

On the Impact of Re-evaluation in 5G NR V2X Mode 2

Alejandro Molina-Galan, *Student Member, IEEE*, Luca Lusvarghi, *Graduate Student Member, IEEE*, Baldomero Coll-Perales, *Member, IEEE*, Javier Gozalvez, *Senior Member, IEEE*, Maria Luisa Merani, *Senior Member, IEEE*

Abstract—5G NR V2X has been designed to support advanced connected and automated driving V2X services. These services are characterized by variable traffic patterns that can generate packet collisions in decentralized systems where vehicles autonomously select their radio resources like 5G NR V2X mode 2. 5G NR V2X introduces a re-evaluation mechanism at the MAC layer to detect and avoid possible packet collisions before a vehicle transmits in selected resources. Most of the studies conducted to date on 5G NR V2X do not consider the re-evaluation mechanism despite being a mandatory MAC feature. This paper advances the state of the art with an in-depth analysis and evaluation of the operation and performance of re-evaluation in 5G NR V2X mode 2 under different traffic patterns and mode 2 configurations. The study shows that re-evaluation is effective in avoiding collisions with periodic traffic but its effectiveness decreases with aperiodic traffic and of variable size. The study also shows that re-evaluation is effective in avoiding collisions generated by the retransmission of packets. However, its overall impact on the performance of 5G NR V2X mode 2 is small, while it can have a relevant implementation cost due to the frequent re-evaluation checks and resource reselections. This raises questions on the current design of the re-evaluation mechanism that is a mandatory feature in 5G NR V2X mode 2.

Index Terms—5G NR V2X, re-evaluation, CAV, C-V2X, cellular V2X, connected automated vehicles, Mode 2, NR V2X, collisions, aperiodic, periodic, distributed scheduling, resource allocation.

I. INTRODUCTION

THE 5G New Radio (NR) Vehicle-to-Everything (V2X) standard published in 3GPP Release 16 is the first 5G NR standard that enables sidelink (SL) or direct Vehicle-to-Vehicle (V2V) communications using the NR PC5 interface [1]. 5G NR V2X (or NR V2X) is designed to complement and not replace LTE V2X. LTE V2X was designed to support basic safety applications using broadcast messages. NR V2X also supports unicast and groupcast transmissions, and includes new features and functionalities to support advanced V2X services with stringent requirements such as cooperative perception and driving, among others. To this aim, NR V2X SL introduces two new operating modes: mode 1 and mode 2. In mode 1, the cellular infrastructure manages and selects the

communication resources for each SL communication, while in mode 2, vehicles autonomously select and manage radio resources without the support of the cellular infrastructure.

NR V2X mode 2 is critical to support connected and automated mobility since safety services should not always depend on the availability of cellular coverage. According to 3GPP [2], these advanced safety services will generate V2X messages of variable size and generation times. Variable traffic patterns were shown to significantly impact the operation and performance of LTE V2X mode 4 [3], which is the counterpart of NR V2X mode 2. This was due to certain Medium Access Control (MAC) inefficiencies when vehicles generate aperiodic messages of variable size that result in packet collisions and require additional solutions [4][5]. NR V2X mode 2 introduces a re-evaluation mechanism at the MAC sublayer to detect and prevent possible collisions caused by aperiodic messages of variable size. The re-evaluation mechanism is a mandatory MAC feature that is executed before a vehicle transmits on selected resources to detect any possible packet collisions [6]. Several studies have recently analyzed the performance of NR V2X mode 2. In [7], authors evaluate different configurations of NR V2X mode 2 parameters under periodic traffic of fixed size. The studies reported in [8], [9] and [10] evaluate NR V2X mode 2 considering also aperiodic traffic of fixed size. However, these studies ([7]-[10]) did not implement re-evaluation despite being a mandatory MAC feature, and did not consider the variability in the generation and size of messages as recommended in 3GPP guidelines for advanced safety services [2]. For the sake of brevity, we refer to NR V2X mode 2 as mode 2 in the rest of the paper.

This paper advances the state of the art with an in-depth analysis and evaluation of the impact of the re-evaluation mechanism on the operation and performance of mode 2. To the authors' knowledge, this is the first study that analyses the effectiveness of re-evaluation to detect and avoid packet collisions when mode 2 operates with the Semi-Persistent Scheduling (SPS) or Dynamic Scheduling (DS) schemes. With SPS, vehicles select and reserve radio resources for the transmission of several consecutive data packets as well as for their possible retransmissions. On the other hand, vehicles

UMH work was supported in part by MCIN/AEI/10.13039/501100011033 (grants IJC2018-036862-I, PID2020-115576RB-I00) by Ministry of Universities (grant FPU18/00691) and by Generalitat Valenciana (CIGE/2021/096). (*Corresponding author: Luca Lusvarghi*).

A. Molina-Galan, B. Coll-Perales and J. Gozalvez are with the UWICORE Laboratory, Universidad Miguel Hernandez de Elche (UMH), 03202 Elche, Spain, e-mail: {alejandro.molinag, bcoll, j.gozalvez}@umh.es.

L. Lusvarghi and M. L. Merani are with the Department of Engineering "Enzo Ferrari", University of Modena and Reggio Emilia, Modena 41125, Italy, e-mail: {luca.lusvarghi5, marialuisa.merani}@unimore.it, and also with Consorzio Nazionale Interuniversitario per le Telecomunicazioni, CNIT.

using DS need to select new radio resources for the transmission of each data packet, and can only reserve resources for the retransmission of these packets. This study performs the analysis considering that vehicles transmit periodic or aperiodic packets of fixed or variable size following 3GPP guidelines in [2]. The study also evaluates the impact of retransmissions on the effectiveness of re-evaluation. Our study shows that re-evaluation is effective in avoiding packet collisions when packets are periodic and of fixed size and are transmitted with SPS. However, these collisions are rare, and hence the impact of re-evaluation for this traffic is small. On the other hand, the effectiveness of re-evaluation to avoid packet collisions decreases with aperiodic traffic of variable size whether using SPS or DS. The capacity of the re-evaluation mechanism to detect and avoid packet collisions improves when retransmissions are considered under both SPS and DS. However, the impact of re-evaluation on the performance of mode 2 is small with SPS and DS since the benefit of retransmissions prevails over the gains obtained with the packet collisions avoided with re-evaluation.

The rest of this paper is organized as follows. Section II provides an overview of mode 2, including the re-evaluation mechanism and a discussion on the impact of packet variability on the MAC. Section III presents an in-depth analysis of the re-evaluation mechanism that identifies and helps understand when re-evaluation can be effective or not in detecting and avoiding packet collisions. Section IV presents the evaluation environment and the metrics utilized. Section V evaluates the impact of re-evaluation on SPS without retransmissions, and Section VI extends the analysis to the scenario where retransmissions are considered. The impact of re-evaluation on DS is analyzed in Section VII, and Section VIII summarizes the main outcomes of this study.

II. RESOURCE ALLOCATION IN 5G NR V2X MODE 2

Mode 2 radio resources are organized in a grid that is made of slots in the time domain and Resource Blocks (RBs) in the frequency domain. The slot duration is $2^{-\mu}$ ms and an RB consists of 12 consecutive subcarriers with a subcarrier spacing (SCS) of $2^\mu \times 15$ kHz, where μ is the OFDM numerology, $\mu = 0, 1, 2$, or 3 . This results in slots of $\{1, 0.5, 0.25, 0.125\}$ ms and RBs of $\{180, 360, 720, 1440\}$ kHz for SCSs of $\{15, 30, 60, 120\}$ kHz, respectively. Vehicles in a particular region communicate over a common set of radio resources, termed resource pool. A resource pool uses a single numerology and its RBs are referred to as physical resource blocks (PRBs). PRBs within the same slots are grouped into sub-channels that represent the smallest unit for SL data transmission or reception (see Fig. 1). The number of PRBs that form a sub-channel (i.e., the sub-channel size) can be configured but it is fixed for a given resource pool.

In mode 2, data packets are transmitted in Transport Blocks (TBs) that are carried on the Physical Sidelink Shared Channel (PSSCH). Note that the terms TB and packet are interchangeable in this paper. A TB can occupy more than one sub-channel depending on the size of the packet, the sub-channel size, and the utilized Modulation and Coding Scheme (MCS). TBs can be transmitted using QPSK, 16-QAM, 64-QAM or 256-QAM modulations, and are encoded using Low-

Density Parity-Check (LDPC) coding. Each TB is associated with Sidelink Control Information (SCI). A TB and its associated SCI are transmitted in the same slot. The SCI in NR V2X is transmitted in two stages. The 1st-stage SCI is carried on the Physical Sidelink Control Channel (PSCCH), while the 2nd-stage SCI is multiplexed together with the TB in the PSSCH. The 1st-stage SCI indicates the resources used by the PSSCH and carries information required for decoding the TB. If retransmissions are employed, the 1st-stage SCI indicates the resource reservation for up to two retransmissions of the TB. The 1st-stage SCI also informs about the *Resource Reservation Interval (RRI)* if the vehicle reserves resources semi-persistently for the PSSCH, as detailed in Subsection II.A. The 2nd-stage SCI carries information used for decoding the PSSCH as well as for supporting retransmissions and mechanisms to report channel state information.

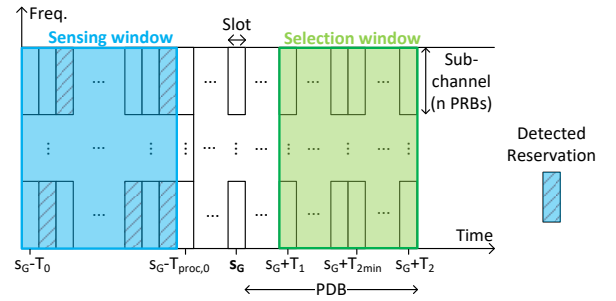


Fig. 1. NR V2X channelization and illustration of resource allocation in mode 2 (when $T_2 = PDB$).

A. Resource allocation

Vehicles using mode 2 autonomously select their resources (one or several sub-channels) to transmit the TBs. The process to select new resources is referred to as *reselection* by the 3GPP standards. Mode 2 can operate using a Dynamic Scheduling (DS) or Semi-Persistent Scheduling (SPS) scheme. Both schemes follow similar procedures to select new resources [11]. However, the DS scheme selects new resources for each TB and can only reserve resources for the retransmissions of that TB. The SPS scheme selects and reserves resources for the transmission of *Reselection Counter* consecutive TBs, and can also reserve resources for the retransmissions of the TBs. It is important to highlight the differences between selected and reserved resources. A selected resource is a resource that a vehicle selects to transmit a TB using the two-step resource allocation algorithm that is described in the remainder of this Section. A reserved resource is a selected resource that the vehicle reserves for a future transmission by notifying neighboring vehicles using the 1st-stage SCI. The *RRI* determines the time period between the resources reserved for the transmission of consecutive TBs in SPS. The *RRI* can be $\{0, [1:99], 100, 200, 300, 400, 500, 600, 700, 800, 900, 1000\}$ ms. A vehicle can be configured with a list of up to 16 different *RRI*s, but it only selects one *RRI* from the list when it selects new resources. The selected *RRI* also determines the value of the *Reselection Counter*. The vehicle randomly sets the *Reselection Counter* within the interval $[5, 15]$ when $RRI \geq 100$ ms, and within the interval $[5 \cdot C, 15 \cdot C]$, where $C = 100 / \max(20, RRI)$, when $RRI < 100$ ms [6].

New resources (for both the DS and SPS schemes) are selected in the so-called selection window portrayed in Fig. 1 [11]. The selection window is defined within the range of slots $[s_G + T_1, s_G + T_2]$, where s_G is the slot at which a new TB is generated. T_1 is the processing time required to identify candidate resources within the selection window to transmit the TB and its associated SCI. $T_1 \leq T_{proc,1}$, where $T_{proc,1}$ is 3, 2.5, 2.25 or 2.125 ms for an SCS of 15, 30, 60 or 120 kHz, respectively. T_2 can be set by the vehicle within $T_{2min} \leq T_2 \leq PDB$. The Packet Delay Budget (PDB) is established by the V2X application generating the TB, and defines the latency deadline by which the TB must be transmitted. T_{2min} can be set by the vehicle to $\{1, 5, 10, 20\}$ ms depending on the priority of the TB. Vehicles sense transmissions performed by other vehicles within the so-called sensing window (Fig. 1) while they are not transmitting. This allows them to identify which candidate resources are available within the selection window. The sensing window range is $[s_G - T_0, s_G - T_{proc,0}]$. T_0 can be 1100 ms or 100 ms, and $T_{proc,0}$ is equal to 1 ms for a SCS of 15 kHz and 0.50 ms for the remaining values of SCS.

DS and SPS schemes follow a two-step algorithm to select new resources [6][11]. During step 1, the vehicle is in charge of excluding resources from the selection window. First, the vehicle excludes resources that it could not sense when it was transmitting due to its half-duplex operation. In particular, if a vehicle could not sense resources at slot s_i within the sensing window, it excludes all resources within the selection window located at an integer number of RRI (in slots) ahead of s_i^1 . The vehicle also decodes the 1st-stage SCI received from other vehicles in the sensing window. For each transmission received in the sensing window, the vehicle also measures the Reference Signal Received Power (RSRP) [11]. A resource in the selection window is considered occupied if the vehicle detected in the 1st-stage SCIs decoded in the sensing window that another vehicle was reserving it and the measured RSRP was higher than an RSRP threshold. If this is the case, these candidate resources within the selection window are excluded. Once the execution of step 1 is completed, the vehicle (with DS or SPS) checks whether the percentage of candidate resources that have not been excluded in the selection window is equal to or higher than a threshold $X\%$; X can be 20, 35, or 50. If not, step 1 is repeated using an RSRP threshold increased by 3 dB.

In step 2 (with DS or SPS), the vehicle randomly selects the resources for the transmission of a TB from the available candidate resources within the selection window. A vehicle can select N candidate resources ($N \leq 32$) within the same selection window for the initial transmission of a TB and its $N-1$ retransmissions. NR V2X supports blind and Hybrid Automatic Repeat Request (HARQ) feedback-based retransmissions. Blind retransmissions are considered in this work when we refer to retransmissions. Each vehicle can select the value of N but it cannot be higher than the number of available candidate resources after step 1. The vehicle considers the limitations of the 1st-stage SCI for the selection and reservation of the N candidate resources. In particular, a 1st-stage SCI can only

notify about a maximum number of N_{SCI} resources (equal to 2 or 3). The selection of candidate resources also takes into account that a 1st-stage SCI can only notify about resource reservations for retransmissions located within a window W of 32 slots, with the first slot of W being the one where the 1st-stage SCI is transmitted. When retransmissions are separated by more than 32 slots from the slot where the 1st-stage SCI is transmitted, they are not reserved with the 1st-stage SCI [6].

With SPS, when the vehicle performs the transmission of a TB, it also reserves resources for the transmission of the next TB using the RRI included in the 1st-stage SCI. The RRI also reserves the resources for the retransmissions of the next TB when the 1st-stage SCI informs about the retransmissions of a TB. The vehicle reserves resources every RRI ms for *Reselection Counter* transmissions. The *Reselection Counter* is decremented by one every time the vehicle transmits a TB and its $N-1$ retransmissions. When *Reselection Counter* depletes, the vehicle decides with probability $(1-P)$ whether it has to select new resources for the transmission of the following TBs; P can be set between 0 and 0.8. If not, the vehicle keeps using the same resources for the next *Reselection Counter* TBs and the same RRI included in the 1st-stage SCI. If the vehicle has to select new resources, it sets to zero the value of the RRI in the 1st-stage SCI of the TB that depleted the *Reselection Counter*. This is done to notify other vehicles that it is not reserving the same resources for the transmission of the next TB. Note that a vehicle using SPS may need to select new resources for the transmission of a new TB even if *Reselection Counter* is not depleted. This happens when the size of a new TB does not fit in the resources previously selected, or when the previously selected resources do not meet the latency requirement of the new TB as detailed in Section II.C.

B. Re-evaluation mechanism

Mode 2 introduces the re-evaluation mechanism to detect and avoid possible collisions in the transmission of a TB. To this aim, vehicles that have selected new resources check whether these resources are still available (i.e., they have not been reserved by another vehicle) before transmitting a TB. If they are not available, they will select new resources to avoid the detected collision. We should note that the re-evaluation mechanism can only operate over selected resources and not reserved ones according to the standard [6]. Re-evaluation applies to both DS and SPS.

The operation of re-evaluation is illustrated in Fig. 2. Let us suppose that the vehicle selects new resources located at slot m . It must then execute again step 1 of the resource allocation process at slot $s_G' = m - T_3$ to check whether the selected resources are still available or they are excluded², where T_3 is equal to $T_{proc,1}$. The execution of step 1 at slot s_G' is referred to as a *re-evaluation check* by the 3GPP standards. The execution of the re-evaluation check results in the definition of a new selection window SW' within the range of slots $[s_G' + T_1, s_G' + T_2]$, where T_2' is defined in the range $T_{2min} \leq T_2' \leq PDB - (s_G' - s_G)$ so that the upper limit of SW' does not overpass the PDB of the TB to be transmitted. Step 1 is executed over the candidate resources in

¹ The resource exclusions due to half-duplex operation have to consider all possible RRI values of the RRI list.

² The standard defines that the execution of step 1 at $m - T_3$ is mandatory. The vehicle could also execute step 1 before $m - T_3$ [6].

SW'. If step 1 reveals that the originally selected resource at slot m is now excluded, then the re-evaluation check has resulted in a *re-evaluation detection* following the 3GPP terminology [11]. The re-evaluation detection triggers the execution of step 2 of the resource allocation algorithm to select new resources among the currently available resources in SW'. As a result, the initially selected resources are replaced by new resources located at, e.g., slot m' in Fig. 2. The execution of step 2 as part of the re-evaluation mechanism is referred to as *resource replacement*.

A vehicle could have selected N resources for the initial transmission of the TB and its retransmissions. If this is the case, when the vehicle performs the re-evaluation check at slot s_G' , it will assess whether the N selected resources are still available. If the re-evaluation detection happens over a subset M of the N selected resources, then the vehicle executes a resource replacement to select M new resources among the available candidate resources in SW'. It is important to note that when the vehicle performs the initial transmission of the TB, it might be announcing the reservation of the $N_{SCI} - 1$ following retransmissions of the TB. As soon as the resources for these retransmissions are reserved, they are no more eligible for a re-evaluation check since re-evaluation only operates over selected and not reserved resources.

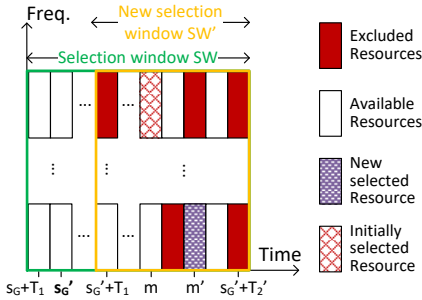


Fig. 2. Operation of re-evaluation under NR V2X mode 2.

It is important to distinguish two different cases where a vehicle can execute a re-evaluation check. The first case happens when the vehicle selects new resources in the selection window. This re-evaluation check is mandatory following 3GPP standards [6] and occurs for both the DS and SPS schemes. The second case only occurs when a vehicle is configured with SPS, and it does not utilize a reservation announced in the 1st-stage SCI. If it later generates a new TB, the transmission of the TB could take place in the resources located RRI ms after the unutilized reservation. Since the reservation of these resources has not been announced, the standard defines that it is 'up to UE implementation' whether the vehicle also executes the re-evaluation check before transmitting the TB [6]. This study considers that vehicles execute the re-evaluation in both cases.

C. Impact of Packet Variability on SPS

SPS reserves the same resources for *Reselection Counter* consecutive TBs with a time gap between reservations of RRI ms. SPS is particularly suited for the transmission of periodic traffic with fixed size. In this case, SPS only reselects resources when the *Reselection Counter* is depleted (depending on $1-P$). We define this event as *counter reselection*. We consider SPS

to be stable if all reservations are utilized to transmit TBs and reselections happen only after the reselection counter depletes. Like in LTE-V2X mode 4 [4], additional reselections may be triggered if the TB size or the inter-arrival time between TBs change. These additional reselections can make SPS more unstable and prone to collisions since neighboring vehicles will not be aware of the newly selected resources until the next TB is transmitted and the transmitting vehicle announces the reservation for the following TBs. Therefore, the probability of packet collisions increases with the number of reselections. It is worth noting that packet variability has no impact on the operation of the DS scheme since vehicles using DS reselect resources for every new TB.

1) Size reselections

SPS triggers a resource reselection when the size of a new TB does not fit in the previously reserved resources. This event is termed *size reselection*. Fig. 3 shows a vehicle V_A that generates a TB (e.g., 200-byte long) at slot s_{G1} , and selects one sub-channel for its transmission at slot s_{R1} . When transmitting the TB at s_{R1} , V_A also announces in the associated SCI that the same sub-channel is reserved at slot $s_{R2} = s_{R1} + RRI$. Let us now suppose that V_A generates at slot s_{G2} a new TB (e.g., 600-byte long) that does not fit the current resource reservation at s_{R2} . V_A is forced to drop the reserved resources at s_{R2} , reselect new resources able to accommodate the size of the new TB (e.g., two sub-channels in Fig. 3), and transmit the TB in the reselected resources at slot s_{R3} .

2) Latency reselections

A vehicle might also need to select new resources if its current reservation is not able to cope with the latency requirements (i.e., the PDB) of a new TB. We refer to this event as *latency reselection*. Latency reselections occur when a vehicle generates aperiodic traffic and the adopted RRI value is larger than the latency deadline or PDB of a TB. The latency reselection is also illustrated in Fig. 3. In this example, V_A has reserved one sub-channel at slot s_{R4} . V_A generates its next TB at slot s_{G3} , and its latency deadline is set at slot s_{LIMIT} . The current reservation at s_{R4} does not respect the latency requirements of the TB since $s_{R4} > s_{LIMIT}$. V_A is then forced to drop the reserved resources at s_{R4} , and select new resources by the latency deadline (e.g., at s_{R5} in Fig. 3).

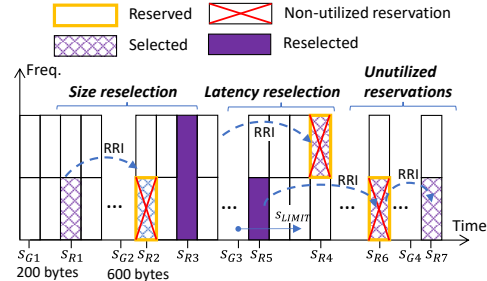


Fig. 3. Impact of packet variability on SPS under NR V2X mode 2.

3) Unutilized reservations

The stability of SPS might be also compromised when reserved resources are not utilized by a vehicle; we term this event as *unutilized reservations* [4]. This can happen because there is no TB ready to be transmitted at the slot where the resources are reserved. Fig. 3 shows that unutilized reservations

occur when the inter-arrival time between the generated TBs is larger than the adopted RRI value. In Fig. 3, V_A has reserved one sub-channel at slot $s_{R6} = s_{R5} + RRI$. However, the next TB is generated at slot s_{G4} with $s_{G4} > s_{R6}$, and V_A leaves the sub-channel at s_{R6} unutilized. We should note that V_A cannot exploit its transmission opportunity at s_{R6} to announce the reservation at slot s_{R7} in the SCI. The transmission of V_A at slot s_{R7} is then prone to packet collisions since it has not been reserved.

III. ANALYSIS OF THE RE-EVALUATION MECHANISM

The re-evaluation mechanism is an important novelty introduced in mode 2 to increase the flexibility in the management of resources and guarantee a more effective scheduling of transmissions. This Section analyzes the operation of the re-evaluation mechanism and discusses the impact that the most relevant mode 2 parameters have on the effectiveness of the re-evaluation check, the re-evaluation detection, and the resource replacement phase.

A. Re-evaluation Check

Vehicles use the re-evaluation check to assess whether selected resources are still available or not right before transmitting the TB. The objective is to detect and avoid potential collisions. 3GPP standards establish that re-evaluation checks are only possible on selected (and not reserved) resources. Accordingly, re-evaluation checks are performed before the transmission of all TBs when using the DS scheme since this strategy selects new resources for each TB. When the SPS scheme is considered, re-evaluation checks affect a smaller number of TBs since SPS only selects new resources as a result of a counter reselection, by design. Once new resources are selected, the remaining TBs are transmitted on reserved resources. If we assume, for example, $P = 0$ and $RRI \geq 100$ ms, only 1 TB out of 10 triggers a counter reselection (the average reselection counter value is 10 in this case), and hence only 10% of the generated TBs are transmitted on selected resources that are eligible for a re-evaluation check. However, we should note that latency reselections, size reselections, and unutilized reservations (see Section II.C) increase the fraction of TBs that are transmitted on selected resources in SPS, and thus increases the number of re-evaluation checks.

Regardless of the scheduling scheme, the fraction of TBs that triggers a re-evaluation check is also affected by the value of T_2 , i.e., by the width of the selection window (see Fig. 1). Let us assume that a vehicle V_A generates a new TB at slot s_{G1} and performs a resource reselection. The selection window is defined by the range of slots $[s_{G1} + T_1, s_{G1} + T_2]$ where $T_1 \leq T_{proc,1}$. In principle, any selected resource included within the selection window shall be eligible for a re-evaluation check. However, a re-evaluation check can be performed only if the vehicle has sufficient processing capabilities to run the entire re-evaluation mechanism before transmitting the TB. According to the 3GPP standard [6], a re-evaluation check can only be performed if the selected resource is included in the $(s_{G1} + T_3, s_{G1} + T_2]$ interval, where T_3 is strictly equal to $T_{proc,1}$. Therefore, the candidate resources included from slot $s_{G1} + T_1$ to slot $s_{G1} + T_3$ are not eligible for a re-evaluation check. Depending on T_2 , the number of resources included in

$[s_{G1} + T_1, s_{G1} + T_3]$ can be a significant fraction of the total number of resources within the selection window. For example, let us assume that $T_2 = PDB = RRI$, and that $\mu = 0$, $T_1 = 1$ slot and $T_3 = 5$ slots. In this case, the percentage of selection window resources that are not eligible for a re-evaluation check is equal to $\{5, 25, 50\}$ % when $RRI = \{100, 20, 10\}$ ms.

B. Re-evaluation Detection

A re-evaluation detection is triggered after a re-evaluation check when the initially selected resources are no longer available. Typically, a re-evaluation detection occurs when the initially selected resources have also been reserved by a neighboring vehicle, and a potential collision is detected. This section sheds light on the circumstances under which a potential collision does and does not trigger a re-evaluation detection. To do so, we separately analyze the re-evaluation detection phase when each TB is transmitted once ($N = 1$) and when it is transmitted twice ($N = 2$, with one blind retransmission) without loss of generality. This section concludes with an insightful discussion about the effectiveness of the re-evaluation detection phase.

1) Single transmission per TB ($N=1$)

We first consider the case where a collision occurs on selected resources. This type of collision cannot be detected by a re-evaluation detection since vehicles do not announce their selection before transmitting on selected resources. This is illustrated in Fig. 4 where V_A and V_B select new resources to transmit their TBs generated at slots s_{G1} and s_{G2} , respectively. If their selection windows (SW_A and SW_B) overlap, the two vehicles may select the same resources at slot s_{R1} , as illustrated in Fig. 4. V_A performs a re-evaluation check at slot $s_{R1} - T_3$, but step 1 does not exclude the resources at slot s_{R1} since V_B has not yet announced its reservation. This is the case because also V_B has performed a reselection after generating the TB at slot s_{G2} and its transmission at slot s_{R1} occurs on selected resources. The re-evaluation mechanism is not capable to detect and avoid the collision at s_{R1} . The same situation occurs when V_B executes its re-evaluation check. If V_A and V_B use the SPS strategy, they will reserve the same resources for transmitting their next TB, at slot s_{R2} , and they will persistently collide until a (counter, latency, or size) reselection occurs if they employ the same RRI (like in Fig. 4). The persistent collision cannot be avoided by the re-evaluation mechanism because it is only executed over selected and not reserved resources.

If vehicles use the DS strategy, they select new resources for every TB when $N=1$. Therefore, they will never be able to detect and avoid a collision. As demonstrated in the remainder of this section, only collisions between selected and reserved resources can be identified by the re-evaluation detection. Since vehicles using DS do not reserve resources, the following discussions apply only to the SPS scheme.

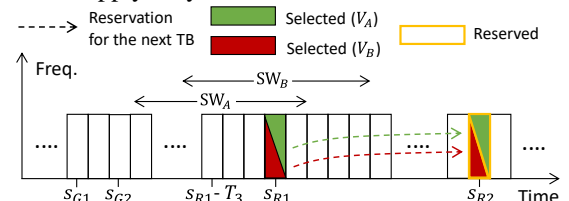


Fig. 4. Persistent collisions not detected by re-evaluation when using SPS.

Depending on the RRI s used by vehicles and the type of generated traffic, we can identify four different cases in which a re-evaluation detects a potential collision. The first case is illustrated in Fig. 5(a) and corresponds to the scenario where vehicles V_A and V_B transmit periodic TBs of fixed size and use the same RRI . V_B selects new resources to transmit the TB generated at slot s_{G1} and selects the resources reserved by V_A at slot s_{R1} . V_B does not exclude the resources reserved by V_A from its selection window (SW_B) during the resource reselection process because V_A announced its reservation in the range of slots $[s_{G1} - T_{proc,0}, s_{G1}]$, i.e., just after the end of V_B 's sensing window. However, V_B can avoid the collision thanks to the re-evaluation check executed at slot $s_{R1} - T_3$. At this time, V_B defines a new sensing window that includes the reservation announced by V_A . Then, the re-evaluation detection identifies the potential collision and V_B excludes the resources reserved by V_A from its new selection window. It is worth pointing out that the probability of this type of re-evaluation detection is very low since the width of the $[s_{G1} - T_{proc,0}, s_{G1}]$ interval is equal to 2 slots for an SCS of 15 or 30 kHz, and equal to 3 slots for a 60 kHz SCS.

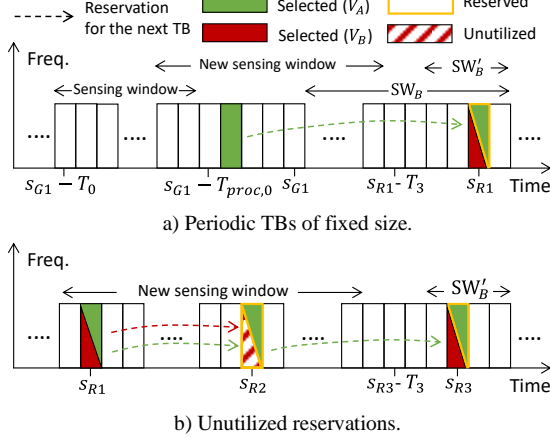


Fig. 5. Re-evaluation detection when vehicles use the same RRI .

The second case is illustrated in Fig. 5(b) and corresponds to the scenario where V_A and V_B are involved in a persistent collision and an unutilized reservation occurs (see Section II.C). In this figure, V_A and V_B initially collide at slot s_{R1} without triggering a re-evaluation detection (like in Fig. 4) and start to persistently collide since they use the same RRI and re-evaluation cannot be applied to reserved resources. Let us suppose that V_B does not generate a TB and leaves the reserved resources at slot $s_{R2} = s_{R1} + RRI$ unutilized. In this case, V_B will not be able to announce the reserved resources at slot $s_{R3} = s_{R2} + RRI$, and will transmit its next TB in selected, rather than reserved, resources. As a result, V_B runs a re-evaluation check at slot $s_{R3} - T_3$, right before transmitting its TB. During the re-evaluation check, the new sensing window of V_B includes the reservation announced by V_A at slot s_{R2} , and V_B excludes the resources reserved by V_A at slot s_{R3} from its new selection window (SW'_B). This triggers a re-evaluation detection that resolves the persistent collision between the two vehicles.

The third and fourth cases where re-evaluation detection successfully detects potential collisions occur when the two vehicles involved use different RRI s. Without loss of

generality, we consider two different RRI values, RRI_1 and RRI_2 , with $RRI_1 < RRI_2$. The third case occurs when vehicles transmit periodic TBs of fixed size, and the two following conditions are satisfied: (i) the resources selected and reserved by a vehicle using the smaller RRI_1 are included within the selection window of a vehicle configured with RRI_2 , and (ii) the vehicle using RRI_2 selects the resources reserved by the vehicle using RRI_1 . This situation is illustrated in Fig. 6(a) where V_A generates a new TB at slot s_{G1} , transmits it on the selected resources at slot s_{R1} , and reserves the same resources at slot $s_{R2} = s_{R1} + RRI_1$. When V_B generates its new TB at slot s_{G2} , it cannot be aware of the reservation announced by V_A at slot s_{R1} due to the overlap between their selection windows (SW_A and SW_B). Let us then suppose that V_B selects the same resources at slot s_{R2} and generates a collision. Note that SW_B is wider than SW_A because V_B uses the largest RRI_2 value. V_B can avoid the collision at slot s_{R2} by executing a re-evaluation check at slot $s_{R2} - T_3$. The new sensing window of V_B will now include the reservation announced by V_A at slot s_{R1} since $s_{R2} - T_3 > s_{R1}$. Then, V_B excludes the resources at slot s_{R2} from its new selection window (SW'_B) and the re-evaluation detection triggers the process to select new resources.

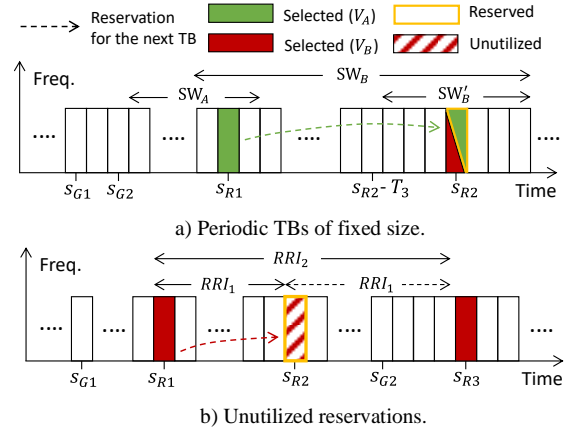


Fig. 6. Re-evaluation detection when vehicles use different RRI s.

The fourth case where re-evaluation detection successfully detects a collision occurs when the two RRI values are multiples of each other (e.g., $RRI_2 = 2 \cdot RRI_1$) and a vehicle leaves one of its reservations unutilized. This is illustrated in Fig. 6(b), where V_B generates a TB at slot s_{G1} , selects the resources for its transmission at slot s_{R1} , and periodically reserves them at slots s_{R2} and s_{R3} using the smallest RRI value, i.e., RRI_1 . Let us suppose that V_B leaves the resources at slot s_{R2} unutilized because it has no TB ready to be transmitted. As a result, V_B cannot reserve the resources at slot s_{R3} , and it will run a re-evaluation check at slot $s_{R3} - T_3$. During the re-execution of step 1, V_B will remove the resources at slot s_{R3} from its selection window due to its half-duplex limitations, as it could not sense the reservations announced from neighboring users at slot s_{R1} , therefore triggering a re-evaluation detection. We should recall from Section II.A that step 1 excludes from the selection window all the slots in which V_A was previously transmitting, considering the entire list of allowed RRI values. Since $s_{R3} = s_{R1} + RRI_2$ and V_B was transmitting at slot s_{R1} , it excludes slot s_{R3} from its selection window.

2) Two transmissions per TB ($N=2$)

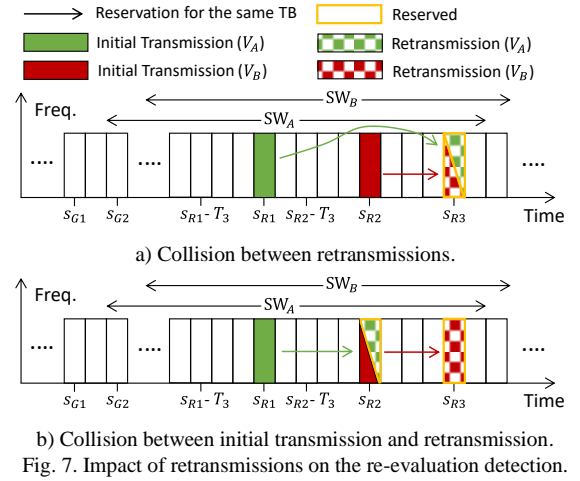
Without loss of generality, this subsection analyzes the impact of retransmissions on the re-evaluation detection considering one blind retransmission per TB (i.e., $N=2$). When $N > 1$, the 1st-stage SCI associated with the TB's initial transmission can reserve the resources used for the retransmission of the same TB if the distance between selected resources is smaller than 32 slots (see Section II.A). In this case, the number of reservations announced by the SCI is indicated with $N_{SCI} = 2$. If the distance between selected resources is larger than 32 slots, the SCI is not able to announce reservations for the retransmission of the same TB and $N_{SCI} = 1$. In $N_{SCI} = 1$ case, the initial transmission and the retransmission of the TB behave as two completely independent events, and no additional collision between selected and reserved resources can occur with respect to the $N=1$ analysis. For this reason, we assume $N_{SCI} = 2$ in the rest of this section. We should also note that vehicles using the DS are allowed to reserve resources for the retransmission of a TB. As a result, the discussion in this section applies to both the SPS and the DS.

We should first note that, like for the $N = 1$ case, re-evaluation cannot detect potential collisions between the initial transmissions of TBs on selected resources when $N > 1$. This is the case because vehicles transmitting on selected resources have not yet announced their selection, and do not allow the re-evaluation mechanism to detect the collision.

In addition to the four cases described when $N = 1$, there are two additional cases when $N = 2$ where the re-evaluation detection can successfully detect a potential collision. These two additional cases originate from potential collisions that involve resources reserved for the retransmission of a TB, and therefore do not depend on the employed RRI values. The first case is illustrated in Fig. 7(a) where a potential collision between the retransmissions of two TBs is considered. In this figure, the initial transmission of V_A and V_B is performed on collision-free resources at slots s_{R1} and s_{R2} , respectively. Due to the overlap between the selection windows of V_A and V_B , let us now assume that the retransmission of both TBs is scheduled on the same resources at slot s_{R3} , potentially leading to a collision. Before transmitting at s_{R2} , V_B runs a re-evaluation check at slot $s_{R2} - T_3$ and senses the reservation announced by V_A for the retransmission of the same TB; this reservation is announced by the SCI associated with the TB's initial transmission. V_B triggers then a re-evaluation detection to select new resources for the retransmission. Note that also V_A runs a re-evaluation check at slot $s_{R1} - T_3$, but it cannot sense the reservation announced by V_B because $s_{R2} > s_{R1} - T_3$.

The second case occurs when there is a potential collision between the retransmission and the initial transmission of TBs, and is illustrated in Fig. 7(b). In the figure, V_A selects resources at slots s_{R1} and s_{R2} for the initial transmission and the retransmission of a TB, while V_B selects resources at slots s_{R2} and s_{R3} for the initial transmission and the retransmission of a TB. V_B runs a re-evaluation check at slot $s_{R2} - T_3$ and senses the reservation announced by V_A at slot s_{R1} . This reservation included the resources initially selected (and now reserved, since $s_{R2} - T_3 > s_{R1}$) at s_{R2} for the retransmission of the TB by V_A . V_B detects the possible collision between its initial transmission and the retransmission of V_A , excludes the

resources initially selected at s_{R2} from its new selection window, and triggers a re-evaluation detection.



3) Effectiveness of re-evaluation detections

This section has identified and analyzed carefully all the circumstances under which a collision can (and cannot) trigger a re-evaluation detection. However, a re-evaluation detection is not always effective in avoiding collisions. An ineffective re-evaluation detection occurs if the reservations that triggered a re-evaluation detection are not finally used for transmitting a TB. To further clarify the notion of effective re-evaluation detection, let us consider the scenario illustrated in Fig. 6(a). In this figure, V_B triggers a re-evaluation detection because it detected the imminent collision with V_A at slot s_{R2} . Then, V_B completes the re-evaluation process to select new resources and avoid the collision. If V_A eventually transmits its next TB using the reserved resources at slot s_{R2} , then the re-evaluation detection triggered by V_B was effective in avoiding the collision with V_A . Conversely, let us now suppose that the next TB of V_A does not fit in the resources reserved at s_{R2} , and V_A must perform a size reselection to reserve new resources able to accommodate the size of the new TB. In this case, the resources at slot s_{R2} are unutilized since both V_A and V_B selected new resources. In this case, the re-evaluation detection has been ineffective since it did not avoid any collision between V_A and V_B . Re-evaluation detections would also be ineffective if V_A performs a latency reselection or leaves unutilized the resources that it has reserved at s_{R2} . It is important to point out that vehicles cannot determine in advance if a re-evaluation detection will be ultimately effective or not, except when it is triggered by a reservation for the retransmission of the same TB. Reservations for the retransmission of the same TB always satisfy the size and latency requirements of the generated TB, and they are not subject to latency reselections, size reselections, or unutilized reservations. Therefore, a re-evaluation detection triggered by a retransmission of the same TB is always effective.

C. Resource Replacement

During a re-evaluation, if a vehicle detects a potential collision it triggers the re-execution of step 2 of the resource allocation algorithm as part of the resource replacement phase.

The objective is to select new collision-free resources and avoid the identified collision; however, the selection of collision-free resources cannot be fully guaranteed, as explained in Section II.C. During the resource replacement phase, a vehicle might select resources that are already occupied by neighboring vehicles and experience a collision on selected resources that cannot be detected by the re-evaluation mechanism.

Therefore, the selection of collision-free resources during the resource replacement phase is instrumental to the effectiveness of the re-evaluation mechanism. Since such collision-free selection cannot be always guaranteed, it is necessary to evaluate the actual effectiveness of the re-evaluation mechanism.

IV. SIMULATION ENVIRONMENT

The operation and impact of the re-evaluation mechanism is evaluated using a 5G NR V2X mode 2 simulator implemented by the authors in ns-3. The implemented simulator is fully compliant with the 5G NR V2X mode 2 standard. 5G NR V2X is configured to operate over a 20 MHz channel with a subcarrier spacing of 30 kHz in the 5.9 GHz frequency band. The sub-channel size is set to 12 RBs, and there are then 4 sub-channels per slot. The transmission power has been set to 23 dBm and the sensitivity to -103.5 dBm, according to the prototype data in [12]. The pathloss is modeled using the reference 3GPP pathloss model [2]. The shadowing effects are modeled using a log-normal distribution with zero mean and a standard deviation of 3 dB. Shadowing spatial correlation is modeled following the 3GPP guidelines in [2]. We consider that each TB is transmitted using 16QAM and a coding rate equal to 0.5. We model the PHY layer performance using lookup tables from 3GPP working documents that relate the Block Error Rate (BLER) vs Signal to Interference to Noise Ratio (SINR). We use the lookup tables from [13] for the transmission of TBs and the ones from [14] for the SCIs.

This study considers the reference 3GPP 5 km highway scenario with 3 lanes in each direction. We analyze densities of 25, 50 and 100 veh/km, and in all these scenarios the vehicle speed is set to 70 km/h. Vehicles transmit TBs following the 3GPP periodic and aperiodic traffic models [2]. The periodic model considers 190-byte TBs generated with a constant inter-packet arrival time; the latency requirement or PDB is set equal to the inter-packet arrival time. We refer to this traffic as periodic of fixed packet size. The aperiodic traffic model considers TBs generated with an inter-packet arrival time $\tau = c + r$, where c is a constant and r is an exponentially distributed random variable. The PDB for the aperiodic traffic is set to c . The size of a TB for the aperiodic traffic is uniformly distributed in the [200,1200] byte range, with a 200-byte step. We refer to this traffic as aperiodic of variable size. For periodic and aperiodic traffic, we consider two different scenarios: single and mixed traffic. In the single traffic scenario, all vehicles generate traffic with an average inter-packet arrival time of 100 ms. For periodic traffic, the inter-packet arrival time is constant. In the aperiodic traffic case, we set $c = \bar{r} = 50$ ms. In the mixed traffic scenario, 80% of vehicles have an average inter-packet arrival time of 100 ms, and the remaining 20% have an average inter-packet arrival time of 20 ms ($c = \bar{r} = 10$ ms for the aperiodic traffic).

We evaluate the performance of the re-evaluation mechanism for the SPS and DS scheduling schemes. For both schemes, we set the processing delay times $T_{proc,0}$, T_0 and T_3 equal to 1 slot, 1100 ms (equivalent to 2200 slots with a subcarrier spacing of 30 kHz) and 5 slots respectively. The limits of the selection window T_1 and T_2 are set equal to 2 slots and to the PDB, respectively. The percentage X of resources that must be available after the execution of step 1 of the resource allocation algorithm is set to 20%. The threshold RSRP is set to its minimum value, i.e. -128 dBm, following the results obtained in [8]. We evaluate the impact of retransmissions on the performance of re-evaluation considering N equal to 2. For the SPS scheme, the probability P to keep the same resources has been set to 0, and we evaluate two different strategies for the selection of the *RRI* [15]:

- Average *RRI*: the *RRI* is set equal to the average inter-packet arrival packet time.
- Minimum *RRI*: the *RRI* is set equal to the minimum of the inter-packet arrival time. This strategy seeks to avoid latency reselections (see Fig. 3 in Section II.C).

Note that the two *RRI* strategies result in the same value of the *RRI* with periodic traffic since the inter-packet arrival time is constant. However, with aperiodic traffic, the average *RRI* strategy sets the *RRI* value equal to $c + \bar{r}$, while the minimum one sets it equal to c . In the single traffic scenario, all vehicles are configured with a single *RRI* value (following the average or minimum *RRI* strategy) to support the 100 ms average inter-packet arrival time. In the mixed traffic scenario, vehicles are configured with two different *RRI* values to support the 100 ms and 20 ms average inter-packet arrival time.

We define the following evaluation metrics:

- Packet Delivery Ratio (PDR): fraction of correctly received TBs over the total number of transmitted TBs. In the case of retransmissions, a TB is labeled as correctly received if at least 1 out of the N transmissions is correctly received.
- PDR-Re-evaluation: PDR of specific TBs for which at least a re-evaluation has been detected. In the case of retransmissions, this PDR is obtained at the MAC level for each of the N TB (re)transmissions.
- Packet Collision Ratio (PCR): fraction of TBs that are incorrectly received due to packet collisions over the total number of transmitted TBs. This error occurs when the TB cannot be correctly decoded because the SINR is too low due to the interference generated by other vehicles.
- Re-evaluation Check Ratio (ReCR): fraction of TBs that have been checked for re-evaluation at least once over the total number of transmitted TBs.
- Re-evaluation Detection Ratio (ReDR): fraction of TBs that experience at least 1 re-evaluation detection over the total number of transmitted TBs.
- Ineffective Re-evaluation Detection Ratio (IReDR): fraction of TBs over which at least 1 re-evaluation was detected but the reservations that triggered the re-evaluation detections are not finally utilized for transmitting a TB (see Section III.B.3).
- Size reselection ratio (SRR): fraction of TBs that produce a size reselection over the total number of transmitted TBs [4].
- Latency reselection ratio (LRR): fraction of TBs that produce a latency reselection over the total number of transmitted TBs.

> REPLACE THIS LINE WITH YOUR MANUSCRIPT ID NUMBER (DOUBLE-CLICK HERE TO EDIT) <

- **Unutilized Reservation Ratio (URR)**: fraction of unused reservations over the total number of reserved resources. URR does not account for unutilized reservations that are considered in the SRR and LRR metrics [4].

V. IMPACT OF RE-EVALUATIONS ON SPS WITHOUT RETRANSMISSIONS

This section analyzes the impact of re-evaluations on the operation and performance of SPS when $N = 1$, i.e., when each TB is transmitted once with no retransmissions. We focus first on the mixed traffic scenario with vehicles transmitting aperiodic traffic of variable size. This is a key target scenario since most V2X services to be supported by NR V2X generate this type of traffic, and this traffic can create instability in the operation of SPS due to frequent unutilized reservations as well as size and latency reselections. This instability increases the probability of packet collisions, and re-evaluation was introduced to avoid such collisions.

The variability introduced by aperiodic traffic of variable size results in that more than 50% of the packets generated by the vehicles are transmitted in selected (and hence not reserved) resources and are hence eligible for a re-evaluation check. This is visible in Table I.a which reports the different metrics for the two *RRI* selection strategies and all traffic densities. We should note that the ReCR, SRR, LRR and URR metrics do not vary with the vehicle density because they only depend on the traffic and on the reservations that each vehicle individually generates. The table shows that the ratio of re-evaluation checks (ReCR metric) is higher than 50% for both *RRI* selection strategies. Vehicles execute a large number of re-evaluation checks because they transmit a large number of packets in selected resources. This is due to a large number of size and latency reselections or unused reservations (see SRR, LRR, URR in Table I.a). The average *RRI* strategy reduces the ratio of unutilized reservations (URR) but augments the size and latency reselections (SRR and LRR), while the minimum *RRI* strategy minimizes SRR and LRR at the cost of increasing URR. Table I.a also shows that re-evaluation is able to detect a larger number of packet collisions (ReDR) as the vehicular density increases. For example, re-evaluation detects collisions on over 16% of the packets with 100 veh/km and the average *RRI* strategy. This percentage increases to over 44% with the minimum *RRI* selection strategy.

Table I shows that re-evaluation detects a large number of potential packet collisions (ReDR). However, Fig. 8 shows that re-evaluation is not fully effective in avoiding collisions and in improving the packet delivery ratio; this is independent of the *RRI* selection strategy. Fig. 8 compares the performance when re-evaluation is implemented and when it is not. Fig. 8(a) and Fig. 8(b) plot the PDR for two traffic densities and Fig. 8(c) the PCR for one of these densities. The figure shows that the performance is nearly identical when utilizing re-evaluation and when not. There are several reasons why re-evaluation is not effective in avoiding packet collisions and improving the PDR with aperiodic traffic of variable size. First, re-evaluation cannot detect collisions between two vehicles that are selecting new resources since these vehicles have not yet announced their selection. The second reason is that packet variability can produce size and latency reselections and increase the

TABLE I. PERFORMANCE METRICS (IN %) OF SPS WHEN $N=1$

<i>RRI</i> strategy	ReCR	SRR	LRR	URR	25 veh/km ReDR	50 veh/km ReDR	100 veh/km ReDR	25 veh/km IReDR	50 veh/km IReDR	100 veh/km IReDR
Avg <i>RRI</i>	60.9	27	57	4	10.7	7.7	14.2	10.2	16.2	11.7
Min <i>RRI</i>	57.6	4	3	55	37.3	23.5	41.9	26.4	44.6	28.1

B) PERIODIC TRAFFIC OF FIXED SIZE

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

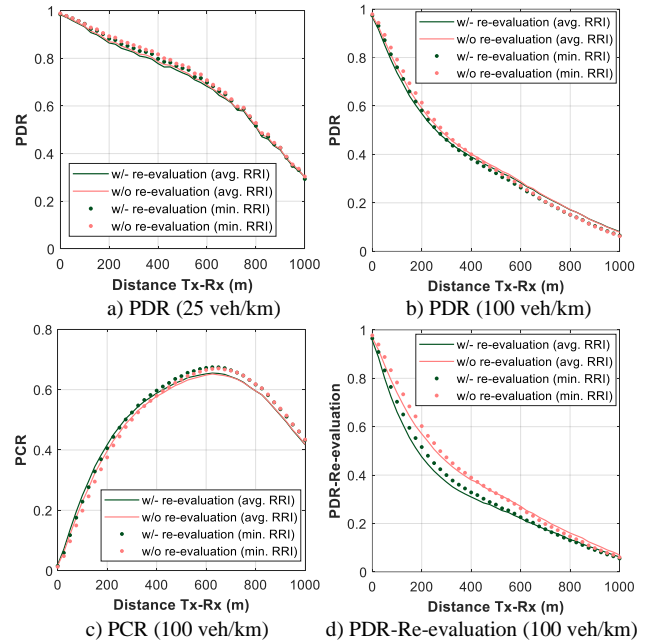


Fig. 8. SPS performance in mixed traffic scenario with aperiodic traffic of variable size, $N = 1$.

probability of having to select new resources. Since re-evaluation cannot detect collisions between vehicles that are selecting new resources, the packet variability increases the probability of having collisions that cannot be detected by re-evaluation. In addition, we should note that re-evaluations may not be effective if the reservations that triggered a re-evaluation detection are not finally used for transmitting a TB. In this case, vehicles change resources to avoid a collision that never happened, and we cannot guarantee when changing resources that an undetectable collision will not happen in the newly selected resources. In our analysis, 72% and 63% of the resource reservations that triggered re-evaluation detections for the average and minimum *RRI* strategies, respectively, were not finally used for transmitting a TB under all evaluated vehicle densities. The ineffectiveness of the re-evaluation mechanism is reflected in the IReDR metric reported in Table I.a., and negatively impacts the PDR of the TBs for which at least a re-evaluation has been detected (PDR-Re-evaluation in Fig. 8(d)). Fig. 8(d) shows that the PDR of the TBs that perform a resource replacement after a re-evaluation detection degrades compared to the PDR measured when re-evaluation is not implemented.

We analyze now the impact of re-evaluations on SPS when vehicles transmit periodic traffic of fixed size. Periodic traffic of fixed size does not generate undetected collisions due to size and latency reselections as it was the case of aperiodic traffic of variable. The impact of these undetected collisions that are not resolved by re-evaluation can be visualized in Fig. 9 which compares the PDR with periodic traffic of fixed size and aperiodic traffic of variable size for the same vehicular density when re-evaluation is implemented. The figure clearly shows how these undetected collisions reduce the PDR under aperiodic traffic of variable size, and their impact increases with the vehicular traffic.

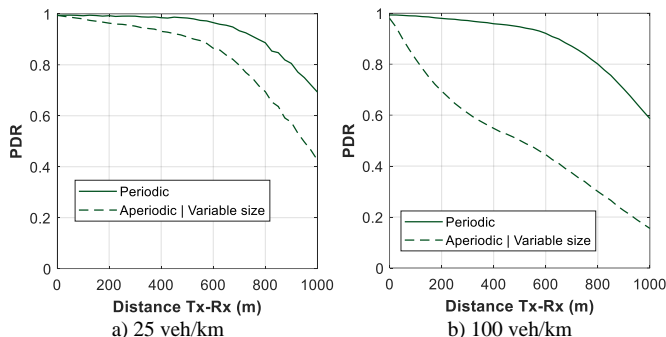


Fig. 9. PDR for periodic traffic of fixed size and aperiodic traffic of variable size in single traffic scenario ($N = 1$ and minimum RRI strategy). Similar trends are observed in mixed traffic scenario and with the average RRI strategy.

Periodic traffic of fixed size can be affected by persistent collisions that occur when various vehicles select the same resources within overlapping selection windows³. These persistent collisions cannot be detected by re-evaluation, since re-evaluation cannot detect collisions between selected resources as explained in Section III.B. With periodic traffic of fixed size, collisions persist until one of the vehicles depletes its *Reselection Counter* and executes a resource reselection. We should note that only TBs transmitted after the *Reselection Counter* depletes are eligible for a re-evaluation check since they are transmitted on selected resources. With $RRI = 100$ ms, the *Reselection Counter* range is $[5, 15]$, and the ReCR is on average equal to 10% for the single traffic scenario (see Table I.b); similar trends are observed for the mixed traffic scenario. Out of the limited set of TBs that are eligible for a re-evaluation check, a vehicle can only use re-evaluations to detect a collision under the conditions illustrated in Fig. 5(a) (Section III.B.1). These conditions require that the reservation that causes the collision is made by a vehicle in a 2-slot time interval just before the generation of the TB. This unlikely condition results in the low ratio of re-evaluation detections (ReDR) reported in Table I.b and the small impact of re-evaluation on the PCR in Fig. 10(a). Nevertheless, the vehicles that did execute re-evaluation avoided the persistent packet collisions generated by an initial collision between a selected and a reserved resource. The avoided persistent packet collisions affected on average the transmission of 5.65 consecutive TBs (100 veh/km, $RRI = 100$ ms). Fig. 10(b) reports the PDR evaluated for the TBs over

which a re-evaluation was detected. Fig. 10(b) shows significant gains compared to the performance obtained if re-evaluations were not implemented. In this case, re-evaluations were effective to avoid the limited set of packet collisions detected with periodic traffic of fixed size.

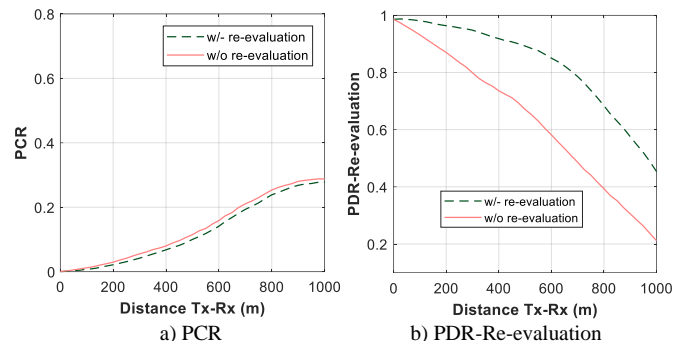


Fig. 10. SPS performance in mixed traffic scenario for periodic traffic, 100 veh/km, $N = 1$. Similar trends are observed under the single traffic scenario.

VI. IMPACT OF RE-EVALUATIONS ON SPS WITH RETRANSMISSIONS

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

Table II reports the performance metrics when $N=2$ and the traffic is aperiodic and of variable size. The table shows that retransmissions generate many more re-evaluation detections: ReDR increases to more than 25% in the single and mixed traffic scenarios compared to 10.7% when $N = 1$ (see Table I.a). Traffic variability can still impact the initial transmission of TBs when $N=2$. However, retransmissions do not generate unutilized reservations or size and latency reselections as the resources reserved for retransmissions always fit the requirements of the retransmitted TB both in size and time. This brings some stability to the operation of SPS which benefits the operation of re-evaluation. In particular, reservations made to transmit the retransmission of the same TB always hold a transmission. In this case, re-evaluation detections are always effective since they avoid an imminent collision. The conducted simulations show that in the single traffic scenario more than 88% of the re-evaluation detections are triggered by reservations made for the retransmission of the same TB. Since re-evaluations are always effective in avoiding this collision,

³ Using [16], we can estimate that around 30% of packets that trigger a resource reselection would experience persistent collisions with 100 veh/km.

> REPLACE THIS LINE WITH YOUR MANUSCRIPT ID NUMBER (DOUBLE-CLICK HERE TO EDIT) <

the PDR for the packets that detected a re-evaluation (PDR-Re-evaluation) significantly outperforms the PDR without re-evaluation (Fig. 11(a)); this was not the case without retransmissions ($N=1$) as shown in Fig. 8(d). Fig. 11(a) shows that re-evaluations improve the PDR for both initial transmissions and retransmissions that detected re-evaluations when the single traffic scenario is considered; for example, the improvement is equal to 53% and 70% when the Tx-Rx distance is 300 m and the density is 50 veh/km. In the mixed traffic scenario, reported in Fig. 11(b), less than 37% of the detected re-evaluations are caused by reservations for the retransmission of a TB (compared to more than 88% in the single traffic scenario). The remaining re-evaluation detections are triggered by reservations for the next TB. Reservations for the next TB do not always hold a transmission in the reserved resources and affect the effectiveness of the re-evaluation mechanism. This explains the higher IReDR values in the mixed traffic scenario compared to the single traffic scenario (Table II) as well as the lower positive impact of re-evaluation in Fig. 11(b) compared to Fig. 11(a).

TABLE II. PERFORMANCE METRICS (IN %) OF SPS FOR APERIODIC TRAFFIC OF VARIABLE SIZE WHEN $N=2$ (AVERAGE RRI STRATEGY)

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

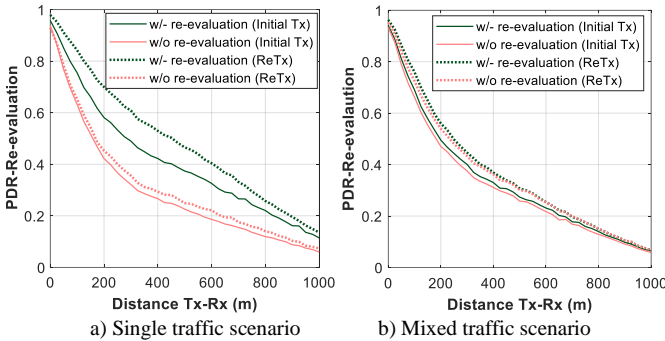


Fig. 11. PDR-Re-evaluation experienced by SPS for aperiodic traffic of variable size when $N = 2$ (50 veh/km, average RRI strategy). Similar trends have been obtained for other densities.

The obtained results show that re-evaluations are effective in avoiding collisions on retransmissions. However, re-evaluation can only improve the PDR with $N=2$ if: 1) both the initial transmission and the retransmission experience a collision (without re-evaluation, a packet is correctly received if just one of the two transmissions is correctly received); 2) re-evaluation can detect at least one of the two collisions; and 3) the resource replacement is effective in avoiding a collision. For the single traffic scenario, 20% and 26% of TBs experienced a collision in their initial transmission and retransmission, and re-evaluation detected at least one of them, for densities of 50 veh/km and 100 veh/km, respectively. Despite these non-negligible percentages, Fig. 12(a) shows that re-evaluation does not significantly improve the PDR. This is because the resource replacements ultimately did not avoid a collision with aperiodic traffic of variable size. We should not forget that following a

resource replacement, a vehicle selects a new resource and is therefore prone to new potential undetected collisions.

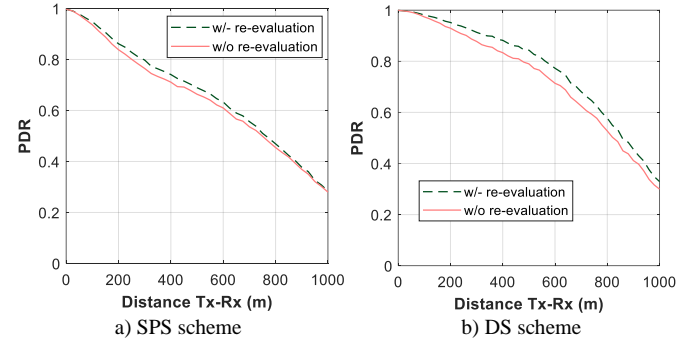


Fig. 12. PDR experienced by mode 2 for aperiodic traffic of variable size when $N = 2$, 50 veh/km (average RRI strategy for SPS scheme) in the single traffic scenario. Similar trends have been obtained for other densities.

For periodic traffic of fixed size, re-evaluation is again effective in avoiding collisions. However, like for $N=1$, the impact on the PDR is small because the fraction of TBs that experience at least one re-evaluation detection (ReDR) is very low (below 2%).

VII. IMPACT OF RE-EVALUATIONS ON DS

Vehicles using the DS scheme always transmit the generated TBs on selected resources when only one transmission per TB is considered ($N = 1$). Collisions that occur between selected (not reserved) resources do not trigger any re-evaluation detection (Section III.B). Re-evaluation has therefore no impact or benefit when using DS with $N = 1$. We then analyze the impact of re-evaluations on the DS when considering retransmissions ($N = 2$), since retransmissions occur on reserved resources and can trigger a re-evaluation detection. This section considers aperiodic traffic of variable size. However, we should note that the performance of DS does not depend on the traffic pattern since DS selects new resources for the initial transmission and the retransmission of every TB. This also entails that DS does not experience any size reselections, latency reselections, or unutilized reservations.

Table III reports the ratio of re-evaluation checks (ReCR) and detections (ReDR) that characterize the DS in the different settings considered in this work. Table III shows that the ReCR is equal to 96.6% in the single traffic scenario, i.e., a much larger value with respect to its SPS counterpart in Table II. Such an increase in the ReCR occurs because almost every TB is transmitted on selected resources and is therefore eligible for a re-evaluation check when the DS is considered. This was not the case with SPS because TBs are transmitted on selected resources only after an unutilized reservation or a (counter, size, latency) reselection. Table III also shows that DS is characterized by fairly large ReDR values in the single and mixed traffic scenarios. The ReDR values increase with the vehicular density. A larger density increases the probability that several vehicles select the same resources, and therefore increases the number of potential collisions. With respect to its single traffic counterpart, the ReCR decreases in the mixed traffic scenario (similarly to the SPS case). During a reselection, vehicles with a smaller RRI have a larger probability of

selecting resources that are not eligible for a re-evaluation check (see Section III.A). With DS, the mixed traffic scenario does not experience additional re-evaluation detection opportunities compared to SPS. As a result, a smaller ReCR implies a reduction in the measured ReDR levels with respect to the single traffic scenario (Table III). Such ReDR reduction is more evident at larger densities.

Fig. 13 depicts the impact of re-evaluations on the PCR when using DS; the PCR is measured separately for the initial transmission of a TB and its retransmission. Fig. 13(a) shows that re-evaluation can improve the PCR of both initial transmissions and retransmissions when the channel is lightly loaded. The figure reveals that re-evaluations are more effective in reducing the PCR experienced by retransmissions since initial transmissions of a TB are accommodated over selected resources and are more prone to experience undetected collisions. This effect is more visible in Fig. 13(b) which corresponds to the highest vehicle density. This figure shows that re-evaluations can have a negative impact on the PCR of initial transmissions under high channel loads whereas it improves the PCR of retransmissions.

With DS, re-evaluation detection is always effective, and re-evaluation improves the PDR for the TBs for which at least a re-evaluation has been detected (PDR-Re-evaluation). However, the impact of re-evaluations on the PDR is limited also in the DS case, as shown in Fig. 12(b): like for SPS, the re-evaluation mechanism can improve the PDR only if both the initial transmission of a TB and its retransmission experience a collision and re-evaluation can detect at least one of them. Despite the large ReDR values reported in Table III, this occurs for only the 0.35%, 2.8% and 9.8% of the TBs for the 25 veh/km, 50 veh/km and 100 veh/km densities, respectively. In addition, the impact of re-evaluation on the PDR is limited by the accuracy of the resource replacement phase. As illustrated in Fig. 13, the selection of collision-free resources during the resource replacement phase is not guaranteed (especially when the channel load is large) and vehicles are prone to experience potentially undetected collisions after the resource replacement.

TABLE III. PERFORMANCE METRICS (IN %) OF DS FOR APERIODIC TRAFFIC OF VARIABLE SIZE WHEN $N=2$

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

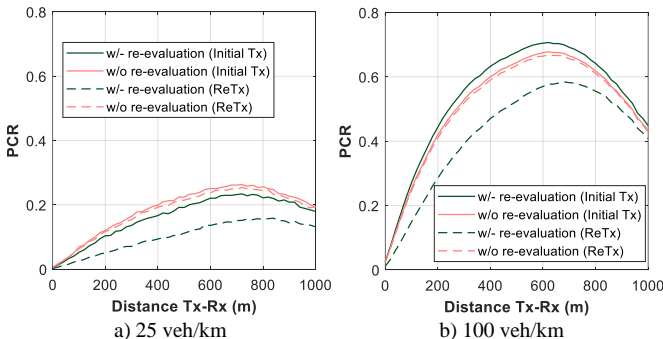


Fig. 13. PCR experienced by DS with aperiodic traffic of variable size in the single traffic scenario when $N=2$.

VIII. CONCLUSIONS

This paper has presented a comprehensive analysis and evaluation of the impact of the re-evaluation mechanism included in 3GPP Release 16 on the operation and performance of 5G NR V2X mode 2.

The study has shown that the impact of re-evaluation and its effectiveness to avoid packet collisions depends on the traffic patterns and NR V2X mode 2 configurations. Without retransmissions, re-evaluations detect collisions only with the semi-persistent scheduling (SPS) scheme. In this case, re-evaluation can detect many collisions due to size and latency reselections as well as unutilized reservations. However, re-evaluation does not improve the performance since traffic variability increases the probability of having to select new resources (due to size and latency reselections), and re-evaluation cannot detect collisions between vehicles that are selecting new resources since these vehicles have not yet announced their selection. Size and latency reselections also reduce the effectiveness of re-evaluation since many reservations that triggered a re-evaluation detection are not finally used for transmitting a TB. In this case, vehicles change resources to avoid a collision that actually never happened, and it is not possible to guarantee that an undetectable collision will not happen in the newly selected resources.

The operation of SPS is more stable when vehicles generate periodic traffic of fixed size, and this reduces significantly the number of collisions and improves the effectiveness of re-evaluation in avoiding detected collisions. However, re-evaluation can only detect a very limited amount of collisions, so its overall impact on the performance of SPS with periodic traffic of fixed size is negligible.

With retransmissions, re-evaluation is more capable to detect collisions in SPS and DS and avoid them effectively, even under the presence of aperiodic traffic of variable size. This is the case because retransmissions do not change the size or latency of the initial transmissions, and hence reservations made for the retransmission of a TB always hold a transmission. However, the impact of re-evaluation on the performance of SPS and DS with retransmissions is low. This is the case because resource replacements are not effective in avoiding a collision after selecting new resources. Without re-evaluation, a packet is correctly received if just one of the two transmissions is correctly received.

The results presented in this study serve as a reference to understand when (scenario and mode 2 configurations) and how re-evaluation is effective in detecting and avoiding collisions. However, this study has demonstrated that re-evaluation does not ultimately provide significant benefits for 5G NR V2X mode 2, and is not that effective in avoiding packet collisions. We should also note that the implementation of re-evaluation (mandatory according to current 3GPP standards) implies a significant computational cost as a result of frequent re-evaluation checks and resource reselections. It is therefore questionable whether re-evaluation (in its current format) is beneficial for improving the reliability of 5G NR V2X mode 2 sidelink communications.

REFERENCES

- [1] 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*, vol. 23, no.3, pp.1972-2026, Feb. 2021.
- [2] 3GPP, "TR 37.885 "Study on evaluation methodology of new Vehicle-to-Everything (V2X) use cases for LTE and NR (v15.3.0, Release 15)", 3GPP, Tech. Rep., June 2019.
- [3] 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.
- [4] 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, July 2020.
- [5] L. Lusvarghi and M. L. Merani, "Machine Learning for Disseminating Cooperative Awareness Messages in Cellular V2V Communications," *IEEE Transactions on Vehicular Technology*, vol. 71, no. 7, pp. 7890-7903, July 2022.
- [6] 3GPP, "TS 38.321 NR; Medium Access Control (MAC) protocol specification (v16.7.0, Release 16)," 3GPP, Tech. Spec., Jan. 2022.
- [7] Z. Ali, S. Lagén, L. Giupponi and R. Rouil, "3GPP NR V2X Mode 2: Overview, Models and System-Level Evaluation," *IEEE Access*, vol. 9, pp. 89554-89579, Nov. 2021.
- [8] 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, Oct. 2021.
- [9] C. Campolo, V. Todisco, A. Molinaro, A. Berthet, S. Bartoletti and A. Bazzi, "Improving Resource Allocation for beyond 5G V2X Sidelink Connectivity," *Proceedings of the 2021 55th Asilomar Conference on Signals, Systems, and Computers*, Nov. 2021, pp. 55-60.
- [10] M. Muhammad Saad, M. Ashar Tariq, M. Mahmudul Islam, M. Toaha Raza Khan, J. Seo and D. Kim, "Enhanced Semi-persistent scheduling (e-SPS) for Aperiodic Traffic in NR-V2X," *Proceedings of the 2022 International Conference on Artificial Intelligence in Information and Communication (ICAIIIC)*, Feb. 2022, pp. 171-175.
- [11] 3GPP, "TS 38.214 NR; Physical layer procedure for data (v16.8.0, Release 16)," 3GPP, Tech. Spec., Jan. 2022.
- [12] 5GAA, "V2X Functional and Performance Test Report; Test Procedures and Results," 5GAA Report, Apr. 2019.
- [13] Huawei, HiSilicon, "R1-1900852. Link level evaluations on sidelink for NR V2X," 3GPP TSG RAN WG1 Ad-Hoc Meeting 1901, Taipei, Jan. 2019.
- [14] Ericsson, "R1-1903180. Link level evaluations of NR PSCCH," 3GPP TSG-RAN WG1 Meeting #96, Athens, Greece, March 2019.
- [15] 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?", *Proceedings of the 2022 IEEE 23rd International Conference on High-Performance Switching and Routing (HPSR 2022)*, Taicang, Jiangsu, China, 6-8 June 2022, virtual conference.
- [16] M. Gonzalez-Martín, M. Sepulcre, R. Molina-Masegosa and J. Gozalvez, "Analytical Models of the Performance of C-V2X Mode 4 Vehicular Communications," *IEEE Transactions on Vehicular Technology*, vol. 68, no. 2, pp. 1155-1166, Feb. 2019.