M. L. Merani, "Out-of-Coverage multi-hop road safety message distribution via LTE-A cellular V2V (C-V2V)," in *Proc. IEEE 88th Veh. Technol. Conf. (VTC-Fall)*, Chicago, IL, USA, Aug. 2018, pp. 1-6.
- [5] F. Romeo, C. Campolo, A. Molinaro and A. O. Berthet, "Asynchronous Traffic on the Sidelink of 5G V2X," in *2020 IEEE Int. Conf. on Commun. Workshops (ICC Workshops)*, 2020, pp. 1-6.
- [6] 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.
- [7] 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.
- [8] M. Drago, T. Zugno, M. Polese, M. Giordani, and M. Zorzi, "MilliCar: An ns-3 Module for mmWave NR V2X Networks," in *Proc. of the 2020 Workshop on ns-3 (WNS3 2020)*, pp. 9-16.
- [9] 3GPP, "Study on evaluation methodology of new Vehicle-to-Everything (V2X) use cases for LTE and NR," TR 37.885, Rel-15 V15.3.0, Jun. 2019.
- [10] P. A. Lopez et al., "Microscopic Traffic Simulation using SUMO," in *21st Int. Conf. on Intell. Transp. Syst. (ITSC)*, 2018, pp. 2575-2582.
- [11] M. H. C. Garcia et al., "A Tutorial on 5G NR V2X Communications," *IEEE Commun. Surveys & Tutorials*.
- [12] Ericsson, "R1-1813644. Link level evaluations of NR PSSCH," 3GPP TSG-RAN WG1 Meeting #95 Spokane, USA, Nov. 2018.
- [13] Ericsson, "R1-1907147. Link level evaluations of NR PSSCH," 3GPP TSG-RAN WG1 Meeting #97 Reno, USA, May 2019.
- [14] 3GPP, "User Equipment (UE) radio transmission and reception; Part 1: Range 1 Standalone," TS 38.101-1, Rel-17 V17.1.0, March 2021.