

UNIVERSIDAD MIGUEL HERNÁNDEZ DE ELCHE

**Programa de Doctorado en Tecnologías Industriales y de
Telecomunicación**



**5G NR V2X Communications for Connected
and Automated Vehicles**

Autor: Alejandro Molina Galán

Director: Prof. Javier Manuel Gozávez Sempere

Codirector: Prof. Baldomero Coll Perales

- 2024 -

This doctoral thesis is supported by a compendium of works that have been previously published, following the regulations of Universidad Miguel Hernández de Elche. The body of this thesis is made up of the articles whose bibliographic references and quality indicators are provided below.

Articles published in high-impact journals, indexed according to JCR Science Edition:

- *A Tutorial on 5G NR V2X Communications*
M. H. C. Garcia, A. Molina-Galan, M. Boban, J. Gozalvez, B. Coll-Perales, T. Sahin, and A. Kousaridas
IEEE Communications Surveys & Tutorials
Vol. 23, Issue 3, pp. 1972-2026 (February 2021).
ISSN: 1553-877X. Editorial: IEEE
Impact Factor JCR-SCI: 33.840, Quartile Q1
DOI: 10.1109/COMST.2021.3057017
- *On the Impact of Re-Evaluation in 5G NR V2X Mode 2*
A. Molina-Galan, L. Lusvarghi, B. Coll-Perales, J. Gozalvez and M. L. Merani
IEEE Transactions on Vehicular Technology
Vol. 73, Issue 2, pp. 2669-2683 (February 2024).
ISSN: 1939-9359. Editorial: IEEE
Impact Factor JCR-SCI: 6.8, Quartile Q1
DOI: 10.1109/TVT.2023.3318235

Articles published in international conferences:

- *How does 5G NR V2X Mode 2 Handle Aperiodic Packets and Variable Packet Sizes?*
A. Molina-Galan, B. Coll-Perales, L. Lusvarghi, J. Gozalvez and M. L. Merani
2022 IEEE 23rd HPSR
06-08 June 2022 (Taicang, Jiangsu, China) (Virtual)
ISBN: 978-1-6654-0607-9. Editorial: IEEE
DOI: 10.1109/HPSR54439.2022.9831314
- *Re-Evaluation Strategies for 5G NR V2X Communications*
A. Molina-Galan, B. Coll-Perales and J. Gozalvez
2022 IEEE 96th Vehicular Technology Conference (VTC2022-Fall)
26-29 September 2022 (London, United Kingdom)
ISBN: 978-1-6654-5468-1. Editorial: IEEE
DOI: 10.1109/VTC2022-Fall57202.2022.10012842



UNIVERSITAS
Miguel Hernández

AUTORIZACIÓN DE PRESENTACIÓN DE TESIS DOCTORAL POR COMPENDIO DE PUBLICACIONES

Director: Prof. Javier Manuel Gozávez Sempere

Codirector: Prof. Baldomero Coll Perales

Título de la tesis: **5G NR V2X Communications for Connected and Automated Vehicles**

Autor: Alejandro Molina Galán

Programa de doctorado en Tecnologías Industriales y de Telecomunicación

Instituto de Investigación en Ingeniería de Elche

Universidad Miguel Hernández de Elche

El director y codirector de la tesis reseñada AUTORIZAN SU PRESENTACIÓN EN LA MODALIDAD DE COMPENDIO DE PUBLICACIONES.

En Elche, a ____de julio de 2024.

Fdo. Dr. Javier Manuel Gozávez Sempere Fdo. Dr. Baldomero Coll Perales



UNIVERSITAS
Miguel Hernández

INFORME DEL COORDINADOR DEL PROGRAMA DE DOCTORADO
EN TECNOLOGÍAS INDUSTRIALES Y DE TELECOMUNICACIÓN

Dr. D. ÓSCAR REINOSO GARCÍA, coordinador del Programa de Doctorado en
Tecnologías Industriales y de Telecomunicación (TECNIT) de la Universidad Miguel
Hernández de Elche,

CERTIFICA

Que el trabajo realizado por D. Alejandro Molina Galán titulado **5G NR V2X Communications for Connected and Automated Vehicles** ha sido dirigido por el Prof. D. JAVIER MANUEL GOZÁLVEZ SEMPERE y codirigido por el Prof. D. BALDOMERO COLL PERALES, y se encuentra en condiciones de ser leído y defendido como Tesis Doctoral ante el correspondiente tribunal en la Universidad Miguel Hernández de Elche.

Lo que firmo para los efectos oportunos.

En Elche, a ____ de julio de 2024.

Fdo. Dr. D. Óscar Reinoso García
Coordinador del Programa de Doctorado en
Tecnologías Industriales y de Telecomunicación (TECNIT)

Funding Agreements

The work conducted in this thesis has been supported by the following funding agreements.

- Ministerio de Universidades (grant FPU18/00691).
- MCIN/AEI/10.13039/501100011033 (PID2020-115576RB-I00) and “European Union NextGenerationEU/PRTR” (TED2021-130436B-I00).
- Ministerio de Ciencia e Innovación (MCI), AEI, and FEDER funds through the Project under Grant TEC2017-88612-R.
- Generalitat Valenciana (CIGE/2021/096 and CIGE/2022/17).
- UMH’s Vicerrectorado de Investigación.
- Spanish Ministry of Science and Innovation and Universities, AEI and FEDER funds under Project EQC2018-004288-P.



Agradecimientos

A lo largo de este viaje académico y personal, he contado con el apoyo y guía de personas a quienes quiero expresar mi más sincero agradecimiento.

En primer lugar, gracias a mis directores de tesis, Javier y Baldo. Ha sido un privilegio contar con su guía en estos años, que ha sido clave para mejorar la calidad de la investigación llevada a cabo durante esta tesis. Gracias también por todos los consejos y comentarios de mejora que me han permitido crecer tanto profesional como personalmente.

A mis compañeros de uwicore, ha sido un lujo trabajar en este grupo de investigación. He aprendido mucho de ellos durante todos estos años y cada uno de ellos ha contribuido a crear un ambiente de trabajo enriquecedor y colaborativo. Gracias también por el apoyo y los momentos de desconexión fuera del trabajo, que también han sido muy importantes en esta etapa.

A Bea, mi mujer, gracias por ser la mejor compañera de vida. Gracias por tu apoyo constante, tanto en los buenos momentos como en los más complicados, que ha sido esencial para poder llegar a este momento.

A mi familia, gracias por todo el apoyo que me ha ayudado a llegar hasta aquí. Incluyo a la familia de Bea, que también se ha convertido en mi familia. Un agradecimiento especial a mis padres por haberme educado con unos valores y un ejemplo que ha sido clave para alcanzar este objetivo. Un agradecimiento especial también para mis abuelos a los que siento cerca aun estando lejos, incluyendo un recuerdo especial para mi abuelo Juan.

A mis amigos, gracias por todos los ratos de diversión y desconexión necesarios que me han brindado en estos años. Su amistad ha sido muy importante para mantener un equilibrio en esta etapa de mi vida.

Finalmente, quiero agradecer a todos los profesores que he tenido y a los centros educativos en los que he estado. Cada uno de ellos ha contribuido a mi educación y me ha preparado para llegar a este momento.

Resumen

Los vehículos autónomos dependen de sensores integrados para conducir con poca o ninguna intervención humana. En los últimos años, el rendimiento de los sensores en los vehículos autónomos ha mejorado significativamente en términos de alcance de percepción y precisión de detección. Sin embargo, la efectividad de los sensores aún puede verse comprometida por condiciones climáticas adversas, obstáculos u otros factores, lo que puede afectar negativamente a la seguridad y eficiencia de los vehículos autónomos. Los vehículos conectados y autónomos pueden mitigar estos problemas a través de comunicaciones V2X (*Vehicle-to-Everything*), que permiten a los vehículos compartir datos entre sí. Gracias a las comunicaciones V2X, los vehículos conectados y autónomos pueden implementar servicios V2X avanzados como la percepción cooperativa y la conducción cooperativa para mejorar tanto la seguridad como la eficiencia del tráfico. La percepción cooperativa implica el intercambio de datos percibidos por los sensores sobre el entorno de conducción a través de las comunicaciones V2X. Esto permite a los vehículos complementar los datos obtenidos con sus sensores a bordo con los datos obtenidos por los sensores de los vehículos que tienen alrededor e intercambiados mediante V2X. La conducción cooperativa permite que los vehículos conectados y autónomos compartan sus intenciones de conducción y coordinen sus maniobras, lo cual es clave para realizar las maniobras de una manera más segura y eficiente.

El *3rd Generation Partnership Project (3GPP)* publicó el estándar 5G NR (*New Radio*) V2X en la *Release 16*. El estándar 3GPP 5G NR V2X (o NR V2X) soporta comunicaciones directas o *sidelink (SL)* de Vehículo a Vehículo (V2V) y de Vehículo a Infraestructura (V2I) y complementa a LTE V2X (es decir, el predecesor de NR V2X) que fue diseñado para servicios V2X básicos de seguridad vial. NR V2X ha sido diseñado para soportar servicios V2X avanzados para la conducción conectada y autónoma (como la percepción y conducción cooperativa) que tienen requisitos estrictos, por ejemplo, en términos de latencia y fiabilidad. Con este objetivo, NR V2X introduce dos modos de operación: modo 1 y modo 2. En el modo 1, la infraestructura celular gestiona y selecciona los recursos de

vii

comunicación para cada comunicación SL, mientras que en el modo 2, los vehículos seleccionan y gestionan los recursos radio de forma autónoma sin el apoyo de la infraestructura celular. En el modo 2, los vehículos también pueden reservar recursos radio de forma autónoma para sus futuras transmisiones de paquetes anunciando la reserva de recursos radio a los otros vehículos.

La fase inicial de esta tesis coincide con la fase final del proceso de estandarización de NR V2X. Cuando se publica un nuevo estándar, es esencial estudiar su funcionalidad y evaluar su rendimiento. La evaluación del rendimiento proporciona los conocimientos necesarios para diseñar mejoras. En este contexto, esta tesis estudia, evalúa y mejora el estándar NR V2X, centrándose en la evaluación a nivel de sistema de las comunicaciones NR V2X modo 2. Esta tesis se centra en NR V2X modo 2 ya que asegura que la provisión de servicios V2X no esté limitada por la disponibilidad de cobertura celular.

En la fase inicial de esta tesis no existía en la comunidad investigadora un simulador de código abierto de NR V2X modo 2. Esta tesis ha implementado un simulador de NR V2X modo 2 que cumple con los estándares 3GPP y sigue las directrices de evaluación del 3GPP. Este simulador ha sido utilizado en esta tesis para evaluar el estándar NR V2X modo 2 y las diferentes propuestas.

Esta tesis primero se enfoca en analizar la eficiencia de NR V2X modo 2 para soportar servicios V2X avanzados. Estos servicios generan tráfico V2X aperiódico de tamaño variable según el 3GPP. Estudios anteriores demostraron que los patrones de tráfico variable impactan significativamente en la operación y el rendimiento de LTE V2X modo 4 (es decir, el análogo de NR V2X modo 2) debido a ciertas ineficiencias de la MAC (Medium Access Control). NR V2X modo 2 introduce nuevas características MAC para soportar servicios V2X avanzados para la conducción conectada y autónoma. Una de estas características es el mecanismo de re-evaluación o *re-evaluation mechanism* diseñado para detectar y evitar colisiones de paquetes. El mecanismo de re-evaluación es un mecanismo obligatorio que chequea antes de cada transmisión si los recursos seleccionados aún están disponibles. Si otro vehículo ha reservado los mismos recursos, el mecanismo de re-evaluación selecciona nuevos recursos. Estudios previos habían evaluado el rendimiento de NR V2X modo 2 con diferentes patrones de tráfico. Sin embargo, estos estudios no implementaban el mecanismo

de re-evaluación. Además, estos estudios solo consideraban tráfico V2X periódico o aperiódico de tamaño fijo. Sin embargo, las directrices de la metodología de evaluación del 3GPP para NR V2X recomiendan modelos de generación de tráfico para servicios V2X avanzados que también incluyen tráfico de tamaño variable. El estudio realizado en esta tesis representa entonces la primera evaluación existente en la comunidad investigadora de una implementación completamente conforme al estándar NR V2X modo 2 con los patrones de tráfico recomendados por el 3GPP para servicios V2X avanzados. Este estudio muestra que NR V2X modo 2 aún experimenta desafíos a nivel MAC para soportar eficientemente tráfico aperiódico de tamaño variable.

Esta tesis también realiza un análisis y evaluación en profundidad del impacto del mecanismo de re-evaluación en la operación y el rendimiento de NR V2X modo 2. Este estudio es el primero que analiza cuándo y por qué el mecanismo de re-evaluación es efectivo o no en la detección y prevención de colisiones de paquetes. El estudio llevado a cabo en esta tesis considera vehículos transmitiendo paquetes periódicos o aperiódicos de tamaño fijo o variable, siguiendo las directrices del 3GPP. El estudio muestra que el mecanismo de re-evaluación evita de forma efectiva las colisiones de paquetes con tráfico periódico de tamaño fijo, pero su impacto es pequeño ya que el número de colisiones de paquetes detectadas por el mecanismo de re-evaluación es bajo con este tráfico. Con tráfico aperiódico de tamaño variable, la efectividad del mecanismo de re-evaluación disminuye. NR V2X modo 2 permite que los vehículos realicen retransmisiones (es decir, más de una transmisión por paquete). La efectividad del mecanismo de re-evaluación para detectar y evitar colisiones mejora cuando se consideran las retransmisiones. Sin embargo, el impacto del mecanismo de re-evaluación en el rendimiento de NR V2X modo 2 es pequeño ya que el beneficio de las retransmisiones prevalece sobre las ganancias obtenidas con las colisiones de paquetes evitadas con el mecanismo de re-evaluación.

El mecanismo de re-evaluación chequea antes de cada transmisión si los recursos seleccionados aún están disponibles (es decir, si otro vehículo no los ha reservado). Si otro vehículo ha reservado los mismos recursos, el mecanismo de re-evaluación selecciona nuevos recursos. El estándar 3GPP permite diferentes estrategias para determinar cuándo y con qué frecuencia realizar estos chequeos asociados con el mecanismo de re-evaluación.

Esta tesis también evalúa el impacto de diferentes estrategias de chequeo del mecanismo de re-evaluación en el rendimiento de NR V2X modo 2 con tráfico aperiódico de tamaño variable. En particular, esta tesis evalúa dos estrategias diferentes propuestas por el estándar 3GPP para NR V2X modo 2 (que esta tesis denomina estrategias *one-slot* y *all-slots*). La evaluación llevada a cabo en esta tesis muestra que las dos estrategias del estándar presentan un equilibrio entre la latencia de transmisión y el coste computacional del mecanismo de re-evaluación. Esta tesis propone entonces una estrategia alternativa (estrategia *first-slot*) que reduce la latencia de transmisión y el coste computacional sin degradar la fiabilidad de las transmisiones de NR V2X modo 2.

Esta tesis muestra que el mecanismo de re-evaluación no es completamente efectivo para evitar las colisiones de paquetes generadas por tráfico aperiódico de tamaño variable porque muchas de las colisiones detectadas finalmente no ocurren, y seleccionar nuevos recursos con el mecanismo de re-evaluación aumenta el riesgo de colisiones de paquetes. Para abordar esta ineficiencia, esta tesis propone un mecanismo de re-evaluación selectivo (*selective re-evaluation mechanism*) que solo selecciona nuevos recursos cuando el vehículo está seguro de que una colisión detectada va a ocurrir, lo cual es el caso cuando la colisión se detecta con una reserva para una retransmisión de un paquete. El mecanismo de re-evaluación selectivo es por lo tanto aplicable solo cuando NR V2X modo 2 está configurado con retransmisiones. Esta tesis muestra que el mecanismo de re-evaluación selectivo propuesto mejora la fiabilidad y la latencia de las comunicaciones NR V2X modo 2.

Esta tesis también propone la técnica *V2X Rebroadcasting*, que puede aplicarse cuando NR V2X modo 2 está configurado para realizar una transmisión por paquete (es decir, sin retransmisiones). *V2X Rebroadcasting* mejora la eficiencia de la MAC de NR V2X para cualquier mensaje y patrón de tráfico V2X. Una de las ineficiencias predominantes que NR V2X modo 2 experimenta con el tráfico aperiódico de tamaño variable es causada por reservas no utilizadas. Las reservas no utilizadas ocurren cuando un vehículo reserva un recurso radio, pero finalmente no lo usa porque no tiene un paquete para transmitir en el momento de la reserva. El vehículo no puede informar a otros vehículos sobre los recursos que reserva para la siguiente transmisión, y esto aumenta el riesgo de colisiones de paquetes. Para reducir este riesgo, *V2X Rebroadcasting* permite a los vehículos retransmitir o

rebroadcastear paquetes de otros vehículos en reservas no utilizadas. Al hacerlo, el vehículo puede utilizar el paquete *rebroadcasteado* para reservar recursos para su próxima transmisión, y por lo tanto reduce el riesgo de colisiones de paquetes. Además, la propuesta aumenta la fiabilidad de los paquetes *rebroadcasteados*. Los resultados demuestran que *V2X Rebroadcasting* reduce la probabilidad de colisiones de paquetes y mejora la fiabilidad de las comunicaciones V2X en comparación con el estándar NR V2X modo 2.



Abstract

Automated vehicles rely on embedded sensors to drive with low or no human intervention. In recent years, the performance of sensors in automated vehicles has greatly improved in terms of perception range and detection accuracy. However, the effectiveness of sensors can still be compromised by adverse weather conditions, obstacles, or other factors, which can negatively impact the safety and efficiency of automated vehicles. Connected and Automated Vehicles (CAVs) can mitigate these issues through V2X (Vehicle-to-Everything) communications, which allow vehicles to share data with each other. Thanks to V2X communications, CAVs can implement advanced V2X services such as cooperative perception and cooperative driving to enhance both traffic safety and efficiency. Cooperative perception involves the exchange of data perceived by the sensors about the driving environment through V2X communications. This allows vehicles to complement the data obtained with their on-board sensors with data obtained by the sensors of surrounding vehicles. Cooperative driving enables CAVs to share their driving intentions and coordinate their manoeuvres which is key to achieve a smoother and safer traffic flow.

The 3rd Generation Partnership Project (3GPP) published the 5G NR (New Radio) V2X standard in Release 16. The 3GPP 5G NR V2X (or NR V2X) standard supports direct or sidelink (SL) Vehicle-to-Vehicle (V2V) and Vehicle-to-Infrastructure (V2I) communications and complements LTE V2X (i.e. the NR V2X predecessor) that was designed for basic awareness services. NR V2X has been designed to support advanced V2X services for connected and automated driving (such as cooperative perception and driving) which have stringent requirements, e.g. in terms of latency and reliability. To this aim, NR V2X introduces two 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. In mode 2, vehicles can also autonomously reserve radio resources for their future packet transmissions by announcing the reservation of radio resources to the other vehicles.

The initial phase of this thesis coincides with the final phase of the NR V2X standardization process. When a new standard is published, it is essential to study its functionality and evaluate its performance. Performance evaluation provides the necessary insights to design improvements. In this context, this thesis studies, evaluates, and improves the NR V2X standard, with a focus on the system-level evaluation of NR V2X mode 2 communications. This thesis focuses on NR V2X mode 2 as it ensures that V2X service provisioning is not limited by the availability of cellular coverage.

At the initial phase of this thesis, there were no open-source NR V2X mode 2 simulators available in the community. This thesis has included the implementation of an NR V2X mode 2 simulator that adheres to the 3GPP standards and follows the 3GPP evaluation guidelines. This simulator has been utilized in this thesis to evaluate the NR V2X mode 2 standard and the different proposals.

This thesis first focuses on analyzing the efficiency of NR V2X mode 2 to support advanced V2X services. These services are expected to generate aperiodic V2X traffic of variable size according to 3GPP. Variable traffic patterns were shown to significantly impact the operation and performance of LTE V2X mode 4 (i.e. the counterpart of NR V2X mode 2) due to certain MAC (Medium Access Control) inefficiencies. NR V2X mode 2 introduces novel MAC features to support advanced V2X services for connected and automated driving. One of these features is the re-evaluation mechanism designed to detect and avoid packet collisions. The re-evaluation mechanism is a mandatory MAC feature that checks before every transmission if the selected resources are still available. If another vehicle has reserved the same resources, the re-evaluation mechanism selects new resources. Previous studies had evaluated the performance of NR V2X mode 2 under different traffic patterns. However, these studies did not implement the re-evaluation mechanism. In addition, these studies only considered periodic or aperiodic V2X traffic of fixed size. However, the 3GPP evaluation methodology guidelines for NR V2X recommend traffic generation models for advanced V2X services that also include traffic of variable size. The study conducted in this thesis is then the first evaluation of a fully standard compliant implementation of NR V2X mode 2 under the traffic patterns recommended by 3GPP for advanced V2X services. This

study shows that NR V2X mode 2 still faces MAC challenges to efficiently support aperiodic traffic of variable size.

This thesis also conducts an in-depth analysis and evaluation of the impact of the re-evaluation mechanism on the operation and performance of NR V2X mode 2. This study is the first that analyzes when and why re-evaluation is effective or not in detecting and avoiding packet collisions. The analysis considers vehicles transmitting periodic or aperiodic packets of fixed or variable size, following 3GPP guidelines. The study shows re-evaluation effectively avoids packet collisions with periodic traffic of fixed size, but its impact is small since the number of packet collisions detected by re-evaluation is low with this traffic. With aperiodic traffic of variable size, the effectiveness of re-evaluation decreases. NR V2X mode 2 allows vehicles to perform retransmissions (i.e. more than one transmission per packet). The effectiveness of re-evaluation to detect and avoid collisions improves when retransmissions are considered. However, the impact of re-evaluation on the performance of NR V2X mode 2 is small since the benefit of retransmissions prevails over the gains obtained with the packet collisions avoided with re-evaluation.

The re-evaluation mechanism checks before every transmission whether the selected resources are still available (i.e., they have not been reserved by another vehicle). If another vehicle has reserved the same resources, the re-evaluation mechanism selects new resources. The 3GPP standard allows different strategies for when and how often to perform these checks associated with the re-evaluation mechanism. This thesis also evaluates the impact of different re-evaluation check strategies on the performance of NR V2X mode 2 under aperiodic traffic of variable size. In particular, it evaluates two different strategies proposed by the 3GPP standard for NR V2X mode 2 (that this thesis names one-slot and all-slots strategies). The evaluation reveals that the two standardized strategies balance transmission latency and the computational cost of re-evaluation. This thesis then proposes an alternative strategy (first-slot strategy) that reduces transmission latency and computational cost without degrading the reliability of NR V2X mode 2 transmissions.

This thesis shows that re-evaluation is not fully effective in avoiding packet collisions generated by aperiodic traffic of variable size because many of the detected collisions ultimately do not happen, and selecting new resources with the re-evaluation mechanism

increases the risk of packet collisions. To address this inefficiency, this thesis proposes a selective re-evaluation mechanism that only selects new resources when the vehicle is certain that a detected collision is going to occur. This is only the case when the collision is detected with a reservation for a retransmission of a packet. The selective re-evaluation mechanism is therefore applicable only when NR V2X mode 2 is configured with retransmissions. This thesis shows that the proposed selective re-evaluation mechanism improves the reliability and latency of 5G NR V2X mode 2 communications.

This thesis also proposes the *V2X Rebroadcasting* scheme, which can be applied when NR V2X mode 2 is configured to perform one transmission per packet (i.e. without retransmissions), while still improving the efficiency of the NR V2X MAC for any V2X message and traffic patterns. One of the predominant inefficiencies that NR V2X mode 2 experiences under aperiodic traffic of variable size is caused by unutilized reservations. Unutilized reservations occur when a vehicle reserves a radio resource, but does not finally use it because it has no packet to transmit at the time of the reservation. The vehicle cannot inform other vehicles about the resources it reserves for the following transmission, and this increases the risk of packet collisions. To reduce this risk, *V2X Rebroadcasting* allows vehicles to rebroadcast packets from other vehicles in detected unutilized reservations. By doing so, the vehicle can utilize the rebroadcasted packet to reserve resources for its next transmission, and hence reduces the risk of packet collisions. In addition, the proposal increases the reliability of rebroadcasted packets. The results demonstrate that *V2X Rebroadcasting* reduces the probability of packet collisions and improves the reliability of V2X communications compared to the standard NR V2X mode 2.

Contents

Resumen	vii	
Abstract	xii	
Contents	xvi	
List of Figures	xix	
List of Tables	xxi	
1	Introduction	1
1.1	Objectives	2
1.2	Thesis structure and outline.....	4
1.3	Results of the thesis	8
2	5G NR V2X Communications	10
2.1	Introduction.....	10
2.2	Resource grid	10
2.3	Resource allocation	11
2.4	Re-evaluation mechanism	14
2.5	Additional details.....	16
3	Simulation Environment	18
3.1	Simulation platform.....	18
3.1.1	ns-3 simulator.....	18
3.1.2	Implementation of 5G NR V2X in ns-3.....	20
3.2	Scenarios and parameters	30
3.3	Evaluation metrics	31
4	5G NR V2X MAC Challenges	33
4.1	Introduction.....	33
4.2	Impact of packet variability on 5G NR V2X.....	34
4.2.1	Size reselections	34
4.2.2	Latency reselections.....	35
4.2.3	Unutilized reservations	35
4.2.4	Unused sub-channels	35
4.3	Evaluation	36
4.4	Conclusions and additional results.....	40

5	Impact of Re-Evaluation on 5G NR V2X	42
5.1	Introduction.....	42
5.2	Analysis of the Re-evaluation Mechanism	43
5.2.1	Re-evaluation check.....	43
5.2.2	Re-evaluation detection	44
5.2.3	Resource replacement	45
5.3	Impact of Re-evaluation without retransmissions	46
5.4	Impact of Re-evaluation with retransmissions.....	49
5.5	Re-evaluation strategies.....	50
5.5.1	One-slot strategy	51
5.5.2	All-slots strategy	51
5.5.3	First-slot strategy.....	52
5.6	Impact of re-evaluation strategies	53
5.7	Conclusions and additional results.....	55
6	Selective Re-evaluation Mechanism	58
6.1	Introduction.....	58
6.2	Effective and ineffective re-evaluations.....	58
6.3	Effectiveness of re-evaluation	60
6.3.1	One transmission per TB	61
6.3.2	Two transmissions per TB	65
6.3.3	Numerical evaluation	67
6.4	Selective re-evaluation mechanism.....	69
6.5	Evaluation	70
6.6	Conclusions.....	76
7	V2X Rebroadcasting.....	77
7.1	Introduction.....	77
7.2	Impact of RRI selection strategy on MAC challenges	79
7.3	V2X Rebroadcasting.....	79
7.3.1	Random	82
7.3.2	Minimum Age	82
7.3.3	Low Age and Random (LAR)	82
7.3.4	Low Age, Minimum Rebroadcasts and Random (LAMR)	83
7.3.5	Low Age, Minimum Rebroadcasts and Power (LAMRP)	84
7.4	Scenario and Metrics	84

7.5	Evaluation of V2X Rebroadcasting	85
7.5.1	Configuration and selection of Low Age algorithms	85
7.5.2	Evaluation of V2X Rebroadcasting algorithms	88
7.6	Comparison of V2X Rebroadcasting and the standard 5G NR V2X mode 2	89
7.7	Conclusions.....	93
8	Conclusions and future work	95
9	Conclusiones y trabajo futuro.....	101
	Bibliography	107
	Annex A.1 Publication.....	110
	Annex A.2 Publication.....	166
	Annex A.3 Publication.....	173
	Annex A.4 Publication.....	189



List of Figures

Figure 1. NR V2X channelization and illustration of resource allocation in NR V2X mode 2 (when $T_2 = PDB$) [17].	11
Figure 2. Operation of re-evaluation under NR V2X mode 2 [16].	15
Figure 3. Architecture of the NR V2X mode 2 implementation in ns-3 [own elaboration].	22
Figure 4. NR-V2X-Mac layer [own elaboration].	25
Figure 5. NR-V2X-Phy layer [own elaboration].	27
Figure 6. NR-V2X-Channel entity [own elaboration].	28
Figure 7. Impact of packet variability on SPS under NR V2X mode 2 [17].	35
Figure 8. Packet Delivery Ratio (PDR) [16].	37
Figure 9. Packet Collision Ratio (PCR) under medium (left) & high (right) traffic intensity; 100 veh/km – similar trends observed for other vehicle densities [16].	39
Figure 10. SPS performance in mixed traffic scenario with aperiodic traffic of variable size, $N = 1$ [17].	47
Figure 11. SPS performance in mixed traffic scenario for periodic traffic, 100 veh/km, $N = 1$. Similar trends are observed under the single traffic scenario [17].	48
Figure 12. 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 [17].	50
Figure 13. PDR experienced by SPS for aperiodic traffic of variable size when $N = 2$, (50 veh/km average RRI strategy) in the single traffic scenario. Similar trends have been obtained for other densities [17].	50
Figure 14. Illustration of re-evaluation check strategies (when T_3 and T_1 are equal to 5 and 2 slots, respectively) [18].	51
Figure 15. CDF of the end-to-end latency for 50 veh/km (similar trends observed for 100 veh/km) [18].	53
Figure 16. PDR of packets that experienced at least a re-evaluation detection (a) and PDR experienced by all packets (b). Medium intensity with 50 veh/km (same trends observed in the other evaluated scenarios) [18].	55

Figure 17. Probability of not utilizing a reservation for a new TB (<i>PNUT</i>) for one transmission per TB [19].	68
Figure 18. Probability of not utilizing a reservation for a new TB (<i>PNUTR</i>) for two transmissions per TB [19].	68
Figure 19. Reliability of the standard and selective re-evaluation mechanisms for aperiodic traffic of variable size in low, medium and high intensity scenarios with 50 and 100 veh/km [19].	71
Figure 20. Ratio of Not Utilized reservations (<i>NUT</i> , Section 6.3.2). This is the ratio of reservations for a new TB that are not utilized to the total number of reservations for new TBs for 50 and 100 veh/km. Average inter-packet arrival times of 20, 100 and 200 ms correspond to the high, medium and low intensity scenarios respectively [19].	72
Figure 21. Reliability of the standard and selective re-evaluation mechanisms for periodic traffic of variable size in the medium intensity scenario with 50 and 100 veh/km. Similar trends are observed in the other intensity scenarios [19].	74
Figure 22. Impact of the re-evaluation mechanisms on the latency: SRO for 100 veh/km (a) and CDF of the latency for low intensity (b), medium intensity (c) and high intensity (d) scenarios with 50 and 100 veh/km [19].	75
Figure 23. MAC operation of V2X Rebroadcasting [20].	81
Figure 24. Illustration of V2X Rebroadcasting [20].	81
Figure 25. PDR of LAMR for different values of Threshold Age and Δ_{Age} . Medium intensity traffic [20].	86
Figure 26. PDR of Low Age algorithms for the medium intensity traffic scenario. PDB = 50ms, Threshold Age = PDB, Δ_{Age} = 20ms [20].	87
Figure 27. PDR and PLR of Random, Minimum Age and LAMR algorithms under the medium intensity traffic scenario and a density of 50 veh/km. Similar trends were observed for high traffic intensity and 25 veh/km [20].	89
Figure 28. PDR achieved by the standard NR V2X mode 2 and V2X Rebroadcasting for (a) medium intensity and (b) high intensity traffic scenario [20].	91
Figure 29. PCR and PLR of V2X Rebroadcasting and standard NR V2X mode 2 for (a) medium intensity and (b) high intensity traffic scenarios with 25 veh/km. Similar trends were observed for a density of 50 veh/km [20].	91
Figure 30. PIR of standard NR V2X mode 2 and V2X Rebroadcasting in the (a) medium and (b) high intensity traffic scenarios [20].	93

List of Tables

Table 1. Probability of LOS and NLOSv link states	29
Table 2. Pathloss model and NLOSv losses depending on the link state.....	29
Table 3. NR V2X mode 2 operation for the medium and high traffic intensity scenarios	38
Table 4. Performance metrics (in %) when $N=1$	47
Table 5. Performance metrics (in %) for aperiodic traffic of variable size when $N=2$ (average RRI strategy)	50
Table 6. Operation and Performance of the Re-evaluation Check Strategies.....	54
Table 7. Probability that TBs are rebroadcasted once ($P(NREB=1)$) or more than once ($P(NREB >1)$) for LAMR. Medium intensity traffic and 50 veh/km. Similar trends were observed for 25 veh/km.	86
Table 8. Impact of Threshold Age on the PDBR for LAMR	86
Table 9. Probability that TBs are rebroadcasted once ($P(NREB=1)$) or more than once ($P(NREB >1)$) for Low Age algorithms. Medium intensity traffic and 50 veh/km. Similar trends were observed for high intensity traffic and 25 veh/km.	87
Table 10. Probability that TBs are rebroadcasted once ($P(NREB=1)$) or more than once ($P(NREB >1)$) for Random, Minimum Age and LAMR. Medium intensity traffic and 50 veh/km. Similar trends were observed for high intensity traffic and 25 veh/km.	89
Table 11. PDBR for LAMR (Threshold Age = PDB), Random and Minimum Age algorithms.....	89

1 Introduction

Automated vehicles rely on embedded sensors to drive with low or no human intervention. The vehicle's planning system utilizes perception and localization data to determine its travel path and driving actions, such as lane changes, acceleration, and braking. These actions are executed by the vehicle's control system. For accurate perception and localization, automated vehicles are equipped with multiple sensors such as lidars, radars, and cameras that allow them to perceive their driving environment locally [1].

In recent years, the performance of sensors in automated vehicles has greatly improved in terms of perception range and detection accuracy [2]. However, the effectiveness of sensors can still be compromised by adverse weather conditions, obstacles, or other factors, which can negatively impact the safety and efficiency of automated vehicles [3]. Connected and Automated Vehicles (CAVs) can mitigate these issues through V2X (Vehicle-to-Everything) communications, which allow vehicles to share data with each other. V2X communications allow vehicles to receive additional information about the driving environment, including data that allow them to extend their perception beyond the field of view (FoV) of their on-board sensors.

Thanks to V2X communications, CAVs can implement advanced V2X services such as cooperative perception [4] and cooperative driving [5] to enhance both traffic safety and efficiency. Cooperative perception involves the exchange of data perceived by the sensors about the driving environment through V2X communications. This allows vehicles to complement the data obtained with their on-board sensors with data obtained by the sensors of surrounding vehicles. Cooperative driving enables CAVs to share their driving intentions and coordinate their manoeuvres which is key to achieve a smoother and safer traffic flow.

The 3rd Generation Partnership Project (3GPP) published the 5G NR (New Radio) V2X standard in Release 16 [6][7]. The 3GPP 5G NR V2X (or NR V2X) standard supports direct or sidelink (SL) Vehicle-to-Vehicle (V2V) and Vehicle-to-Infrastructure (V2I) communications and complements LTE V2X [8][9] (i.e. the NR V2X predecessor) that was designed for basic awareness services [10]. NR V2X has been designed to support advanced V2X services for connected and automated driving (such as cooperative perception and driving) which have stringent requirements, e.g. in terms of latency and reliability [11]. To this aim, NR V2X introduces two 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. In mode 2, vehicles can also autonomously reserve radio resources for their future packet transmissions by announcing the reservation of radio resources to the other vehicles. This thesis focuses on NR V2X mode 2 as it ensures that V2X service provisioning is not limited by the availability of cellular coverage.

The initial phase of this thesis coincides with the final phase of the NR V2X standardization process. When a new standard is published, it is essential to study its functionality and evaluate its performance. Performance evaluation provides the necessary insights to design improvements. In this context, the general purpose of this thesis is to study, evaluate, and improve the NR V2X standard, with a focus on the system-level evaluation of NR V2X mode 2 communications. To achieve this purpose, this thesis has accomplished the objectives detailed in the following section.

1.1 Objectives

Objective 1: The first objective of this thesis is to achieve a deep understanding of the 5G NR V2X standard and the related literature. During the standardization process of a 3GPP standard, various institutions and organizations' delegates that make up the 3GPP meet regularly to make proposals and reach agreements. Therefore, it is essential to periodically track the meetings and study the agreements reached by the 3GPP during the NR V2X standardization process. Once the standard documents are published, it is necessary to extend the study to these 3GPP standard documents as well as to the related literature.

Objective 2: At the initial phase of this thesis, there was no open-source NR V2X mode 2 simulator available in the community to evaluate the standard and potential improvement proposals. Therefore, the second objective of this thesis is to develop an NR V2X mode 2 simulator that adheres to the 3GPP standards [6][7] and follows the 3GPP evaluation guidelines [12].

Objective 3: The third objective of this thesis is analyzing the efficiency of NR V2X mode 2 to support advanced V2X services. These services are expected to generate V2X aperiodic traffic of variable size according to 3GPP [12]. Variable traffic patterns were shown to significantly impact the operation and performance of LTE V2X mode 4 (i.e. the counterpart of NR V2X mode 2) [13]. This was due to certain MAC (Medium Access Control) inefficiencies when vehicles generate aperiodic messages of variable size. Therefore, this third objective involves verifying whether NR V2X mode 2 exhibits similar inefficiencies to LTE V2X mode 4 when vehicles generate the aperiodic traffic of variable size characteristic of advanced V2X services.

Objective 4: One of the main novelties of NR V2X mode 2 compared to LTE V2X mode 4 is the introduction of the re-evaluation mechanism. The re-evaluation mechanism is a new mandatory feature designed to detect and avoid packet collisions. The fourth objective is to conduct an in-depth analysis and evaluation of the impact of the re-evaluation mechanism on the operation and performance of NR V2X mode 2. The goal of this analysis is to determine when and why re-evaluation is effective or not in detecting and avoiding packet collisions.

Objective 5: A vehicle with the re-evaluation mechanism checks before transmitting on a resource if another vehicle has announced that it reserves the same resource. If this is the case, the re-evaluation mechanism detects a potential collision and changes the transmission resource to avoid it. The 3GPP standard allows different strategies for when and how often to perform these checks associated with the re-evaluation mechanism. The fifth objective of the thesis is to evaluate the impact of these different strategies on the performance of NR V2X mode 2 and potentially propose an alternative strategy that outperforms the standardized strategies.

Objective 6: The sixth objective of the thesis is to design an improved re-evaluation mechanism to enhance the performance of NR V2X mode 2. This new design will be based on the knowledge acquired from objective 4 about when and why re-evaluation is effective

or not in detecting and avoiding packet collisions. NR V2X mode 2 can be configured to perform one or more than one transmission per packet. This new design of the re-evaluation mechanism will be applicable to NR V2X mode 2 when it is configured to perform more than one transmission per packet, following the configuration recommended by the Society of Automotive Engineers (SAE) [14].

Objective 7: The seventh objective of the thesis is to design an improvement for the NR V2X mode 2 standard applicable when the configuration of NR V2X mode 2 is set so that vehicles perform only one transmission per packet. This solution will complement the new re-evaluation design from objective 6, which will be only applicable when vehicles perform more than one transmission per packet.

1.2 Thesis structure and outline

The main body of this thesis consists of a compendium of four publications included in the annexes (from Annex A.1 to A.4). These publications are also summarized in the chapters 2, 4 and 5 of the thesis. Additionally, other chapters are included in this thesis to describe the simulation environment and the design and evaluation of two improvement proposals for NR V2X mode 2. The remaining chapters of the thesis are organized as follows:

Chapter 2 includes an overview of the NR V2X standard. NR V2X has two modes of operation: mode 1 and mode 2. Chapter 2 focuses on NR V2X mode 2 since it is the focus of this thesis as it ensures that V2X service provisioning is not limited by the availability of cellular coverage. Chapter 2 is a summary of this thesis contribution published in [15] (found in Annex A.1) focused on NR V2X mode 2. This publication is an in-depth and comprehensive tutorial on the NR V2X standard.

Chapter 3 describes the simulation environment used to evaluate the NR V2X mode 2 standard and the improvement proposals presented in this thesis. The simulation environment includes a detailed description of the standard-compliant simulator implemented and used in this thesis. Chapter 3 also details the evaluated driving scenarios, the NR V2X parameter configurations, and the metrics used in the conducted evaluations, which are identified following the 3GPP evaluation guidelines in [12].

Chapter 4 focuses on analyzing the efficiency of NR V2X mode 2 to support advanced V2X services. These services are expected to generate aperiodic V2X traffic of variable

size according to 3GPP. Variable traffic patterns were shown to significantly impact the operation and performance of LTE V2X mode 4 (i.e. the counterpart of NR V2X mode 2) due to certain MAC inefficiencies. NR V2X mode 2 introduces novel MAC features to support advanced V2X services for connected and automated driving. One of these features is the re-evaluation mechanism designed to detect and avoid packet collisions. The re-evaluation mechanism is a mandatory MAC feature that checks before every transmission if the selected resources are still available. If another vehicle has reserved the same resources, the re-evaluation mechanism selects new resources. Previous studies had evaluated the performance of NR V2X mode 2 under different traffic patterns. However, these studies did not implement the re-evaluation mechanism. In addition, these studies only considered periodic or aperiodic V2X traffic of fixed size. However, the 3GPP evaluation methodology guidelines for NR V2X recommend traffic generation models for advanced V2X services that also include traffic of variable size in line with the message patterns characteristic of advanced V2X services such as cooperative perception or cooperative driving. The study conducted in Chapter 4 is then the first evaluation of a fully standard compliant implementation of NR V2X mode 2 using the re-evaluation mechanism under periodic and aperiodic traffic of fixed and variable packet size, in accordance with the 3GPP traffic model recommendations. The evaluation is conducted in scenarios where vehicles perform one transmission per packet. The results of this chapter show that NR V2X mode 2 also faces MAC challenges to handle aperiodic packets of variable size. These MAC challenges are described and quantified in this chapter. Chapter 4 shows that NR V2X mode 2 performance degrades with aperiodic traffic of variable size compared to periodic traffic of fixed size due to increased packet collisions. This occurs despite utilizing the re-evaluation mechanism designed to detect and avoid collisions. It should be noted that the evaluation of NR V2X mode 2 included in Chapter 4 is a summary of this thesis contribution published in [16] (found in Annex A.2).

Chapter 4 provides an initial system-level evaluation of NR V2X mode 2, but it does not explain why re-evaluation is not effective in certain scenarios. Understanding this is crucial for optimizing NR V2X mode 2. It is also necessary to determine if the observations made in Chapter 4 hold when retransmissions are used. In this context, Chapter 5 provides an in-depth analysis and evaluation of the impact of the re-evaluation mechanism on the operation and performance of NR V2X mode 2. This study is the first that analyzes when and why re-evaluation is effective or not in detecting and avoiding packet collisions. The

analysis considers vehicles transmitting periodic or aperiodic packets of fixed or variable size, following 3GPP guidelines [12]. The study shows re-evaluation effectively avoids packet collisions with periodic traffic of fixed size, but its impact is small since the number of packet collisions detected by re-evaluation is low with this traffic. With aperiodic traffic of variable size, the effectiveness of re-evaluation decreases. The capacity of re-evaluation to detect and avoid collisions improves when retransmissions are considered. However, the impact of re-evaluation on the performance of NR V2X mode 2 is small since the benefit of retransmissions prevails over the gains obtained with the packet collisions avoided with re-evaluation. A vehicle with the re-evaluation mechanism checks, before transmitting on a resource, if another vehicle has announced that it reserves the same resource. If so, the re-evaluation mechanism detects a packet collision and selects a new resource to avoid the collision. The 3GPP standard allows different strategies for when and how often to perform these checks associated to the re-evaluation mechanism. Chapter 5 also includes a study to evaluate the impact of different re-evaluation check strategies on the performance of NR V2X mode 2 under aperiodic traffic of variable size. In particular, it evaluates two different strategies proposed by the 3GPP standard for NR V2X mode 2 (that this thesis names one-slot and all-slots strategies). The evaluation reveals that the two standardized strategies balance transmission latency and the computational cost of re-evaluation. The chapter then proposes an alternative strategy (first-slot strategy) that reduces transmission latency and computational cost without degrading the reliability of NR V2X mode 2 transmissions. It should be noted that the part of Chapter 5 that analyzes the impact of re-evaluation and its effectiveness in detecting and avoiding packet collisions in NR V2X mode 2 is a summary of this thesis contribution published in [17] (found in Annex A.3). The part of Chapter 5 that evaluates the impact of different re-evaluation check strategies on NR V2X mode 2 is a summary of this thesis contribution published in [18] (found in Annex A.4).

The analysis in Chapter 5 shows that re-evaluation is not fully effective in avoiding packet collisions generated by aperiodic traffic of variable size because many of the detected collisions ultimately do not happen, and selecting new resources with the re-evaluation mechanism increases the risk of packet collisions. To address this inefficiency, Chapter 6 proposes a selective re-evaluation mechanism that only selects new resources when the vehicle is certain that a detected collision is going to occur, which is the case when the collision is detected with a reservation for a retransmission of a packet. The proposal builds from a thorough analysis and characterization of the challenges faced by the standard re-

evaluation mechanism to avoid packet collisions when dealing with aperiodic packets of variable size. Chapter 6 shows that the proposed selective re-evaluation mechanism improves the reliability and latency of 5G NR V2X mode 2 communications. It should be noted that Chapter 6 is related to this thesis contribution submitted for publication [19].

The selective re-evaluation mechanism proposed in this thesis in Chapter 6 demonstrates improvements in terms of reliability for NR V2X mode 2 communications under aperiodic traffic of variable size when the MAC is configured with retransmissions. However, the study conducted in Chapter 4 shows that NR V2X mode 2 also faces MAC challenges under aperiodic traffic of variable size without retransmissions (i.e. when vehicles only perform one transmission per packet) resulting in increased packet collisions. Chapter 7 then proposes the *V2X Rebroadcasting* scheme, which can be applied in scenarios where NR V2X mode 2 is configured to perform one transmission per packet (i.e. without retransmissions), while still improving the efficiency of the NR V2X MAC for any V2X message and traffic patterns. One of the predominant inefficiencies of NR V2X is caused by unutilized reservations. Unutilized reservations occur when a vehicle reserves a radio resource, but does not finally use it because it has no packet to transmit at the time of the reservation. The vehicle cannot inform other vehicles about the resources it reserves for the following transmission, and this increases the risk of packet collisions. To reduce this risk, *V2X Rebroadcasting* enables vehicles to rebroadcast packets from other vehicles in detected unutilized reservations. By doing so, the vehicle can utilize the rebroadcasted packet to reserve resources for its next transmission, and hence reduces the risk of packet collisions. In addition, the proposal increases the reliability of rebroadcasted packets as it can help resolve potential packet collisions and mitigate propagation losses experienced by the original transmission of the packet. Chapter 7 proposes, evaluates and optimizes different algorithms to select the packets to rebroadcast, and compares the performance achieved with *V2X Rebroadcasting* to the performance experienced with the standard NR V2X mode 2. The results demonstrate that *V2X Rebroadcasting* eliminates unutilized reservations, reduces the probability of packet collisions, improves the reliability of V2X communications, and reduces the packet inter-reception time. The benefits can diminish under very high channel load levels, but *V2X Rebroadcasting* still improves the reliability of NR V2X communications at short to medium distances where safety-critical V2X data is more necessary and relevant. It should be noted that Chapter 7 is related to a journal

paper that we have elaborated in the framework of this thesis and that will be submitted for publication just after the submission of this thesis [20].

Finally, Chapter 8 summarizes the conclusions derived from the work conducted in this thesis and outlines potential future research directions. The content of Chapter 8 is included in Chapter 9 in Spanish.

1.3 Results of the thesis

The research presented in this thesis has been published so far in two journal articles and two conference papers, which are included in the annexes. The references for these publications are as follows:

- 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.
- A. Molina-Galan, B. Coll-Perales, L. Lusvarghi, J. Gozalvez and M. L. Merani, "How does 5G NR V2X Mode 2 Handle Aperiodic Packets and Variable Packet Sizes?," *2022 IEEE 23rd HPSR*, Taicang, Jiangsu, China, 2022, pp. 183-188.
- A. Molina-Galan, L. Lusvarghi, B. Coll-Perales, J. Gozalvez and M. L. Merani, "On the Impact of Re-Evaluation in 5G NR V2X Mode 2," in *IEEE Transactions on Vehicular Technology*, vol. 73, no. 2, pp. 2669-2683, Feb. 2024.
- A. Molina-Galan, B. Coll-Perales and J. Gozalvez, "Re-Evaluation Strategies for 5G NR V2X Communications," *2022 IEEE 96th Vehicular Technology Conference (VTC2022-Fall)*, London, United Kingdom, 2022, pp. 1-5.

In addition to these publications, this thesis has resulted in a journal paper that has been submitted for publication and is currently under review. The research included in this journal paper is presented in Chapter 6 of this thesis. The reference for this journal paper is as follows:

- A. Molina-Galan, J. Gozalvez and B. Coll-Perales, "A Selective Re-Evaluation Mechanism for 5G NR V2X Mode 2 Communications," in *IEEE Transactions on Vehicular Technology* (submitted).

Furthermore, this thesis has resulted in another journal paper, which will be submitted for publication just after the submission of this thesis. The research included in this journal paper is presented in Chapter 7 of this thesis. The reference for this journal paper is as follows:

- A. Molina-Galan, J. Gozálvez and B. Coll-Perales, "V2X Rebroadcasting for C-V2X Communications".



2 5G NR V2X Communications

2.1 Introduction

The 3GPP 5G NR V2X (or NR V2X) standard supports direct or sidelink (SL) Vehicle-to-Vehicle (V2V) and Vehicle-to-Infrastructure (V2I) communications, and complements LTE V2X that was designed for basic awareness V2X services. NR V2X has been designed to support advanced V2X (eV2X) services for connected and automated driving with stringent requirements such as cooperative perception and driving, among others [11]. To this aim, NR V2X defines two 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 [17]. This chapter focuses on NR V2X mode 2 since it is the focus of this thesis as it ensures that V2X service provisioning is not limited by cellular coverage. NR V2X mode 2 inherits the sensing-based scheduling scheme from LTE V2X mode 4, which is its counterpart. However, NR V2X mode 2 introduces new features to support advanced V2X services. This chapter describes NR V2X mode 2, indicating the new features compared to LTE V2X mode 4. In particular, Section 2.2 describes the resource grid of 5G NR V2X. Section 2.3 explains the resource allocation of NR V2X mode 2. Section 2.4 presents the re-evaluation mechanism, one of the main new features introduced in NR V2X mode 2. Finally, Section 2.5 provides additional details about the 5G NR V2X standard included in our related publication [15] (found in Annex A.1).

2.2 Resource grid

NR V2X structures the radio resources into slots in the time domain and Resource Blocks (RBs) in the frequency domain. NR V2X supports a flexible subcarrier spacing (SCS) configured by the OFDM numerology (μ), but all vehicles must utilize the same SCS under NR V2X mode 2 within a given area. μ can be equal to 0, 1, 2 or 3, and the SCS is equal to

$2^{\mu} \times 15$ kHz¹. The SCS determines the slot duration (equal to $2^{-\mu}$ ms) and the RB bandwidth (equal to $12 \times \text{SCS}$ kHz). 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. RBs within the same slot are grouped into sub-channels that represent the smallest unit for SL data transmission or reception (see Figure 1). The number of grouped RBs determine the sub-channel size. Vehicles transmit data packets in Transport Blocks (TBs). Note that the terms TB and packet are interchangeable in this thesis. The number of sub-channels utilized for each TB depends on the packet and sub-channel sizes, and the 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 [17]. In NR V2X, each TB has an associated Sidelink Control Information (SCI) that is transmitted in the same slot as the TB. The SCI provides information on the sub-channels used by the TB, and information necessary for decoding the TB. The SCI can also indicate the reservation of resources for future TBs and retransmissions of the associated TB.

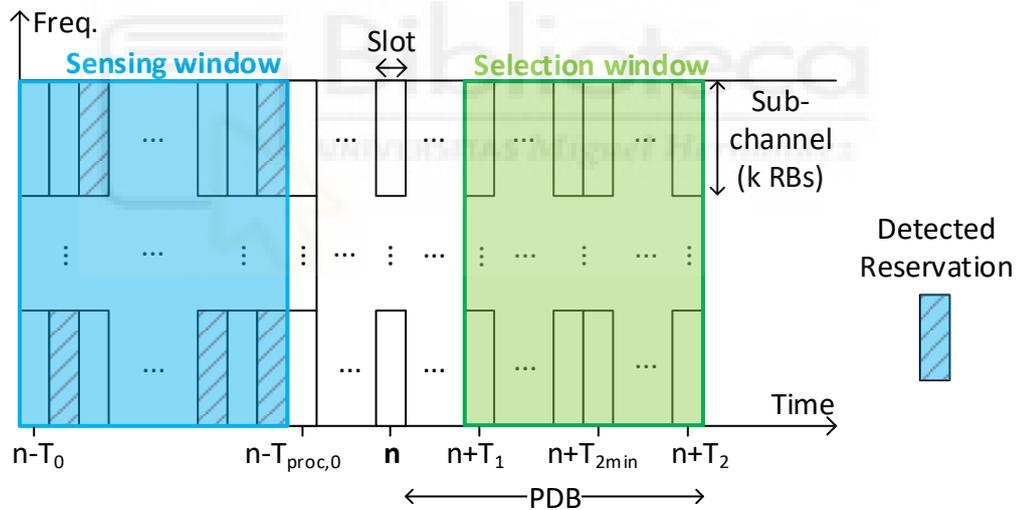


Figure 1. NR V2X channelization and illustration of resource allocation in NR V2X mode 2 (when $T_2 = PDB$) [17].

2.3 Resource allocation

LTE V2X mode 4 operates following a semi-persistent scheduling (SPS) scheme. On the other hand, NR V2X mode 2 can operate using a dynamic scheduling (DS) or an SPS

¹ This is in contrast to LTE V2X that uses a fixed SCS equal to 15 kHz.

scheme that differs from the one designed for LTE V2X mode 4. The DS scheme selects and reserves resources for a single TB, whereas the SPS scheme selects and reserves resources for the transmission and possible retransmissions of *Reselection Counter* consecutive TBs. We should highlight the important difference between a reserved and a selected resource. A reserved resource is a selected resource whose reservation has been announced with the transmission of a previous SCI. 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. Selected resources are hence not reserved. The *Resource Reservation Interval (RRI)* determines the time period between the resources selected for the transmission of consecutive TBs in SPS. The *RRI* can be equal to $\{0, [1:99], 100, 200, 300, 400, 500, 600, 700, 800, 900, 1000\}$ ms. It should be emphasized that NR V2X mode 2 provides higher flexibility to fit the requirements and characteristics of different eV2X services by allowing any integer *RRI* between 1 and 99 ms. This is in contrast to LTE V2X mode 4 that only considered the following *RRI* values: $\{0, 20, 50, 100, 200, 300, 400, 500, 600, 700, 800, 900, 1000\}$ ms [15]. A vehicle can be configured with a list of up to 16 different *RRIs*, but it only selects one *RRI* from the list when it selects new resources. The process to select new resources is referred to as *reselection* by the 3GPP standards. The selected *RRI* also determines the value of the *Reselection Counter*. According to 3GPP standards [6], the vehicle randomly sets the *Reselection Counter* within the interval $[5,15]$ when $RRI \geq 100$ ms, and within the interval $[5*C, 15*C]$, where $C = 100/\max(20, RRI)$, when $RRI < 100$ ms [17].

A vehicle that must select new resources to transmit a TB (for both the DS and SPS schemes) first defines the selection window SW that includes every resource within the range of slots $[n+T_1, n+T_2]$ (Figure 1). n is the slot where the vehicle generates the TB and selects new resources. T_1 is the processing time (in slots) needed by a vehicle to select new resources. The vehicle sets the value of T_2 within $T_{2min} \leq T_2 \leq PDB$, where PDB is the Packet Delay Budget or latency requirement to transmit a TB. Possible values of T_{2min} include $\{1,5,10,20\} * 2^m$ slots, resulting in T_{2min} values equal to $\{1,5,10,20\}$ ms. NR V2X mode 2 can then guarantee minimum latency levels of 1 ms compared to 10 ms in LTE V2X mode 4. After defining the selection window, the vehicle identifies all candidate resources within the selection window. A candidate resource must include K_{SCH} adjacent sub-channels within a slot that fit both the TB and SCI to be transmitted.

To select new resources among the identified candidate resources, DS and SPS schemes define a sensing window within the range of slots $[n-T_0, n-T_{proc,0})$ (Figure 1) [7]. T_0 can be equal to $100 \cdot 2^\mu$ or $1100 \cdot 2^\mu$ slots, while $T_{proc,0}$ can take values from 1 to 4 slots depending on the SCS. During $[n-T_0, n-T_{proc,0})$, vehicles decode the SCIs sent by neighboring vehicles to identify the resources they reserved for their following transmissions. The vehicles also measure the Reference Signal Received Power (RSRP) of the TBs associated with the decoded SCIs [21].

The vehicle that must select new resources to transmit a TB (with DS or SPS) uses then a two-step algorithm to select the resources. In step 1, the vehicle excludes candidate 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^2 [17]. Second, the vehicle also excludes candidate resources from the selection window if: 1) they have already been reserved by other vehicles (through their SCIs), and 2) the measured RSRP exceeds a pre-defined RSRP threshold. Once the execution of step 1 is completed, the vehicle (with DS or SPS) checks the percentage of remaining available candidate resources in the selection window. If this percentage is below a threshold, the RSRP threshold is increased by 3dB and the vehicle repeats step 1. In step 2 (with DS or SPS), the vehicle randomly selects N available candidate resources ($N \leq 32$)³ for the initial transmission of a TB and its $N-1$ retransmissions⁴. If $N > 1$, the vehicle selects the N candidate resources so that all retransmissions of the TB can be reserved by the initial transmission or a previous retransmission of the TB. To do so, the vehicle selects the first candidate resource randomly. For the following ones, the vehicle must consider that an SCI can only reserve a maximum of 2 candidate resources for retransmissions of the TB [6], and the SCI can only reserve resources for retransmissions that take place up to 31 slots later than the SCI [7]. NR V2X mode 2 eliminates the third step included in LTE V2X mode 4 for selecting resources. This third step estimated the average Received Signal

² The resource exclusions due to half-duplex operation have to consider all possible RRI values of the RRI list.

³ This is in contrast to LTE V2X mode 4 that only allows one blind retransmission.

⁴ NR V2X supports blind and Hybrid Automatic Repeat Request (HARQ) feedback-based retransmissions. We consider blind retransmissions when referring to retransmissions in this thesis.

Strength Indicator (RSSI) of the candidate resources to select one from a list of candidate resources with the lowest RSSI.

With SPS, when the vehicle performs the transmission of a TB, it also reserves resources for the transmission of a new TB using the *RRI* included in the SCI. The *RRI* also reserves the resources for the retransmissions of a new TB when the SCI informs about the retransmissions of the 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 $(1-P_k)$ whether it has to select new resources for the transmission of the following TBs; P_k 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 SCI. If the vehicle has to select new resources, it sets to zero the value of the *RRI* in the 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 a new TB. Note that a vehicle using SPS may need to select new resources (i.e. execute a reselection) 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 reserved, or when the previously reserved resources do not meet the latency requirement of the new TB [17].

2.4 Re-evaluation mechanism

NR V2X mode 2 introduces the re-evaluation mechanism to detect and avoid possible packet or TB collisions. With re-evaluation, vehicles verify before transmitting a TB whether the selected resources remain available or not. To do so, the vehicle must again execute step 1 of the resource selection algorithm T_3 slots before the transmission of a TB. This new execution of step 1 is referred to as *re-evaluation check* in 3GPP standards [6]. The vehicle can optionally perform additional re-evaluation checks prior to the one executed T_3 slots before the transmission of a TB. T_3 is the maximum processing time (in slots) allowed for a vehicle to select new resources. T_3 is equal to 3, 5, 9 or 17 slots for a SCS of 15, 30, 60 or 120 kHz, respectively. It is important to note that the re-evaluation mechanism can only be applied over selected (and not reserved) resources. Re-evaluation applies to both DS and SPS.

Figure 2 illustrates the operation of the re-evaluation mechanism. The figure depicts a scenario in which a vehicle originally selects a resource located at slot m , and performs the re-evaluation check at slot n' . The vehicle must then define a new selection window SW' within the range of slots $[n'+T_1, n'+T_2']$. The vehicle configures T_2' within $T_{2min} \leq T_2' \leq PDB - (n'-n)$ so that the upper limit of SW' does not violate the PDB of the TB to be transmitted. The vehicle checks which resources are currently available by executing step 1 over SW' at slot n' . The 3GPP standard specifies that a vehicle performs a *re-evaluation detection* if the originally selected resource at slot m is not available anymore (i.e. another vehicle has reserved it) when the vehicle executes step 1 at slot n' [7]. In this case, the re-evaluation detection triggers the execution of step 2 for selecting new resources (at m' in Figure 2) among the ones available in SW' . 3GPP refers to this new execution of step 2 as *resource replacement* [6].

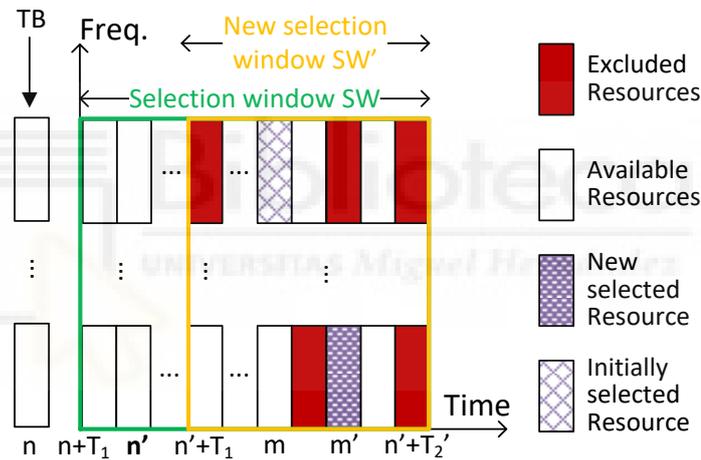


Figure 2. Operation of re-evaluation under NR V2X mode 2 [16].

We should note that a vehicle that has selected $N > 1$ resources for the transmission (and retransmissions) of a TB, conducts the re-evaluation check prior to the initial transmission of the TB. At this time, it verifies the availability of all N selected resources. This is the case because all N resources are eligible for a re-evaluation check before the initial transmission of the TB since they are still selected (and not reserved) resources. This is not the case as soon as resources are reserved for the retransmissions of the TB. If the re-evaluation check identifies a re-evaluation detection for all N resources, the resource replacement selects the first new resource among all available resources within SW' . The resource replacement selects the other new resources so that all retransmissions of the TB can be reserved by the initial transmission or a previous retransmission of the TB. To do so, it considers that the SCI can only reserve resources for retransmissions that take place

up to 31 slots later than the SCI [7]. If the number of re-evaluation detections affects a number M of resources smaller than N , the resource replacement must again select new resources considering the SCI limitations so that all retransmissions of the TB can be reserved. In this case, the new resources must be in a window of ± 31 slots around the selected resources that have not experienced a re-evaluation detection.

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 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][17]. It is important to highlight that the re-evaluation is an important novelty introduced in NR V2X mode 2 compared to LTE V2X mode 4.

2.5 Additional details

For more details about the 5G NR V2X standard, readers are referred to our related publication [15] (found in Annex A.1). In particular, our publication includes the following sections: Section II briefly reviews the LTE V2X standard. NR V2X has been designed to support advanced V2X services or use cases for connected and automated driving. Section III presents these use cases following the work done at the 3GPP and the 5G Automotive Association (5GAA). Section IV presents a high-level overview of the 5G NR system architecture for V2X communication for SL and uplink/downlink communications. Section V provides an in-depth summary of the NR V2X SL physical layer including its structure, the physical sidelink channels, as well as physical layer sidelink procedures. Section VI describes the resource allocation in NR V2X SL for its two communication modes (mode 1 and mode 2). Section VII describes the framework and mechanisms defined to manage quality of service (QoS) in NR V2X communications. Section VIII presents the major enhancements introduced to the uplink/downlink interface and to the mobility management in order to support Vehicle-to-Network (V2N) communications. NR V2X has been designed to complement LTE V2X. To that end, Section IX explains how

the two technologies can co-exist. Section X presents the evaluation methodology defined in 3GPP for NR V2X that includes new channel models and assumptions for system and link level simulations. Finally, Section XI presents future possible enhancements for NR V2X communications [15].

Section VI.B of our related publication [15] (found in Annex A.1) specifically describes NR V2X mode 2. Section VI.B includes the description of another novelty of NR V2X mode 2 compared to LTE V2X mode 4. In particular, NR V2X mode 2 introduces a pre-emption mechanism to prioritize traffic. With pre-emption, a vehicle with low priority traffic must free its reserved resource if it detects that another vehicle with higher priority will use the reserved resource. To do so, vehicles execute again the step 1 of the resource allocation algorithm before a TB transmission on a reserved resource to check if another vehicle with higher priority has reserved the same resource. It should be noted that the pre-emption mechanism is an optional feature of NR V2X mode 2, as it can be enabled or disabled. This contrasts with the re-evaluation mechanism, which is mandatory and must always be enabled according to 3GPP standards [6].



3 Simulation Environment

This thesis uses simulations to evaluate the performance of the standard NR V2X mode 2 and the different thesis contributions. Simulations offer the essential flexibility and precision required to evaluate large-scale scenarios, providing a suitable balance between accuracy, scalability, and repeatability. At the beginning of the work for this thesis, there was no open-source simulation environment available to conduct our evaluations, due to the recent standardization of 5G NR V2X communications. Therefore, one of the most significant tasks at the start of this thesis was the implementation of a new and standard-compliant NR V2X mode 2 simulator that did not previously exist in the community. This chapter describes the simulation platform that has been implemented and utilized (Section 3.1). It also details the scenarios employed and the configuration of the key parameters considered (Section 3.2), in addition to the evaluation metrics used throughout this thesis (Section 3.3).

3.1 Simulation platform

The NR V2X mode 2 simulator developed in this thesis adheres to the 3GPP standards [6][7] and guidelines [12]. The ns-3 simulator has been utilized for the implementation of NR V2X mode 2. This section provides a brief overview of the main features of ns-3 and details the standard-compliant implementation of the NR V2X mode 2 simulator employed in this thesis. It should be noted that the proposals to enhance NR V2X mode 2 have been built upon the standard-compliant implementation described in this section.

3.1.1 ns-3 simulator

3.1.1.1 Description

ns-3 is a discrete event network simulator [22]. It has been developed for research in network protocols. ns-3 is an open-source software project, licensed under the GNU GPLv2.

ns-3 is built on a core framework and a set of layers that add network functionality to the simulator. It provides an Internet protocol stack with implementations of protocols such as TCP and UDP, as well as lower-level protocols like 802.11. Various components and applications can be added to nodes to enable communication between them. ns-3 uses C++ as its primary programming language but also supports Python.

The ns-3 simulator is maintained and developed through the contributions of three main entities: ns-3 developers, ns-3 maintainers, and the ns-3 consortium. As part of the long-term planning and continuous growth and maintenance of ns-3, INRIA (Institut National de Recherche en Informatique en Automatique) and the University of Washington formed a consortium to support the ns-3 open-source project.

ns-3 has a significant presence in the research community, evidenced by the substantial number of publications related to the ns-3 simulator. A search for articles (last accessed in May 2024) on Google Scholar using the keyword 'ns-3' yields 18,400 results since 2020, excluding citations and patents. The same search on the IEEE Xplore Digital Library (last accessed in May 2024) results in 839 ns-3 related articles since 2020.

ns-3 is continuously evolving and adapting to new devices, networks, and models that emerge with advances in network technologies. Periodically (three times a year), a new version (release) of ns-3 is published, incorporating the latest advancements made by its extensive community of users and developers.

3.1.1.2 Configuration of simulations

Within the ns-3 software project, there is a series of directories containing all the necessary software components for running simulations. One such folder is “scratch”, which holds the global configuration files for the simulations. Thanks to the files in the scratch folder, various aspects can be configured for each simulation, including the number of nodes, network interfaces installed on the nodes, network traffic generation, node mobility, initial simulation setup, simulation scenario, and data collection.

A file called “NRV2X.cc” has been used as the scratch file to configure the simulations in this thesis. Specifically, it has been useful for configuring the scenarios and parameters, as described in Section 3.2, and for collecting data during the simulation to calculate the evaluation metrics presented in Section 3.3.

3.1.1.3 ns-3 modules

ns-3 modules constitute another fundamental component of the ns-3 software project. These ns-3 modules are directories within the "src" (source code) folder of ns-3, serving as "containers" for source code files that implement functionalities associated with various network protocols and standards in ns-3. Each ns-3 module corresponds to a network technology and typically includes a type of network interface (NetDevice) associated with that technology. This organization simplifies the installation of network technologies on the different nodes configured in the simulation. It is noteworthy that in ns-3, several network interfaces can be installed on each node.

Each ns-3 module contains a set of subdirectories: "model", "helper", and "examples". "Model" is arguably the most important element of the ns-3 modules as it includes all the functionality associated with each network technology. This subdirectory includes the network layers, network interface, and the entities necessary for the correct operation of the implemented protocol. "Helper" contains files that assist in the installation of network interfaces on each node. "Examples" includes examples to test the functionality implemented in each ns-3 module. These examples are actually configuration files that should be included in the scratch directory for execution. Some ns-3 modules are "Wifi" and "LTE".

In this thesis, a new ns-3 module called "NR-V2X" has been created and defined in ns-3 to implement NR V2X mode 2. This ns-3 module is detailed below in Section 3.1.2.

3.1.2 Implementation of 5G NR V2X in ns-3

This section describes the software implementation of the NR V2X mode 2 standard in the ns-3 simulator. It covers the network entities and layers that make up the "NR-V2X" ns-3 module developed within this PhD thesis.

3.1.2.1 Architecture of the implementation

This section describes the functionality of the "NR-V2X" ns-3 module created to implement the NR V2X mode 2 standard in ns-3. The overview provided here is high-level, with more detailed descriptions of each component in the following sections (sections 3.1.2.2, 3.1.2.3 and 3.1.2.4).

Figure 3 shows a general diagram corresponding to the NR V2X mode 2 implementation in ns-3. This diagram mainly illustrates the different network layers and entities, as well as the procedure followed by each packet to be transmitted and received by the nodes (in this case, vehicles) in the network.

The implemented layers are: NR-V2X-Mac (medium access or MAC layer), and NR-V2X-Phy (physical or PHY layer). Additionally, two entities necessary for the ns-3 simulation environment have been implemented: NR-V2X-NetDevice (network interface), and NR-V2X-Channel (radio channel).

From top to bottom in the diagram of Figure 3, we start with the NR-V2X-NetDevice entity. This entity is directly related to the ns-3 simulator and functions as a network interface between the Upper Layers (e.g. the application layer) and the lower layers (i.e. the medium access and physical layers).

The first layer is the medium access layer, named NR-V2X-Mac (Figure 3). Its primary function is the management of radio resources (in time and frequency domains) so that all vehicles can transmit their packets to the medium in a coordinated and distributed manner, without assistance from the network infrastructure. Two key modules of the NR-V2X-Mac layer are "Resource Allocation" and "Reservations Detection" (Figure 3). The "Resource Allocation" module manages the selection and reservation of each vehicle's resources based on the detected available resources. To determine available resources, the vehicle considers reservations made by other vehicles. The "Reservations Detection" module detects resource reservations made by other vehicles.

The NR-V2X-Phy layer (Figure 3) is responsible for configuring transmission parameters such as transmission power, modulation (e.g., QPSK, 16-QAM), and coding rate (e.g., r0.5, r0.7). Upon reception, it decodes the received packets by managing the received signal level, noise, and interference. Additionally, in NR V2X mode 2, the physical layer performs another crucial function: "Sensing" (Figure 3). "Sensing" provides updated channel state information, useful for the NR-V2X-Mac layer. Specifically, the "Sensing" functionality provides the MAC layer with information about the RSRP measured over the detected transmissions from other vehicles.

The lowest entity is associated with the radio channel, NR-V2X-Channel (Figure 3). This entity is responsible for delivering the packets transmitted by a vehicle to all nodes

configured as receivers. The channel implementation also models the shadowing and pathloss using the reference 3GPP models in [12].

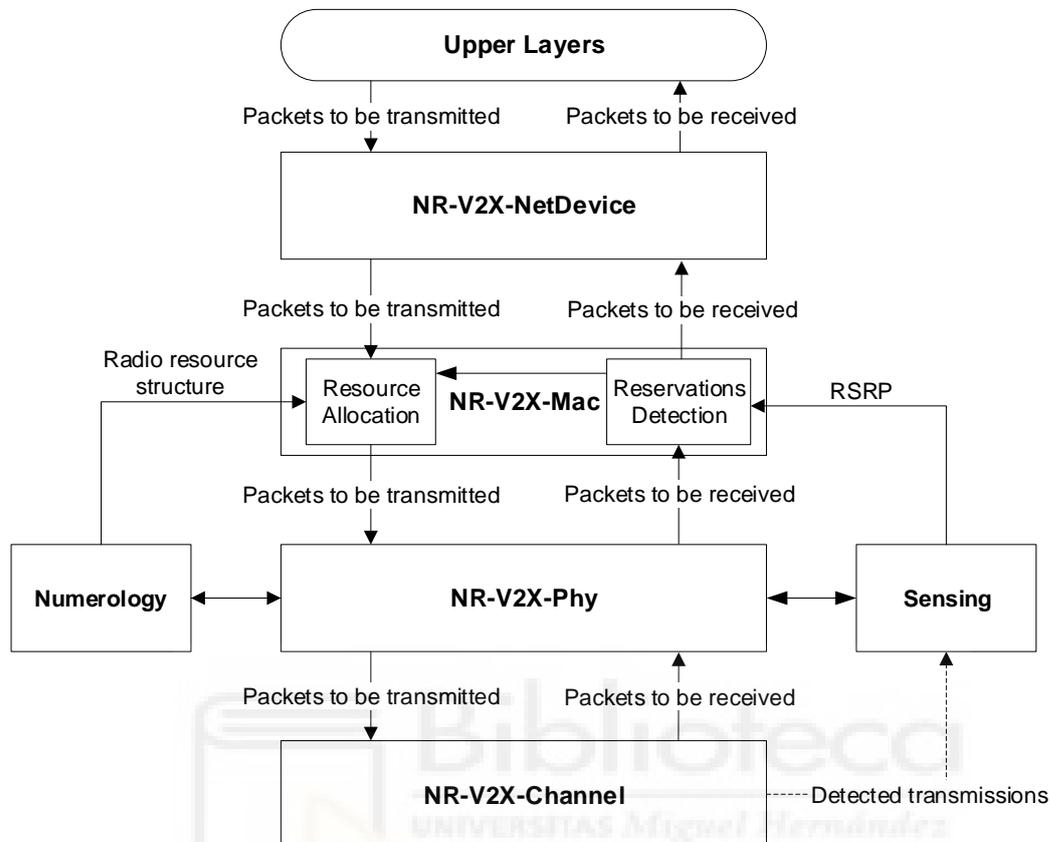


Figure 3. Architecture of the NR V2X mode 2 implementation in ns-3 [own elaboration].

Next, the procedure that a packet follows in the simulator from its generation in the Upper Layers of a vehicle to its arrival at the Upper Layers in the receiving vehicle is described, using Figure 3 as a reference.

Transmission process:

1. Packet generation in the Upper Layers of the transmitting vehicle.
2. The packet to be transmitted arrives at the NR-V2X-NetDevice and is directly forwarded to the NR-V2X-Mac layer.
3. NR-V2X-Mac applies the "Resource Allocation" algorithm, utilizing "Sensing" information previously received from the NR-V2X-Phy layer and considering detected reservations from other vehicles using "Reservations Detection". Additionally, it takes into account the resource structure to be used, which is configured via the "Numerology" module of the NR-V2X-Phy layer. After applying "Resource Allocation", the packet is scheduled for transmission on a specific radio resource.

4. The packet is sent to the NR-V2X-Phy layer to be transmitted to the NR-V2X-Channel with specific transmission parameters.

Channel:

1. The packet is passed to the channel model entity that includes the pathloss and shadowing.

Reception process:

1. The packet is first received by the NR-V2X-Phy layer. This layer checks if the packet and/or the SCI associated with the packet can be decoded with the reception parameters (received power, modulation, coding, interference).
2. The packet is passed to the NR-V2X-Mac layer, where it is verified for correct decoding. If not, it is discarded. The SCI is also checked for correct decoding. If the SCI is decoded correctly, the vehicle checks if the SCI contains any resource reservations using the "Reservations Detection" module. If so, the vehicle forwards this reservation information to the "Resource Allocation" module for consideration in the next resource allocation.
3. If the packet is correctly received, it is forwarded to the NR-V2X-NetDevice, which delivers it to the Upper Layers.

3.1.2.2 MAC layer implementation

Before delving into the description of the MAC layer, it is important to highlight the nomenclature used in ns-3 to indicate which layer or network entity a function belongs to. This nomenclature, inherited from the C++ programming language, will be used throughout this section and the following sections (Sections 3.1.2.3 and 3.1.2.4). The nomenclature is (layer/entity name)::(function name). It should be noted that for this section, all functions not explicitly referenced with their corresponding layer or entity will be considered to belong to the NR-V2X-Mac layer.

Figure 4 shows the diagram of the MAC layer. The part of the diagram that starts with the block "NR-V2X-NetDevice::Send" and ends with "NR-V2X-Phy::SendPhy" represents the part used for packet transmission. The part that starts with the block "NR-V2X-Phy::RxEnd" and ends with "NR-V2X-NetDevice::Receive" is associated with reception. The part of the diagram that begins with the block "NR-V2X-Phy::SensingRSRP" and ends

with the function "ResourceAllocation" is related to the sensing procedure, where sensing data is received from the NR-V2X-Phy layer.

Starting with transmission (Figure 4), the generated packets at the upper layers arrive at this layer via the "NR-V2X-NetDevice::Send" function and are directed to the "ResourceAllocation" function. This function applies the NR V2X mode 2 resource allocation scheme in use (Semi-Persistent Scheme, SPS or Dynamic Scheme, DS). The "ResourceAllocation" function first communicates with "rbsPerPacket" function to determine the numbers of RBs and therefore the number of sub-channels needed by the vehicle to transmit the packet with a specific size and with a particular MCS. The "ResourceAllocation" function also takes into account the radio resource structure defined by the "NR-V2X-Phy::Numerology" function. Then it checks the detected reservations in the received SCIs and the sensing information received from the NR-V2X-Phy layer (RSRP). With this information (number of sub-channels needed, structure of resources, detected reservations, and sensing data), the "ResourceAllocation" function schedules the transmission of each packet for a specific slot and one or more sub-channels ([Slot, Sub-channels] for TX) following the resource allocation algorithm of NR V2X mode 2. In the case of retransmissions, the "ResourceAllocation" function schedules each of the N transmissions of the packet in a different slot with the same number of sub-channels. It should be noted that the "ResourceAllocation" function also includes the implementation of the re-evaluation mechanism⁵. The next step involves configuring the PHY layer through the "configPhy" function. This configuration sets transmission parameters associated with the packet such as the transmission power and the MCS. After the "configPhy" function is executed, the packet is redirected to the "NR-V2X-Phy::SendPhy" function. If it is the last transmission of a packet in SPS, the reselection counter is decremented by one (resCounter--).

⁵ The pre-emption mechanism has not been implemented in the simulator because evaluating the transmission of traffic with different priorities within the same scenario is beyond the scope of this thesis.

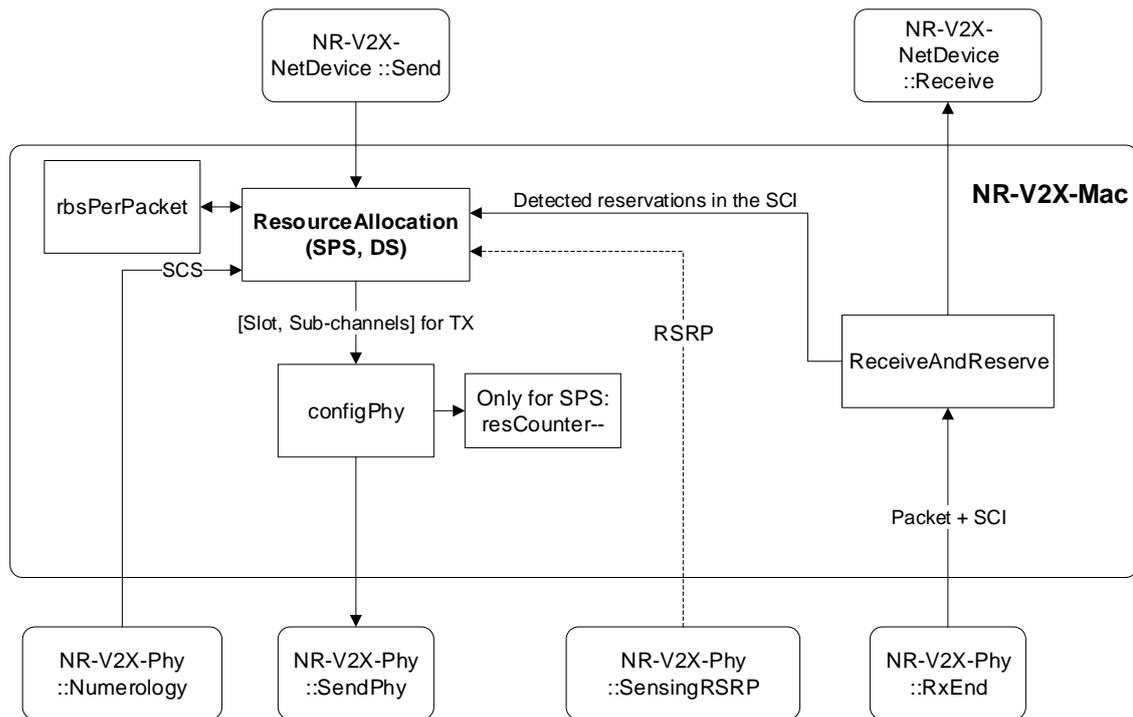


Figure 4. NR-V2X-Mac layer [own elaboration].

In the reception part of the diagram in Figure 4, packets are collected by the "NR-V2X-Phy::RxEnd" function and processed in the "ReceiveAndReserve" function. In this implementation, among other tasks, the NR-V2X-Phy layer checks whether packets are received correctly or incorrectly and flags them accordingly. The "ReceiveAndReserve" function checks this flag and discards the packet if the flag indicates that it has been received incorrectly. Additionally, this function logs resource reservations with the information received from the SCI associated with the packet. This information is sent to the "ResourceAllocation" function for the transmitting vehicle to consider in the next resource selection.

Finally, Figure 4 also shows the functionality associated with sensing. The "ResourceAllocation" function receives the RSRP measurement from the NR-V2X-Phy layer.

3.1.2.3 PHY layer implementation

It should be noted that for this section, all functions not explicitly referenced with their corresponding layer or entity will be considered to belong to the NR-V2X-Phy layer.

Figure 5 shows the diagram of the PHY layer. The part of the diagram that starts with the block "NR-V2X-Mac::configPhy" and ends with "NR-V2X-Channel::Send" represents the part used for packet transmission. The part of the diagram that starts with the block "NR-V2X-Channel::Receive" and ends with "NR-V2X-Mac::ReceiveAndReserve", using solid arrows, symbolizes reception. The part of the diagram that begins with "NR-V2X-Channel::Receive" and ends with the block "NR-V2X-Mac::ResourceAllocation", using dashed arrows, is associated with the sensing procedure. The state machine of the NR-V2X-Phy layer is also depicted, showing two states: transmission (TX) and reception (RX). Additionally, the block for the function "NR-V2X-Phy::Numerology" is included. This function configures the radio resources (i.e. slot duration and RB bandwidth) based on the configured numerology.

Starting with transmission, there is a single function, "SendPhy" (Figure 5). This function is responsible for forwarding packets from the MAC layer to the radio channel. It uses a transmission configuration set by the MAC layer via the function "NR-V2X-Mac::configPhy". Before sending the packet to the channel, to manage half-duplex communication, this function changes the state of the NR-V2X-Phy layer from RX to TX and schedules a return to RX state at the end of the slot where the packet is going to be transmitted.

The function that receives packets directly from the channel ("NR-V2X-Channel::Receive") is "RxStart", executed almost at the beginning of the slot (it only has the channel delay). This function stores the packets that will later attempt to be decoded in "RxMiddle" function, which is executed in the middle of the slot. "RxMiddle" communicates with the function "blockSuccessRate", which is used to obtain a Block Error Rate (BLER) value based on the measured SINR (Signal to Interference plus Noise Ratio). It should be noted that the simulator implements lookup tables (LUTs) from 3GPP working documents to model the transmissions of the packet or TB [23] and SCI [24] at the PHY layer. These LUTs relate the BLER vs the SINR. The obtained BLER values are used to model the correct or incorrect reception of the SCI and the packet. The final goal of the "RxMiddle" function is to flag the packet and its associated SCI, indicating whether they have been received correctly or not. Once flagged, each packet is passed to the "RxEnd" function, executed at the end of the slot. "RxEnd" aims to pass the packets to the MAC layer and remove the packets stored in the NR-V2X-Phy layer for the current slot.

In Figure 5, we can highlight the sensing functions of this layer. The channel is sensed to obtain the RSRP. To obtain the RSRP, the NR-V2X-Phy layer senses the resources used to transmit the packets (Resources of packets) for which the SCIs have been correctly decoded.

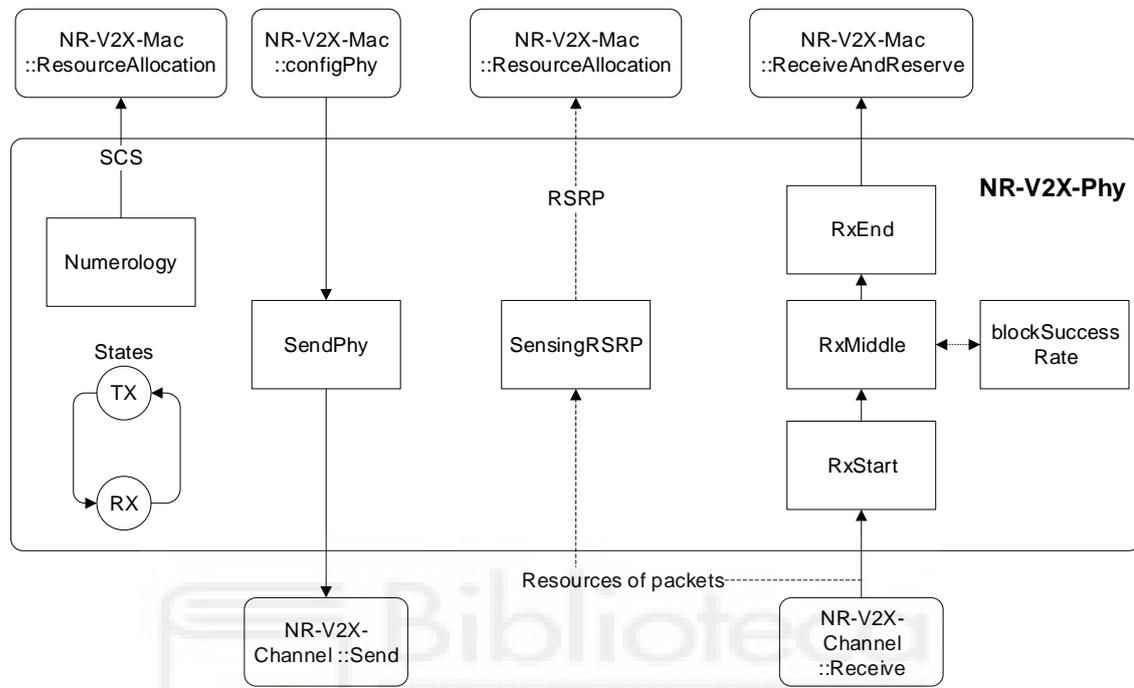


Figure 5. NR-V2X-Phy layer [own elaboration].

3.1.2.4 Radio channel implementation

It should be noted that for this section, all functions not explicitly referenced with their corresponding layer or entity will be considered to belong to the NR-V2X-Channel entity.

Figure 6 shows the diagram of the NR-V2X-Channel entity. The diagram begins with sending packets to the channel from the physical layer, specifically from the function "NR-V2X-Phy::SendPhy". The packets first arrive at the function "NR-V2X-Channel::Send", which schedules the transmission of the packets to all receivers whose PHY layers are "connected" to the NR-V2X-Channel.

The packets in the "NR-V2X-Channel::Send" function (Figure 6) are processed using the channel model. This channel model processes each packet using the channel model proposed by the 3GPP in TR 37.885 [12] for evaluating 5G NR V2X. This 3GPP channel model includes a pathloss model and a shadowing model. It also models Non-Line-of-Sight-vehicular (NLOS_v) losses, or signal losses caused by a vehicle located between the

transmitter and receiver. Additionally, a delay is applied based on the propagation, which depends on the distance between the transmitter and receiver. After this delay, the function "NR-V2X-Channel::Receive" is executed to send the received packets to the PHY layer of each receiving vehicle, achieved by passing the received packets to the function "NR-V2X-Phy::RxStart".

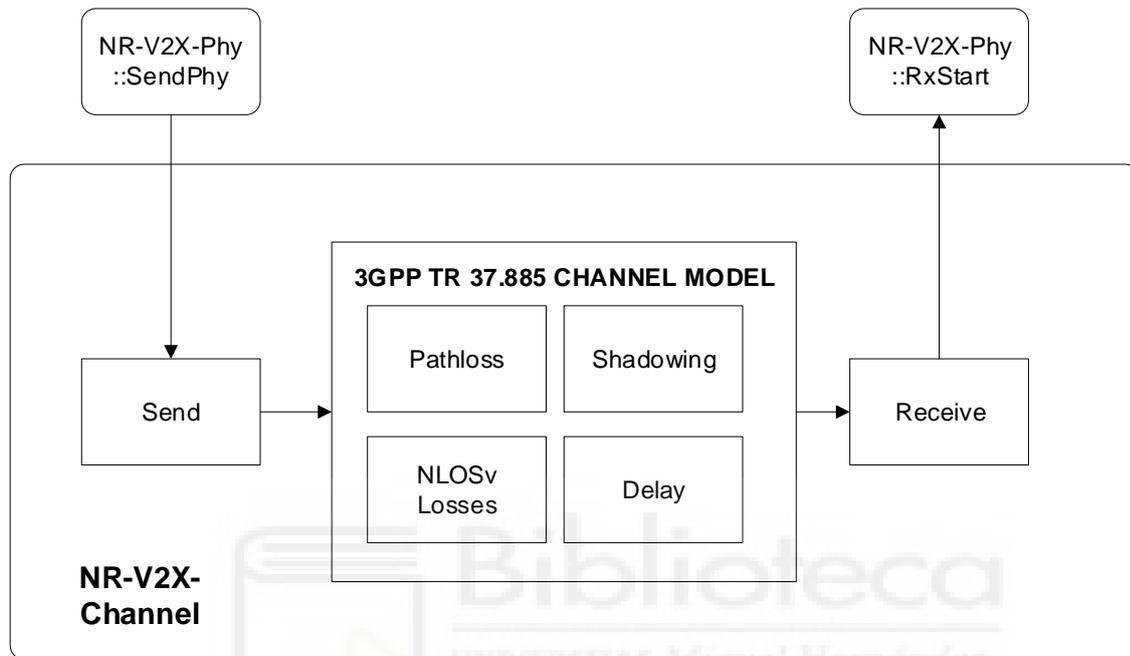


Figure 6. NR-V2X-Channel entity [own elaboration].

The channel modeling is primarily implemented through the "NR-V2X-Channel::Send" function. This function has two fundamental objectives: to send the packet being transmitted through the channel to all nodes except the transmitting vehicle, and to apply the channel model to the packet being transmitted. For the first objective, the ns-3 propagation delay model is used to deliver the packet to each receiver at a given time. By calculating the distance between the transmitter and each receiver, a variable delay is applied to the packet reception based on the distance.

For the second objective, the 3GPP channel model defined in TR 37.885 is applied, which include a pathloss and a shadowing model. In particular, the shadowing effects are modeled using a log-normal distribution with zero mean and a standard deviation of 3 dB with spatial correlation. With respect to the pathloss model [12], it should be noted that the one defined in the 3GPP channel model was not included in the class defining the different propagation models available in ns-3, "ns3::PropagationLossModel" [22]. As part of the

implementation, the new pathloss model has been implemented in that ns-3 class to associate it with the NR-V2X-Channel entity. It should be noted that the 3GPP channel model, when applied to the highway scenario included in the implementation, considers two possible link states between each pair of transmitting and receiving vehicles: Line-of-Sight (LOS) and Non-Line-of-Sight vehicular (NLOSv). The 3GPP channel model defines a probability that a link is in the NLOSv state based on the distance between the transmitter and receiver. This probability and the probability that a link is in the LOS state are shown in Table 1. These probabilities have been implemented in the NR-V2X-Channel entity so that the state of each link between two vehicles can be modeled at any given moment. Table 2 shows the pathloss model and NLOSv losses to be applied according to the 3GPP channel model depending on the link state (LOS or NLOSv). Both link states in the highway scenario use the same pathloss model; however, if the link state is NLOSv, additional NLOSv losses are applied. The simulator implements the pathloss model and the NLOSv losses included in Table 2.

Table 1. Probability of LOS and NLOSv link states

Link state	Computation of P(NLOSv) and P(LOS)
LOS	If $d \leq 475$ m, $P(\text{LOS}) = \min\{1, ad^2 + bd + c\}$ where $a=2.1013 \cdot 10^{-6}$, $b=-0.002$ and $c=1.0193$ If $d > 475$ m, $P(\text{LOS}) = \max\{0, 0.54 - 0.001(d-475)\}$
NLOSv	$P(\text{NLOSv}) = 1 - P(\text{LOS})$

Where:

- d is the distance between the transmitting and receiving vehicles (meters).
- $P(\text{LOS})$ is the probability that the link is in the LOS state.
- $P(\text{NLOSv})$ is the probability that the link is in the NLOSv state.
- a , b and c are constants.

Table 2. Pathloss model and NLOSv losses depending on the link state

Link state	Pathloss [dB]	NLOSv losses [dB]
LOS	$PL = 32.4 + 20 \log_{10}(d) + 20 \log_{10}(f_c)$	0
NLOSv	$PL = 32.4 + 20 \log_{10}(d) + 20 \log_{10}(f_c)$	Modeled by a log-normal distribution: Mean = $5 + \max(0, 15 \cdot \log_{10}(d) - 41)$ dB Standard deviation = 4 dB

Where:

- PL is the pathloss.
- d is the distance between the transmitting and receiving vehicles (meters).

- f_c is the operating frequency (Hz).

3.2 Scenarios and parameters

To evaluate the standard NR V2X mode 2 and the different proposals of this thesis, we have used the following simulation scenarios and parameters:

We consider the 3GPP 5 km highway scenario with 3 lanes in each direction [12], vehicles driving at 70 km/h, and vehicle densities of 25, 50 and 100 veh/km. Vehicles generate packets or TBs following the 3GPP aperiodic and periodic traffic models [12]. In the aperiodic model, vehicles generate TBs with an inter-packet arrival time $\tau = \bar{x} + x(\bar{x})$, where \bar{x} is a constant minimum inter-packet arrival time and x is an exponential random variable of mean equal to \bar{x} . We evaluate low ($\bar{x} = 100$ ms), medium ($\bar{x} = 50$ ms) and high intensity ($\bar{x} = 10$ ms) scenarios, and the PDB is set equal to \bar{x} in all three scenarios. For aperiodic traffic of variable size, the size of TBs is uniformly distributed in the [200, 1200] bytes range (with a 200-byte step). For aperiodic traffic of fixed size, the size of TBs is equal to 200 bytes. We also evaluate periodic traffic of fixed or variable size. For periodic traffic of fixed size, the size of TBs is equal to 190 bytes. For periodic traffic of variable size, vehicles generate TBs with a size that follows the pattern {300, 190, 190, 190, 190} bytes and a random starting point. The inter-packet arrival time for the low, medium and high intensity periodic scenarios is set equal to 200 ms, 100 ms and 20 ms, respectively. The PDB for periodic scenarios is equal to the inter-packet arrival time. Two additional scenarios are defined and evaluated based on the traffic generated by vehicles: single traffic and mixed traffic scenarios. In the single traffic scenario, all vehicles generate traffic with the same intensity (low, medium or high intensity). In the mixed traffic scenario, 80% of vehicles generate traffic with medium intensity and the remaining 20% generate traffic with high intensity.

NR V2X mode 2 is set to operate over 20 MHz in the 5.9 GHz band using a subcarrier spacing of 30 kHz. The sub-channel size is set to 12 RBs, and there are 4 sub-channels per slot. The sensitivity has been set to -103.5 dBm according to [25], and the transmission power has been set to 23 dBm. Vehicles use the MCS index 13 (16QAM modulation and coding rate equal to 0.5) to transmit each TB. In this case, TBs of 190, 200, 300, 400, 600, 800, 1000 and 1200 bytes require 1, 1, 2, 2, 3, 3, 4 and 4 sub-channels, respectively, to be

transmitted along with their associated SCI. In all simulations, we consider broadcast transmissions.

For both DS and SCS scheduling schemes, T_0 and $T_{proc,0}$ define the bounds of the sensing window, and they are configured equal to 2200 slots and 1 slot, respectively. T_1 , T_{2min} and T_2 define the bounds of the selection window. T_1 and T_{2min} are set equal to 2 slots, and T_2 is set equal to the PDB. T_2' is set so that the last slot of the new selection window defined with the re-evaluation mechanism SW' ($n'+T_2'$) and the last slot of the original selection window SW ($n+T_2$) coincide. The RSRP threshold used in step 1 to decide if a resource is excluded or not is set to -128 dBm (i.e. its minimum value) following [26]. The threshold for the percentage of resources that vehicles must have available after executing the step 1 is set to 20% (i.e. its minimum value [27]). We consider that vehicles select 1 or 2 resources in step 2 (i.e. N is equal to 1 or 2) to perform 1 or 2 transmissions per TB, respectively. T_3 is configured equal to 5 slots following [7]. For the SPS scheme, the probability P_k to keep the same resources has been set equal to 0, and two different strategies for the selection of the *RRI* are evaluated:

- Average *RRI*: the *RRI* is set equal to the average inter-packet arrival time.
- Minimum *RRI*: the *RRI* is set equal to the minimum of the inter-packet arrival time.

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 $2\bar{x}$, while the minimum one sets it equal to \bar{x} .

Finally, vehicles execute a re-evaluation check in all selected resources, including the selected resources available with SPS after an unutilized reservation.

3.3 Evaluation metrics

The following metrics are used to evaluate the standard NR V2X mode 2 and the different proposals in this thesis:

- Packet Delivery Ratio (PDR): ratio 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 or TB transmissions that have experienced at least a re-evaluation detection.
- Propagation Losses Ratio (PLR): ratio of TBs that cannot be successfully decoded due to propagation losses to the total number of transmitted TBs. These losses occur when the received signal level is below the sensitivity threshold or the Signal-to-Noise Ratio (SNR) is too low.
- Packet Collision Ratio (PCR): ratio 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. PCR excludes propagation errors.
- Size Reselection Ratio (SRR): ratio of TBs that produce a reselection due to the size of the TBs over the total number of transmitted TBs [13].
- Latency Reselection Ratio (LRR): ratio of TBs that produce a reselection due to the latency requirement of the TBs over the total number of transmitted TBs [13].
- Unutilized Reservation Ratio (URR): ratio of reserved resources that are not utilized because the TB is generated after the reserved resource over the total number of reserved resources. This metric does not account for unutilized reservations that are considered in the size and latency reselection ratios⁶ [13].
- Re-evaluation Check Ratio (ReCR): ratio of TBs that are checked for re-evaluation at least once over the total number of transmitted TBs.
- Re-evaluation Detection Ratio (ReDR): ratio of TBs that experience at least one re-evaluation detection and trigger the selection of a new resource over the total number of transmitted TBs.

⁶ Size and latency reselections also produce unutilized reservations when the vehicle has reserved a resource and the new TB cannot be transmitted in that reserved resource because the resource does not fulfil the size or latency requirement of the new TB. Further details about the impact of these procedures on NR V2X mode 2 are described below in Section 4.2.

4 5G NR V2X MAC Challenges

4.1 Introduction

NR V2X has been designed to complement, and not replace, LTE V2X whose aim is to support basic safety applications. This chapter is focused on analyzing the efficiency of NR V2X mode 2 to support advanced V2X services. These services are expected to generate V2X aperiodic traffic of variable size according to 3GPP [12]. Variable traffic patterns were shown to significantly impact the operation and performance of LTE V2X mode 4 [28]. This was due to certain MAC inefficiencies when vehicles generate aperiodic messages of variable size that result in packet collisions and require additional solutions [13][29]. NR V2X mode 2 can operate using SPS with similar procedures as those present in LTE V2X mode 4. A key mandatory feature introduced in the MAC of NR V2X mode 2 is the re-evaluation mechanism that is designed with the objective of detecting and avoiding detected packet collisions (see Section 2.4) [6]. Previous studies have evaluated the performance of NR V2X mode 2 under different traffic patterns [26][30]-[32]. However, these studies do not implement the re-evaluation mechanism. In addition, these studies only consider periodic or aperiodic traffic of fixed size. However, the 3GPP evaluation methodology guidelines for NR V2X [12] recommend traffic generation models for advanced V2X services that also include traffic of variable size in line with the message patterns characteristic of Day 2 or Day 3 V2X services such as cooperative perception [33] or maneuver coordination [34]. In this context, our study published in [16] (Annex A.2) and summarized in this chapter is the first evaluation of a fully standard compliant implementation of NR V2X mode 2 using the re-evaluation mechanism. The evaluation considers different traffic patterns including periodic and aperiodic traffic with fixed and variable packet sizes based on the models reported in Section 3.2 and in the 3GPP evaluation guidelines [12]. The evaluation is conducted in scenarios where vehicles perform one transmission per TB and execute the re-evaluation check in the mandatory slot indicated by the standard that is located T_3 slots before the initially selected resource (see Section 2.4). This chapter focuses on the SPS scheme which, unlike DS, does not require the implementation of retransmissions to generate reserved resources and reserved resources provide stability to the operation of NR V2X mode 2 (see Section 2.3). The

results of this chapter show that NR V2X mode 2 also faces MAC challenges to handle aperiodic packets of variable size resulting in increased packet collisions. This occurs despite utilizing the re-evaluation mechanism designed to detect and avoid collisions. The MAC challenges faced by NR V2X mode 2 are described in Section 4.2 and then quantified in Section 4.3. Finally, Section 4.4 presents the conclusions of the chapter and summarizes additional results included in the related publication [16] (Annex A.2).

4.2 Impact of packet variability on 5G NR V2X

SPS reserves resources for a consecutive number of TBs defined by the *Reselection Counter*, with a time gap of RRI ms between reservations. SPS is well-suited for transmitting periodic traffic of fixed size. In this case, SPS only reselects resources when the *Reselection Counter* is depleted (based on $1-P_k$), an event termed *counter reselection*. Stability in SPS occurs when all reservations are used for transmitting TBs, and reselections occur only after the *Reselection Counter* depletion. Like in LTE V2X mode 4 [13], additional reselections may occur if the TB size or inter-arrival time between TBs changes. These additional reselections can increase instability and the risk of packet collisions. This is the case because neighboring vehicles are not aware of new selected resources until the next TB is transmitted, and the transmitting vehicle announces reservations for subsequent TBs. Therefore, the probability of 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.

4.2.1 Size reselections

SPS triggers a resource reselection when the size of a new TB exceeds the capacity of previously reserved resources in terms of number of sub-channels. We refer to this MAC challenge as *size reselection* and depict it in Figure 7. Let's suppose that a vehicle, V_A , generates a TB (e.g., 200 bytes) at slot s_{G1} and selects a sub-channel for transmission at slot s_{R1} . While transmitting the TB at s_{R1} , V_A announces in the associated SCI that the same sub-channel is reserved at slot $s_{R2}=s_{R1}+RRI$. If V_A generates a new TB (e.g., 600 bytes) at slot s_{G2} that cannot fit within the current resource reservation at s_{R2} , V_A must discard the reserved resources at s_{R2} . Subsequently, V_A needs to select new resources capable of accommodating the new TB size (e.g., two sub-channels in Figure 7). Finally, V_A transmits the TB using the reselected resources at slot s_{R3} .

4.2.2 Latency reselections

A vehicle may need to select new resources if its current reservation cannot meet the latency requirements (i.e., PDB) of a new TB. We refer to this MAC challenge as *latency reselection* and depict it in Figure 7. A latency reselection occurs when a vehicle generates aperiodic traffic and the selected *RRI* exceeds the TB's latency deadline or PDB. In this case, V_A reserves one sub-channel at slot s_{R4} and generates its next TB at slot s_{G3} , with its latency deadline set at slot s_L . The current reservation at s_{R4} fails to meet the TB's latency requirements since $s_{R4} > s_L$. Consequently, V_A must discard the reserved resources at s_{R4} . Finally, V_A needs to select new resources by the latency deadline of the new TB (e.g., at s_{R5} in Figure 7).

4.2.3 Unutilized reservations

The stability of SPS may also be compromised when reserved resources are not utilized by a vehicle. This occurs when no TB is ready for transmission at the slot where resources are reserved. This MAC challenge is termed *unutilized reservations*. Illustrated in Figure 7, unutilized reservations arise when the inter-arrival time between generated TBs exceeds the selected *RRI*. For example, if V_A reserves one sub-channel at slot $s_{R6} = s_{R5} + RRI$, but the next TB is generated at slot s_{G4} with $s_{G4} > s_{R6}$, V_A leaves the reserved resource at s_{R6} unused. Importantly, V_A cannot leverage its transmission opportunity at s_{R6} to announce the reservation at slot s_{R7} in the SCI. This leaves the transmission at s_{R7} prone to packet collisions, as the resource at s_{R7} is a selected (i.e. not reserved) resource.

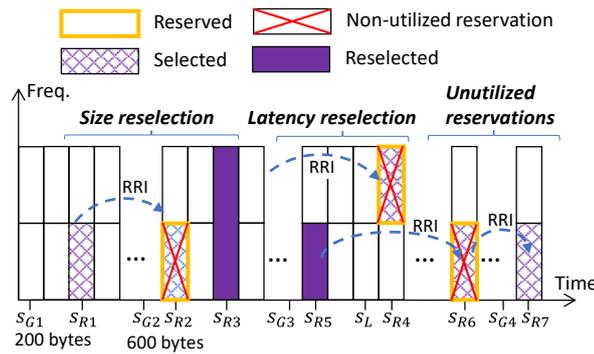


Figure 7. Impact of packet variability on SPS under NR V2X mode 2 [17].

4.2.4 Unused sub-channels

Another MAC challenge stemming from the variation of the TB size that might impact SPS occurs when a vehicle reserves a resource with a specific number of sub-channels, but the subsequent TB generated requires fewer sub-channels for transmission. This MAC

challenge is termed *unused sub-channels*. Let's consider the case where a vehicle generates a 600-byte TB and selects two sub-channels for its transmission at slot s_R . With this transmission, the vehicle reserves a resource at slot $s_R + RRI$ consisting of two sub-channels. Subsequently, the vehicle generates a smaller TB (200 bytes) that only requires a single sub-channel for transmission at slot $s_R + RRI$, leaving one sub-channel unused. However, other vehicles perceive this resource at slot $s_R + RRI$ as reserved for two sub-channels, thereby decreasing the availability of candidate resources and consequently increasing the risk of packet collisions.

4.3 Evaluation

It should be noted that, since this chapter focuses on SPS, the term NR V2X mode 2 will always refer to SPS throughout the rest of the chapter. In this section, the evaluation of NR V2X mode 2 is conducted in the single traffic scenario (see Section 3.2). Figure 8 compares the PDR achieved by NR V2X mode 2 when vehicles transmit periodic traffic with fixed packet size and aperiodic traffic with fixed or variable packet size. Results are also reported in Figure 8 for the two implemented *RRI* strategies that set the *RRI* value to the average and minimum inter-packet arrival time, and for the medium⁷ (Figure 8-left) and high (Figure 8-right) traffic intensity scenarios (see Section 3.2). The results reported in Figure 8 show that NR V2X mode 2 achieves a better PDR performance when vehicles generate TBs periodically compared to the case when they generate TBs aperiodically. The obtained results show that the differences in the PDR achieved by NR V2X mode 2 with periodic and aperiodic traffic increase with the vehicle density and traffic intensity [16].

Table 3 reports a set of key metrics that show that NR V2X mode 2 achieves better performance under the presence of periodic traffic than aperiodic traffic. These metrics are: the ratio of TBs that produce a reselection due to the latency requirement of the TB (LRR), the ratio of TBs that produce a reselection due to the size of the TB (SRR), the ratio of reserved resources that are not utilized because the TB is generated after the reserved resources (URR⁸), and Unused Sub-channels Ratio (USR) that is the ratio of unused sub-channels in the resources used to transmit the TBs. It should be noted that the SRR, LRR,

⁷ The medium traffic intensity scenario defined for this thesis in Section 3.2 corresponds to the low traffic intensity scenario defined in our publication [16] (Annex A.2).

⁸ This metric does not account for unutilized reservations that are considered in the size and latency reselection ratios.

URR and USR metrics do not vary with the vehicle density because they only depend on the traffic characteristics and on the reservations that each vehicle individually generates.

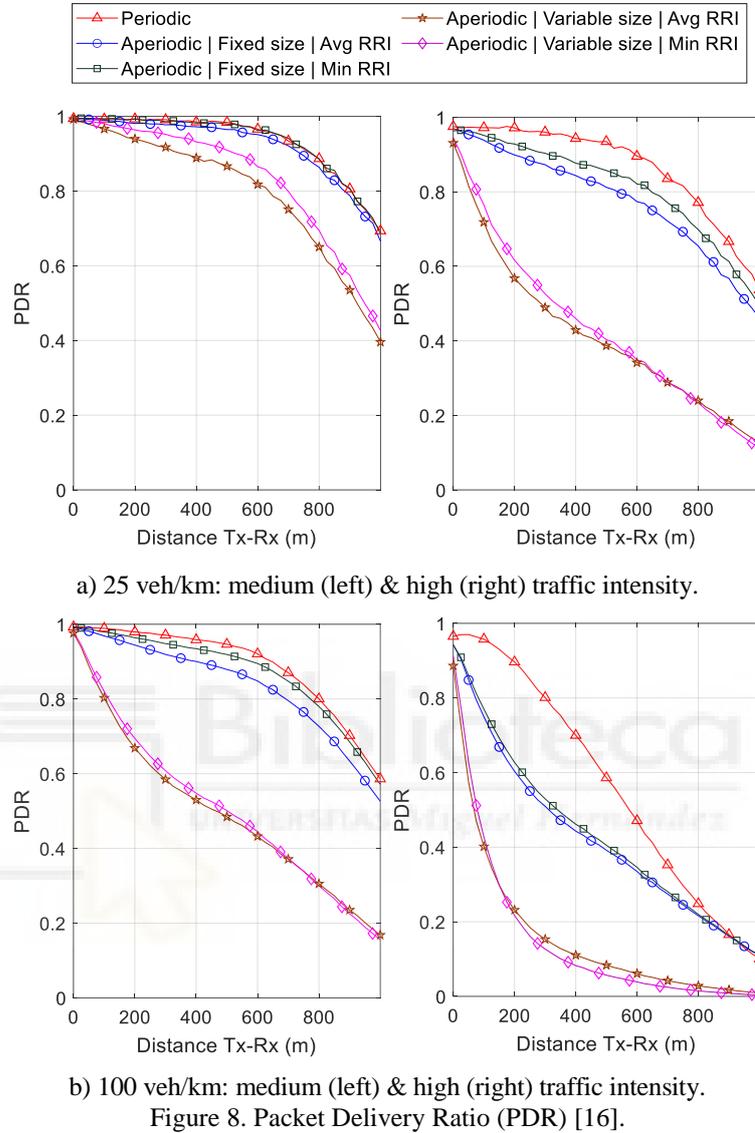


Figure 8. Packet Delivery Ratio (PDR) [16].

Table 3 shows that NR V2X mode 2 does not experience any of the MAC challenges described in Section 4.2 under periodic traffic of fixed size. This behavior brings significant stability to the operation of NR V2X mode 2 that results in the better PDR performance reported in Figure 8. Table 3 also shows that this is not the case under the presence of aperiodic traffic. For example, the variations in the time between TBs can cause latency reselections measured by LRR in Table 3.a). As it is shown in Table 3.a), the chosen *RRI* strategy highly impacts the LRR generated by aperiodic traffic. When the average *RRI* strategy is implemented, more than 50% of the TBs trigger a latency reselection [16]. The higher number of resource reselections makes NR V2X mode 2 more unstable and likely to provoke packet collisions (see Section 4.2). The minimum *RRI* strategy significantly

reduces the LRR compared to the average *RRI* strategy (see Table 3.a). However, the side effect of utilizing the minimum *RRI* strategy is the higher values of unutilized reservations measured by URR (Table 3.c). The unutilized reservations can also negatively impact the operation of NR V2X mode 2 (see Section 4.2.3).

Table 3. NR V2X mode 2 operation for the medium and high traffic intensity scenarios

<i>RRI</i> strategy	<i>Periodic</i>		<i>Aperiodic Fixed size</i>		<i>Aperiodic Variable Size</i>	
	Medium	High	Medium	High	Medium	High
<i>a) Latency Reselection Ratio (LRR) in %</i>						
Min RRI	0	0	1	5	1	5
Avg RRI			55	58.1	55	57
<i>b) Size Reselection Ratio (SRR) in %</i>						
Min RRI	0	0	0	0	5	5
Avg RRI			0	0	26	27
<i>c) Unutilized Reservation Ratio (URR) in %</i>						
Min RRI	0	0	60	59	56	57
Avg RRI			6	6	4	4
<i>d) Unused Sub-channels Ratio (USR) in %</i>						
Min RRI	0	0	0	0	29	29
Avg RRI			0	0	27	27

Figure 8 shows that NR V2X mode 2 sees its PDR degrade when vehicles generate aperiodic traffic with variable packet size with respect to aperiodic traffic with fixed packet size. Aperiodic traffic with variable packet size experiences not only the latency reselections and unutilized reservations, but also size reselections and unused sub-channels that are computed in Table 3 by means of the SRR and USR (see Section 4.2). Table 3.b) shows that the implemented *RRI* strategy also impacts the SRR. In particular, the obtained results show a significant reduction in the SRR when NR V2X mode 2 is configured with the minimum *RRI* strategy with respect to the average *RRI* strategy [16]. However, this is at the cost of leaving some sub-channels unused (see Table 3.d) that other vehicles cannot utilize since they are reserved. Table 3 shows that there are two trade-offs when NR V2X mode 2 is configured with the minimum or average *RRI* strategies: one in terms of LRR and URR, and another one in terms of SRR and USR. Figure 8 shows that these trade-offs compensate each other. In particular, Figure 8 shows that these trade-offs result in a slightly higher PDR with the minimum *RRI* strategy under aperiodic traffic of fixed size and in slight differences in the PDR under aperiodic traffic of variable size.

The analysis performed above has highlighted the MAC challenges that NR V2X mode 2 faces to support aperiodic traffic with fixed and variable packet sizes. The analysis is

revealing that these MAC challenges end up making the NR V2X mode 2 more unstable with the potential risk of causing additional packet collisions. Figure 9 reports the Packet Collision Ratio (PCR) as a function of the distance between the transmitter and the receiver for all the evaluated traffic patterns and *RRI* strategies when the vehicle density is 100 veh/km. The results reported in Figure 9 help identify that the degradation in PDR between the scenarios with periodic and aperiodic traffic (Figure 8) is due to higher packet collisions experienced with aperiodic traffic [16].

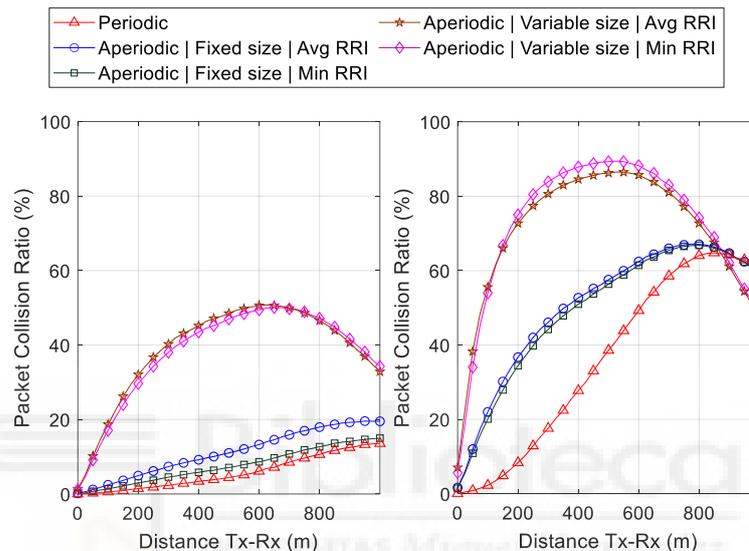


Figure 9. Packet Collision Ratio (PCR) under medium (left) & high (right) traffic intensity; 100 veh/km – similar trends observed for other vehicle densities [16].

As it was demonstrated in [13], LTE V2X mode 4 shows similar inefficiencies and high PCR values when transmitting aperiodic traffic with variable packet size. NR V2X mode 2 faces similar challenges despite introducing the re-evaluation mechanism originally designed with the objective to reduce the packet collisions. The operation of re-evaluation is measured with the ratio of TBs that experience one re-evaluation detection and trigger the selection of a new resource because a potential collision is identified (ReDR). The ReDR experiences levels up to 30% with aperiodic traffic of variable size. Higher ReDR levels are measured with increasing vehicle densities and thus with a higher likelihood that two vehicles select the same resources. The obtained results show that even although re-evaluation is avoiding the detected collisions for a non-negligible percentage of TBs, it is not capable of counteracting the high packet collisions caused by the instability that aperiodic traffic is introducing in the operation of NR V2X mode 2 [16].

4.4 Conclusions and additional results

This chapter has described the MAC challenges that NR V2X mode 2 can face under aperiodic traffic of variable size. These MAC challenges have been quantified with the evaluation of NR V2X mode 2 presented in this chapter. This evaluation is the first fully standard compliant evaluation of NR V2X mode 2 under periodic and aperiodic traffic of fixed and variable packet size in accordance with the 3GPP traffic model recommendations. The configuration of NR V2X mode 2 has considered two strategies that set the *RRI* to the minimum or to the average inter-packet arrival times. Both strategies result in multiple trade-offs in terms of different MAC challenges, but none reduces the resource management instability of NR V2X mode 2 observed when transmitting aperiodic traffic with variable packet size. These instabilities reduce the PDR and increase packet collisions with respect to scenarios where vehicles generate TBs periodically. The obtained results present similar trends as those observed with LTE V2X mode 4 when transmitting aperiodic traffic, despite the new and mandatory re-evaluation mechanism of NR V2X mode 2 that identifies and should, in principle, avoid possible packet collisions [16].

Our publication related to the evaluation of NR V2X mode 2 presented in this chapter [16] (found in Annex A.2) includes the following additional results:

- Evaluation of PDR, PCR and ReDR metrics in the mixed traffic scenario (see Section 3.2). Remarkably, similar conclusions are drawn from this scenario, i.e. aperiodic traffic of variable size significantly increases the packet collisions experienced by NR V2X mode 2 with respect to periodic traffic of fixed size. This is the case despite achieving ReDR levels of up to 44.6% in the mixed traffic scenario.
- Computation of the Reselection Counter Depletion Ratio (RCDR) metric, which quantifies the fraction of Reselection Counters that deplete over the total number of initiated Reselection Counters. This metric serves as an alternative measure for assessing the stability of NR V2X mode 2 with both *RRI* selection strategies under the different traffic patterns.
- Computation of the Channel Busy Ratio (CBR) metric, which quantifies the fraction of sub-channels experiencing an RSSI higher than a threshold within an observation window of $100 \cdot 2^u$ slots. CBR escalates with both vehicle density and

traffic intensity, reflecting the increased number of transmissions occurring in the scenarios.

- Computation of the ratio of TBs that are checked for re-evaluation (ReCR) under aperiodic traffic of variable size.



5 Impact of Re-Evaluation on 5G NR V2X

5.1 Introduction

Chapter 4 has analyzed the performance of NR V2X mode 2 with the Semi-Persistent Scheduling (SPS) scheme under various data traffic patterns when there are no retransmissions. The study has shown that NR V2X mode 2 performance degrades with aperiodic traffic of variable size compared to periodic traffic of fixed size due to increased packet collisions. This occurs despite the re-evaluation mechanism designed to detect and avoid these collisions. Chapter 4 provides an initial system-level evaluation of NR V2X mode 2 but does not explain why re-evaluation is not effective in certain scenarios. Understanding this is crucial for optimizing NR V2X mode 2. It is also necessary to determine if these observations hold when retransmissions are used. In this context, this chapter provides an in-depth analysis and evaluation of the impact of the re-evaluation mechanism on the operation and performance of NR V2X mode 2. This study is the first one that analyzes when and why re-evaluation is effective or not in detecting and avoiding packet collisions. This chapter focuses on SPS because aperiodic traffic of variable size mostly impacts SPS, since the Dynamic Scheduling (DS) scheme selects new resources for each TB (see Section 4.2). The effectiveness of re-evaluation is then particularly relevant when utilizing the SPS scheme. The analysis considers vehicles transmitting periodic or aperiodic packets of fixed or variable size, following 3GPP guidelines [12]. The study shows re-evaluation effectively avoids packet collisions with periodic traffic of fixed size, but its impact is small since the number of packet collisions detected by re-evaluation is low with this traffic. With aperiodic traffic of variable size, the effectiveness of re-evaluation decreases. The capacity of re-evaluation to detect and avoid collisions improves when retransmissions are considered. However, the impact of re-evaluation on the performance of NR V2X mode 2 is small since the benefit of retransmissions prevails over the gains obtained with the packet collisions avoided with re-evaluation.

The evaluation presented in Chapter 4 considers vehicles executing re-evaluation checks only in the mandatory slot according to the standard that is located T_3 slots before the selected resources (see Section 2.4). The 3GPP standard, however, allows different strategies for when and how often to perform these checks [6]. One strategy, used in

Chapter 4, executes one re-evaluation check just before packet transmission, while other strategies perform multiple checks. This chapter also evaluates the impact of these different re-evaluation check strategies on the performance of NR V2X mode 2 under aperiodic traffic of variable size. In particular, it evaluates two different strategies proposed by the 3GPP standard for NR V2X mode 2. The evaluation reveals that the two standardized strategies balance transmission latency and the computational cost of re-evaluation. The chapter then proposes an alternative strategy that reduces transmission latency and computational cost without degrading the reliability of NR V2X mode 2 transmissions.

The rest of this chapter is organized as follows. Section 5.2 presents an 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 5.3 evaluates the impact of re-evaluation without retransmissions, and Section 5.4 extends the analysis to the scenario where retransmissions are considered [17]. The evaluations presented in Section 5.3 and Section 5.4 consider that vehicles execute the re-evaluation check only in the slot that is mandatory in the standard. Section 5.5 describes the different re-evaluation check strategies evaluated, and Section 5.6 evaluates the impact of using the different re-evaluation check strategies. Finally, Section 5.7 summarizes the main outcomes of this chapter and some additional results included in the related publications [17][18] (found in Annexes A.3 and A.4).

5.2 Analysis of the Re-evaluation Mechanism

This Section analyzes the operation and discusses the effectiveness of the three phases of the re-evaluation mechanism: the re-evaluation check, the re-evaluation detection, and the resource replacement.

5.2.1 Re-evaluation check

Vehicles use the re-evaluation check to assess whether selected resources are still available or not 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. By design, SPS only selects new resources as a result of a counter reselection (i.e. when the reselection counter is depleted). 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 4.2) increase the fraction of TBs that are transmitted on selected resources, and thus increase the number of re-evaluation checks [17].

5.2.2 Re-evaluation detection

A re-evaluation detection is triggered after a re-evaluation check when the initially selected resources are no longer available. If this is the case, the re-evaluation mechanism detects a packet collision. This section identifies the packet collisions that can be detected or not by re-evaluation and discusses the effectiveness of the re-evaluation detection phase.

5.2.2.1 Detectable and undetectable packet collisions

The re-evaluation mechanism can only detect a subset of packet collisions. When a vehicle V_A executes the re-evaluation check before the transmission of a TB, it can avoid a collision if it detects that another vehicle V_B has reserved the same resources it selected. On the other hand, V_A cannot detect a potential collision with V_B on the same selected resources if V_B did not announce previously its reservation. This can occur because V_B is transmitting the first TB after a new reselection or because of a previous unutilized reservation (see Section 4.2). In the latter case, V_B could not announce in the unutilized resources that it was reserving the resources where a collision could take place with V_A . We should remember that a new reselection of resources can occur when the reselection counter is depleted and after a size or latency reselection. Overall, unutilized reservations and latency and size reselections increase the number of transmissions in selected (not reserved) resources, and therefore the probability of undetectable collisions by re-evaluation. It should be noted that undetectable collisions can occur when each TB is transmitted once and also when retransmissions are considered. With retransmissions, re-evaluation cannot detect packet collisions between the initial transmissions of TBs on selected resources. In fact, the resource for the initial transmission of the TB is always not reserved (i.e. it is a selected resource) after a new reselection or an unutilized reservation. This is in contrast to the

resource for the retransmission of a TB that can be reserved with the reservation announced in the initial transmission of the TB (see Section 2.3).

5.2.2.2 Effectiveness of Re-Evaluation Detections

Unutilized reservations and size and latency reselections reduce the effectiveness of re-evaluation detections. Let's suppose that V_B has announced the reservation of a resource to transmit a new TB, but this TB generates a size or latency reselection. In this case, V_B drops its transmission and does not utilize its initially reserved resource and selects a new resource. In parallel, let's suppose that V_A had selected the same resources as V_B and executes a re-evaluation check before the transmission of its TB. In this case, re-evaluation detects the possible collision with V_B , and V_A executes the resource replacement to avoid the detected collision. However, this collision was not actually going to occur since V_B does not finally transmit in the reserved resource due to the size or latency reselection. We refer to this event as an ineffective re-evaluation detection. An ineffective re-evaluation detection also occurs if V_B does not utilize its initially reserved resource, and V_A executes a resource replacement because it detected a possible collision with the initially reserved resource by V_B that would actually not have happened since V_B did finally not utilize its reserved resource.

We should note that ineffective re-evaluation detections do not affect resources reserved for retransmissions of a TB. This is the case because a vehicle always utilizes a reservation for a retransmission since it is adapted in size and latency to the TB, and there is no risk of size or latency reselection.

5.2.3 Resource replacement

During a re-evaluation, if a vehicle detects a potential collision it triggers the re-execution of step 2 of the resource reselection 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. During the resource replacement phase, a vehicle might select resources that are already occupied by neighboring vehicles and experience an undetectable collision on selected resources. 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 always be guaranteed, it is necessary to evaluate the actual effectiveness of the re-evaluation mechanism [17].

5.3 Impact of Re-evaluation 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 [17].

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 4.a which reports the different metrics for the two *RRI* selection strategies and all traffic densities evaluated. Table 4.a shows that the ratio of re-evaluation checks (ReCR⁹ 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 4.a). Table 4.a also shows that re-evaluation detects a large number of potential packet collisions (ReDR). However, Figure 10 shows that re-evaluation is not fully effective in avoiding collisions and improving the packet delivery ratio; this is independent of the *RRI* selection strategy. Figure 10 compares the performance when re-evaluation is implemented and when it is not. Figure 10(a) and Figure 10(b) plot the PDR for two traffic densities and Figure 10(c) the PCR for one of these densities [17].

⁹ Like SRR, LRR and URR metrics, ReCR does not vary with the vehicle density because it only depends on the traffic and on the reservations that each vehicle individually generates.

Table 4. Performance metrics (in %) when $N=1$

a) Aperiodic traffic of variable size and mixed traffic scenario

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

b) Periodic traffic of fixed size

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

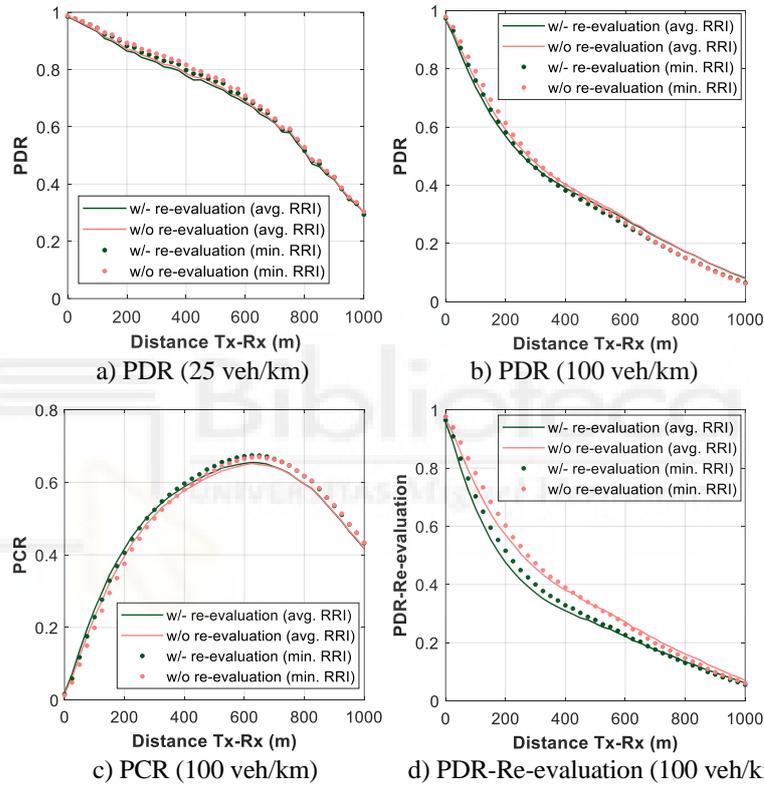


Figure 10. SPS performance in mixed traffic scenario with aperiodic traffic of variable size, $N = 1$ [17].

Figure 10 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 (see Section 5.2.2.1). The second reason is that packet variability can produce size and latency reselections and increase the probability of having to select new resources. 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 [17]. 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. The ineffectiveness of the re-evaluation mechanism is reflected in Table 4.a with the Ineffective Re-evaluation Detection Ratio (IReDR) metric that measures the 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 5.2.2.2). The obtained IReDR values negatively impact the PDR of the TBs for which at least a re-evaluation has been detected (PDR-Re-evaluation in Figure 10(d)). Figure 10(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 [17].

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 size and latency reselections nor unutilized reservations as it was the case with aperiodic traffic of variable size. Therefore, only TBs transmitted after the *Reselection Counter* depletes are eligible for a re-evaluation check since they are transmitted on selected resources (see Section 5.2.1). 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 4.b); similar trends are observed for the mixed traffic scenario. The low ReCR values result in the low ratio of re-evaluation detections (ReDR) reported in Table 4.b and the small impact of re-evaluation on the PCR in Figure 11(a), where the PCR is reported in the mixed traffic scenario for periodic traffic. However, Figure 11(b) shows significant gains in the PDR-Re-evaluation compared to the performance obtained if re-evaluations were not implemented. In this case, the IReDR is always equal to 0 (Table 4.b), and therefore re-evaluations are effective to avoid the limited set of packet collisions detected with periodic traffic of fixed size [17].

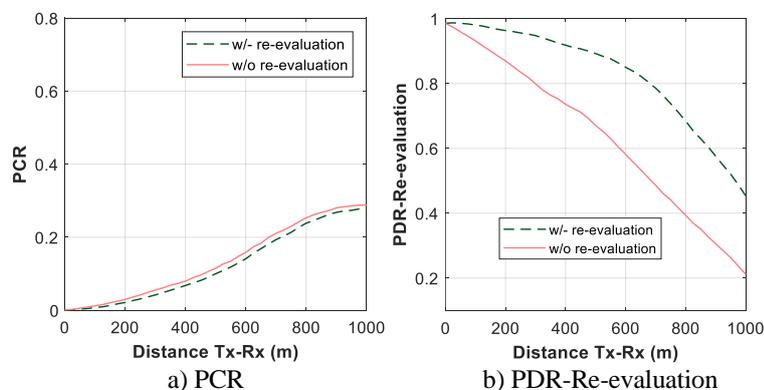


Figure 11. SPS performance in mixed traffic scenario for periodic traffic, 100 veh/km, $N = 1$. Similar trends are observed under the single traffic scenario [17].

5.4 Impact of Re-evaluation 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. Traffic variability can still impact the initial transmission of TBs when $N=2$ (see Section 5.2.2.1). 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. 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 more than 88% of the re-evaluation detections in the single traffic scenario (with vehicles generating medium intensity traffic) are triggered by reservations made for the retransmission of the same TB. Since re-evaluations are always effective in avoiding this collision, the PDR for both initial transmissions and retransmissions that detected re-evaluations (PDR-Re-evaluation) significantly outperforms the PDR without re-evaluation (Figure 12(a)); this was not the case without retransmissions ($N=1$) as shown in Figure 10(d). In the mixed traffic scenario, reported in Figure 12(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 a new TB. Reservations for a new TB do not always hold a transmission in the reserved resources and affect the effectiveness of the re-evaluation mechanism [17]. This explains the higher IReDR values in the mixed traffic scenario compared to the single traffic scenario (Table 5) as well as the lower positive impact of re-evaluation in Figure 12(b) compared to Figure 12(a).

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, Figure 13 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 undetectable collisions [17].

Table 5. Performance metrics (in %) 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

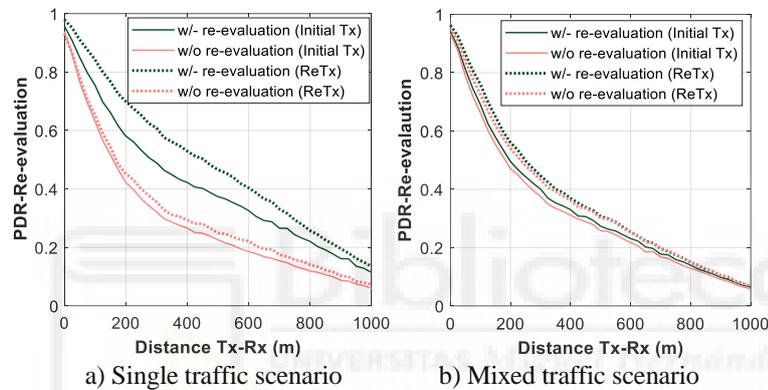


Figure 12. PDR-Re-evaluation experienced by SPS for aperiodic traffic of variable size when $N = 2$ (50 veh/km, average RRI strategy). Similar trends have been obtained for other densities [17].

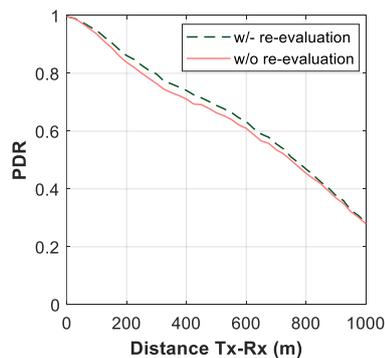


Figure 13. PDR experienced by SPS for aperiodic traffic of variable size when $N = 2$, (50 veh/km average RRI strategy) in the single traffic scenario. Similar trends have been obtained for other densities [17].

5.5 Re-evaluation strategies

The evaluation presented in Section 5.3 and Section 5.4 has considered that vehicles execute the re-evaluation check only in the slot that is mandatory in the standard (i.e. T_3 slots before the slot where the selected resources are located). This section describes the different re-evaluation check strategies that are then evaluated in Section 5.6.

The re-evaluation check must be executed before the slot where the initially selected resources are located (i.e. before slot m in Figure 14). In particular, the standard [6] indicates that the re-evaluation check can be executed at any slot from the slot n at which the packet is generated until the slot $m-T_3$ (see Section 2.4). The re-evaluation check can then be performed at any slot in the range $[n+1, m-T_3]$ in Figure 14. This range defines the different re-evaluation check strategies allowed by the 3GPP standard [6]. We evaluate in this chapter the two strategies analyzed by 3GPP for 5G NR V2X. In the first one, the vehicle only executes once a re-evaluation check at $m-T_3$. In the second strategy, the vehicle executes a re-evaluation check in all slots in the range $[n+1, m-T_3]$. The standard leaves up to UE implementation which re-evaluation check strategy to utilize. In the rest of the chapter, we refer to these strategies as one-slot and all-slots, respectively [18].

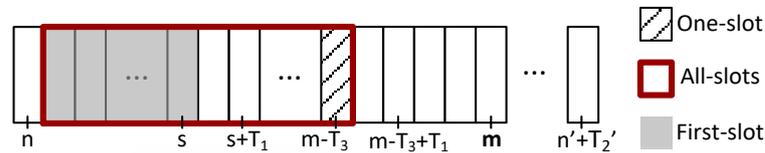


Figure 14. Illustration of re-evaluation check strategies (when T_3 and T_1 are equal to 5 and 2 slots, respectively) [18].

5.5.1 One-slot strategy

The one-slot strategy executes a single re-evaluation check before the initially selected resource. The check is executed at slot $n'=m-T_3$ (Figure 14), i.e. at the last possible moment within the range $[n+1, m-T_3]$. Delaying the re-evaluation check until $m-T_3$ increases the likelihood that the new selected resource after a resource replacement is located at a slot later than the initially selected resource, and this delays the transmission of the TB. If a vehicle selects a new resource at slot m' using the one-slot strategy, the standard forces the vehicle to execute a new re-evaluation check at slot $m'-T_3$. This process is repeated until it is not possible to define a new selection window (e.g., because $T_{2min}>T_2'$), or when the slot at which the vehicle should execute the re-evaluation has passed [18].

5.5.2 All-slots strategy

The standard also allows for the possibility to conduct re-evaluation checks in all slots in the range $[n+1, m-T_3]$. We refer to this strategy as all-slots (Figure 14). This strategy increases the computational cost compared to the one-slot strategy since it significantly augments the number of re-evaluation checks per TB. However, the all-slots strategy has

the potential to reduce the transmission latency compared to the one-slot strategy. This is the case because it is more likely with the all-slots strategy that a new selected resource (in the event of a resource replacement) anticipates the transmission of the TB with respect to the initially selected resource. Let's assume that a re-evaluation check at slot $n'=s$ within the range $[n+1, m-T_3]$ detects a potential collision in the initially selected resource allocated at the slot m (Figure 14). When a resource replacement process is triggered, a new selection window SW' is defined in the range of slots $[s+T_1, s+T_2']$ to select new resources. If s is any slot before $m-T_3$ (as represented in Figure 14), the new selection window SW' defined by the all-slots strategy starts earlier than the one defined with the one-slot strategy since $(s+T_1) < (m-T_3+T_1)$. This increases the likelihood that the new selected resources by the all-slots strategy advance the TB transmission compared to the one-slot strategy. A vehicle that has selected a new resource at slot m' with the all-slots strategy continues performing re-evaluation checks in every slot until the slot $m'-T_3$. This process is repeated iteratively until it is not possible to define a new selection window (e.g., because $T_{2min} > T_2'$), or when the slot at which the vehicle should execute the re-evaluation has passed [18].

5.5.3 First-slot strategy

The one-slot and all-slots re-evaluation check strategies are defined in 3GPP standards, and it is up to UE implementation to decide to use one of them [6]. A new re-evaluation check strategy is proposed with the objective to reduce the transmission latency while controlling the computational cost of the re-evaluation mechanism. This new proposal is referred to as the first-slot strategy. A vehicle implementing the first-slot strategy performs re-evaluation checks sequentially in all slots in the range $[n+1, m-T_3]$ until it detects the first re-evaluation. Figure 14 shows an example in which a vehicle implementing the first-slot strategy performs re-evaluation checks from the slot $n+1$ until the slot $n'=s$ where it detects a re-evaluation. It then executes the resource replacement process and selects new resources at slot m' . After this resource replacement, the vehicle does not perform any additional re-evaluation checks when implementing the first-slot strategy. Compared to the one-slot and all-slots strategies, the first-slot strategy does not execute a re-evaluation check T_3 slots before the selected resources if the vehicle detects a re-evaluation in a previous slot [18].

5.6 Impact of re-evaluation strategies

Results reported in this section are obtained in the single traffic scenario. In particular, vehicles generate aperiodic traffic of variable size, perform one transmission per TB, and use the minimum *RRI* selection strategy (see Section 3.2).

First, we analyze the impact of the re-evaluation check strategies on the end-to-end latency experienced by the packets that detect at least one re-evaluation. The end-to-end latency is measured from the time a vehicle generates a TB to the time it is received correctly. Figure 15 shows the cumulative distribution function (cdf) of this latency for the medium and high intensity scenarios when the vehicle density is 50 veh/km. We should note that the upper-bound of the latency (independently of the re-evaluation check strategy) corresponds to the PDB in each scenario (see Section 3.2). Figure 15 shows that the all-slots and first-slot strategies reduce considerably the latency with respect to the one-slot strategy. The one-slot strategy delays the start of the new selection window created after the re-evaluation detection on average 56 slots (or 28 ms) and 11 slots (or 5.5 ms) with respect to the slot where the packet is generated for the medium and high traffic scenarios, respectively. All-slots and first-slot only delay the start of the new selection window on average 3 slots (or 1.5 ms) and 4 slots (or 2 ms), respectively. Then, first-slot and all-slots can select resources that are located earlier than the ones selected by one-slot. This justifies the lower latencies reported in Figure 15 for all-slots and first-slot strategies with respect to one-slot [18].

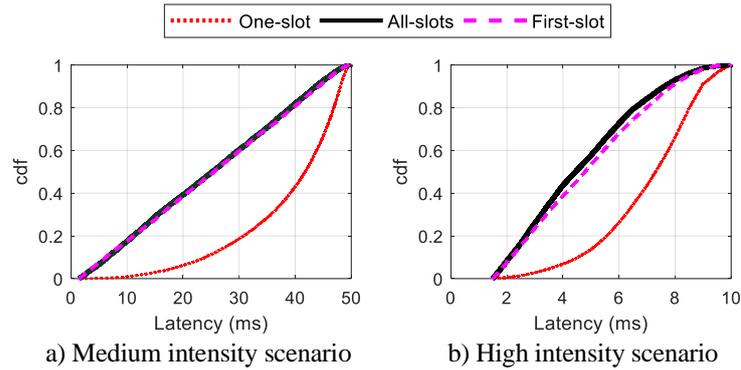


Figure 15. CDF of the end-to-end latency for 50 veh/km (similar trends observed for 100 veh/km) [18].

The latency differences observed in Figure 15 affect the packets that experience at least a re-evaluation detection. As a result, the impact of the re-evaluation check strategies on the NR V2X mode 2 latency increases with the ReDR ratio. Table 6 shows that the ratio of packets that experience at least a re-evaluation detection (ReDR) increases with the vehicular density and the traffic intensity.

The previous results show that all-slots and first-slot strategies significantly reduce the latency compared to the one-slot strategy. However, the one-slot strategy significantly reduces the computational cost compared to the all-slots strategy. This is visible in Table 6 that reports the number of re-evaluation checks that vehicles execute per packet (NReC). Table 6 also shows that the first-slot strategy results in NReC values very similar to those observed with the one-slot strategy [18]. In fact, the first-slot strategy reduces by ~98% and ~75% the average number NReC of re-evaluation checks compared to the all-slots strategy in the medium and high intensity scenarios, respectively (Table 6).

Table 6. Operation and Performance of the Re-evaluation Check Strategies

Scenario	PDB	Veh/km	Strategy	ReDR	NReC
Medium intensity	50 ms	50	One-slot	19 %	1
			All-slots	19 %	46
			First-slot	19 %	1
		100	One-slot	31 %	1
			All-slots	35 %	46.1
			First-slot	35 %	1.1
High intensity	10 ms	50	One-slot	31 %	1.04
			All-slots	38 %	6.4
			First-slot	38 %	1.5
		100	One-slot	30 %	1.04
			All-slots	37 %	6.5
			First-slot	37 %	1.6

The first-slot strategy can achieve the same low latency values as the all-slots strategy, while reducing the computational cost to the same level as the one-slot strategy. These gains are obtained without sacrificing the reliability of NR V2X mode 2. Actually, Figure 16.a shows that the three re-evaluation check strategies result in nearly identical PDRs for the packets transmitted after a re-evaluation detection. The resources selected after a re-evaluation detection and resource replacement might impact the transmissions of other vehicles. However, Figure 16.b shows that three strategies result again in nearly the same PDR for all transmitted packets [18].

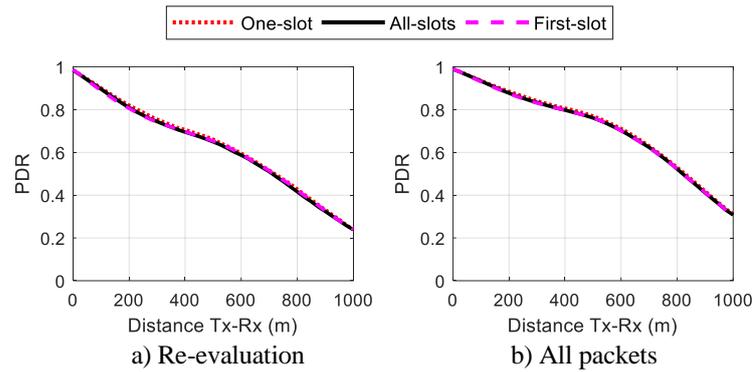


Figure 16. PDR of packets that experienced at least a re-evaluation detection (a) and PDR experienced by all packets (b). Medium intensity with 50 veh/km (same trends observed in the other evaluated scenarios) [18].

5.7 Conclusions and additional results

This chapter has presented an in-depth analysis of the impact of the re-evaluation mechanism on the operation and performance of NR V2X mode 2 sidelink communications. The re-evaluation mechanism has been introduced in 3GPP Release 16 standards to reduce packet collisions. This chapter shows that the effectiveness of re-evaluation to avoid collisions depends on the data traffic patterns and mode 2 configurations. In particular, the study shows that re-evaluation is effective in detecting collisions when vehicles transmit periodic traffic of fixed size. However, the impact on the performance of NR V2X mode 2 is small since the number of packet collisions detected by re-evaluation is low under periodic traffic of fixed size. The effectiveness of re-evaluation can decrease under the presence of aperiodic traffic of variable size because traffic variability increases the probability of selecting new resources, and re-evaluation cannot detect collisions on new selected resources. This is particularly the case when there are no retransmissions. With retransmissions, re-evaluation is more effective in detecting packet collisions, even with aperiodic traffic of variable size. However, the impact of re-evaluation on the performance with retransmissions is low since, without re-evaluation, a TB is correctly received if just one of the two transmissions is correctly received [17].

This chapter also has presented and evaluated a set of re-evaluation check strategies for NR V2X mode 2. The 3GPP standard defines two strategies (one-slot and all-slots) for vehicles to decide when to perform a re-evaluation check and how many times to execute a re-evaluation check. The conducted evaluation has shown that the all-slots strategy can significantly reduce the transmission latency at the expense of a high computational cost compared to the one-slot strategy. This study presents an alternative re-evaluation check

strategy (first-slot) that reduces the latency to values similar to the all-slots strategy with a low computational cost similar to the one of one-slot strategy. These benefits are achieved without degrading the reliability of the NR V2X mode 2 transmissions compared to the standardized re-evaluation check strategies [18].

Sections 5.2, 5.3 and 5.4 of this chapter have analyzed the impact of re-evaluation and its effectiveness in detecting and avoiding packet collisions in NR V2X mode 2. Our publication related to these sections is [17] (found in Annex A.3). This publication includes the following additional details and results:

- Detailed description of the different cases where re-evaluation can detect a collision with one transmission per TB (with single RRI and different RRIs) and with two transmissions per TB.
- Comparative analysis of the PDR of SPS with aperiodic traffic of variable size and periodic traffic of fixed size to show the impact of undetectable collisions generated by the aperiodic traffic of variable size that are not present with periodic traffic of fixed size.
- Computation of Half-Duplex Losses Ratio (HDLR) and Propagation Losses Ratio (PLR) metrics. HDLR represents the fraction of TBs that are incorrectly received because of the half-duplex limitation. HDLR computes the errors that occur when a TB cannot be received because the receiver was transmitting in the same slot. PLR represents the fraction of TBs that cannot be correctly decoded because the received power level is below the sensitivity level or the Signal to Noise Ratio (SNR) is not sufficiently high. HDLR and PLR metrics are used to justify that the differences observed in the performance of NR V2X mode 2 with and without re-evaluation are exclusively due to the re-evaluation mechanism.
- Detailed analysis of the number of TBs affected by a persistent collision that can be avoided by the re-evaluation mechanism under periodic traffic of fixed size.
- Apart from the analysis focused on SPS that has been summarized in this chapter, our related publication also includes an analysis of the effectiveness of re-evaluation to detect and avoid packet collisions when NR V2X mode 2 operates with the Dynamic Scheduling (DS) scheme. The DS scheme can only generate reservations for retransmissions (see Section 2.3) and the re-evaluation mechanism needs the presence of reservations to detect and avoid packet collisions (see Section

5.2.2.1). For this reason, the effectiveness of the re-evaluation mechanism with the DS scheme has been evaluated only with retransmissions. In this case, the re-evaluations are effective with DS under aperiodic traffic of variable size. This is the case because only re-evaluation detections triggered by a reservation for a retransmission are possible under DS, and these re-evaluation detections are always effective. However, like with SPS, the impact of the re-evaluation mechanism on the PDR is low because without re-evaluation, a TB is correctly received if just one of the two transmissions is correctly received.

Sections 5.5 and 5.6 of this chapter have evaluated the impact of different re-evaluation check strategies on NR V2X mode 2. Our publication related to these sections is [18] (found in Annex A.4). This publication includes the following additional results:

- Computation of the Multiple Re-evaluation Detection Ratio (MReDR) metric, which measures the ratio of transmitted TBs experiencing more than one re-evaluation detection. MReDR demonstrates that the first-slot design ensures no TB experiences more than one re-evaluation detection. This contrasts with the standardized strategies, which allow for multiple re-evaluation detections.

6 Selective Re-evaluation Mechanism

6.1 Introduction

The analysis in Chapter 5 has shown that re-evaluation is not fully effective in avoiding packet collisions generated by aperiodic traffic of variable size because many of the detected collisions ultimately do not happen (due to size or latency reselections and unutilized reservations), and selecting new resources with the resource replacement increases the risk of collisions undetectable by re-evaluation. To address this inefficiency, this chapter proposes a selective re-evaluation mechanism that only selects new resources when the vehicle is certain that a detected collision is going to occur, which is the case when a re-evaluation detection is triggered by a reservation for a retransmission of a TB. The proposal builds from a thorough analysis and characterization of the challenges faced by the standard re-evaluation mechanism to avoid packet collisions when dealing with aperiodic packets of variable size. This chapter also focuses on SPS because aperiodic traffic of variable size mostly impacts SPS, since the DS scheme selects new resources for each TB (see Section 4.2). The effectiveness of re-evaluation is then particularly relevant when utilizing the SPS scheme. This chapter shows that the proposed selective re-evaluation mechanism improves the reliability and latency of 5G NR V2X mode 2 communications.

The chapter is organized as follows. Section 6.2 provides an overview of the effectiveness of the standard re-evaluation mechanism. Section 6.3 analytically quantifies the effectiveness of the standard re-evaluation mechanism. Section 6.4 describes the proposed selective re-evaluation mechanism. Section 6.5 details the evaluation of 5G NR V2X mode 2 using both the standard and selective re-evaluation mechanisms. Lastly, Section 6.6 summarizes the main outcomes of this chapter. It should be noted that this chapter is related to our submitted journal paper [19].

6.2 Effective and ineffective re-evaluations

A collision detected by re-evaluation that ultimately does not happen is referred to as ineffective re-evaluation detection (see Section 5.2.2.2). An ineffective re-evaluation

detection can occur when a re-evaluation detection is triggered by a reservation for a new TB. This is because reservations for a new TB may not be utilized due to size or latency reselections and unutilized reservations produced by the new TB. In the case of an ineffective re-evaluation detection, a vehicle detects a collision that ultimately does not happen and executes a resource replacement to select new resources, increasing the probability of an undetectable collision (see Section 5.2.2.1). Ineffective re-evaluation detections do not affect resources reserved for retransmissions of a TB. This is because a vehicle always utilizes a reservation for a retransmission since it is adapted in size and latency to the TB. In this case, a resource replacement triggered by a reservation for a retransmission will always avoid the detected collision. For the sake of brevity, ineffective re-evaluation detections are referred to as ineffective re-evaluations in the rest of this chapter.

Ineffective re-evaluations can also negatively impact the latency experienced since a resource replacement increases the probability to transmit a TB at a later slot. This is the case because the selection window SW' used to select a new resource after a resource replacement starts later than the original selection window SW (see Section 2.4). The potential negative impact of ineffective re-evaluations on latency increases if all transmissions of a TB undergo a resource replacement. In this case, the resource for the first transmission following a resource replacement can be located at any slot of the SW' , and all the following new resources for the retransmissions of a TB must be after the resource for the first transmission considering the SCI limitations (see Section 2.4). If only a subset of transmissions of a TB are affected by a resource replacement, the selection of these resources is limited to 31 slots after the resources that did not experience a re-evaluation detection, which reduces the risk of increasing the latency. The potential negative impact of ineffective re-evaluations on latency also depends on which transmissions are affected by the replacement. Let's suppose that a vehicle originally selects two resources for the initial transmission and the retransmission of a TB ($N=2$), and the resources for the initial transmission are located at slot m_1 . The vehicle executes a re-evaluation check¹⁰ at m_1-T_3 , and SW' is defined by the range of slots $[m_1-T_3+T_1, m_1-T_3+T_2]$. The resource replacement can decrease the transmission latency if the new resource is

¹⁰ It should be noted that vehicles execute the re-evaluation check T_3 slots before the selected resources in this chapter because this is the re-evaluation check that is mandatory according to the standard (see Section 2.4).

selected in the range $[m_1 - T_3 + T_1, m_1 - 1]$. However, this is highly unlikely since the length of this range is, for example, only 3 slots if we consider $T_3 = 5$ slots and $T_1 = 2$ slots with SCS = 30 kHz. The probability to decrease the latency is higher for a resource replacement over a retransmission of the TB. Let's suppose that this retransmission was originally planned at slot m_2 ($m_2 > m_1$). SW' is also defined by the range of slots $[m_1 - T_3 + T_1, m_1 - T_3 + T_2']$ in this case, and the latency is reduced if the new resource is selected in the range $[m_1 - T_3 + T_1, m_2 - 1]$. This range has $m_2 - m_1$ more slots than the range over which the initial transmission of a TB can reduce the latency following a resource replacement (i.e. $[m_1 - T_3 + T_1, m_1 - 1]$).

6.3 Effectiveness of re-evaluation

The effectiveness of re-evaluation strongly depends on the probability of not utilizing a reservation for transmitting a new TB. If this probability is higher than 0.5, there will be more ineffective re-evaluations than effective ones, and resource replacements following a re-evaluation detection are more likely to degrade the reliability than improve it since the selection of new resources increases the probability of packet collisions. To understand whether re-evaluation can be effective in avoiding packet collisions, it is necessary to quantify the probability of not utilizing a reservation for a new TB (P_{NUT}). This section analytically quantifies the probability of not utilizing a reservation for transmitting a new TB (P_{NUT}). We first calculate it when vehicles do one transmission per TB and then extend it to the case of retransmissions, and in particular to the case of two transmissions per TB (the initial one and a retransmission). We should note that two transmissions per TB is the configuration recommended by SAE [14], and is also the mode 2 configuration used in our evaluation. The probability is quantified considering the scenario where vehicles generate the TBs aperiodically since this is the scenario where re-evaluation is more necessary due to the higher packet collisions (see Chapter 4). Following the 3GPP aperiodic traffic model [12], we consider that vehicles generate TBs with an inter-packet arrival time $\tau = \bar{x} + x(\bar{x})$, where \bar{x} is a constant minimum inter-packet arrival time and x is an exponential random variable of mean equal to \bar{x} .

6.3.1 One transmission per TB

A reservation to transmit a new TB is not utilized when there is a latency reselection, a size reselection or an unutilized reservation because the vehicle has generated the new TB after the reserved resource. We can then define the probability of not utilizing a reservation as:

$$P_{NUT} = P(LRR \cup SR \cup URR) \quad (1)$$

where $P(y)$ is the probability of y , LRR is a latency reselection, SR is a size reselection without a simultaneous latency reselection, and URR is an unutilized reservation. A TB can simultaneously generate a size and latency reselection. URR does not include possible unutilized reservations caused by a latency or size reselection triggered by the TB that is generated after the reserved resource since these are already included in LRR and SR . LRR , SR and URR are then exclusive events, and P_{NUT} can be expressed as [35]:

$$P_{NUT} = P_{LRR} + P_{SR} + P_{URR} \quad (2)$$

where P_{LRR} , P_{SR} and P_{URR} are the probabilities of LRR , SR and URR respectively. To compute these probabilities, we define n as the slot where the vehicle generates a TB and selects new resources, and m as the slot where the new selected resource is located. We define \bar{s}_m as the slot where the new selected resource is located on average with respect to the slot n . \bar{s}_m depends on the length of the selection window (and hence on T_2 and T_1) and can be computed as:

$$\bar{s}_m = \bar{m} - \bar{n} = \frac{\sum_{i=T_1}^{T_2} i}{T_2 - T_1 + 1} \quad (3)$$

We define \bar{s}_r as the slot where the resource reserved for a new TB is located on average, and n_2 as the slot where a new TB is generated. \bar{s}_r is equal to the sum of \bar{s}_m and the RRI (expressed in number of slots). We define \bar{s}_s as the slot where the selected resource available after an unutilized reservation is located on average. \bar{s}_s is equal to the sum of \bar{s}_m and $2 * RRI$. $n_2 - n$ is the inter-packet arrival time τ (in number of slots):

$$n_2 - n = \tau = (\bar{x} + x(\bar{x})) * 2^\mu \quad (4)$$

where x is an exponential random variable of mean \bar{x} ms and μ represents the numerology configured¹¹.

P_{LRR} is the sum of the probability P_{LRR1} that the vehicle reserves a resource that does not fulfill the PDB (Packet Delay Budget) of the new TB, and the probability P_{LRR2} that the vehicle generates a TB after the reserved resource and the following selected resource (located RRI slots after) does not fulfill the PDB of the new TB. P_{LRR1} can be expressed as:

$$P_{LRR1} = P((\tau + PDB) < \bar{s}_r) = cdf_{\tau}(\bar{s}_r - PDB) \quad (5)$$

where $cdf_{\tau}(y)$ is the cumulative distribution function of τ evaluated at y . P_{LRR2} can be expressed as:

$$\begin{aligned} P_{LRR2} &= P((\tau > \bar{s}_r) \cap ((\tau + PDB) < \bar{s}_s)) \\ &= cdf_{\tau}(\bar{s}_s - PDB) - cdf_{\tau}(\bar{s}_r) \end{aligned} \quad (6)$$

P_{SR} represents the probability of a size reselection without simultaneous latency reselections. It can be estimated as:

$$P_{SR} = P_{SRR} * (1 - P_{LRR}) \quad (7)$$

where P_{SRR} is the probability of a size reselection with possible simultaneous latency reselections. P_{SRR} can be computed as:

$$P_{SRR} = \sum_{i=\min(i)}^{\max(i)} P_{SRR}(i) \quad (8)$$

where $P_{SRR}(i)$ is the probability of size reselection when the vehicle is using i sub-channels. A size reselection happens when the vehicle generates a TB that needs j sub-channels to be transmitted and $j > i$. $P_{SRR}(i)$ can then be computed as the multiplication of the probability that the vehicle uses a resource of i sub-channels ($P_{RES}(i)$) and the probability that it generates a TB that needs $j > i$ sub-channels ($P_G(j > i)$).

$$P_{SRR}(i) = P_{RES}(i) * P_G(j > i) \quad (9)$$

$P_{RES}(i)$ can be computed as the sum of the following probabilities:

$$P_{RES}(i) = P_{LRR}(i) + P_{RC>0}(i) + P_{RC=0}(i) + P_{SR}(i) \quad (10)$$

¹¹ In the case of periodic traffic, τ should be set equal to the traffic periodicity multiplied by 2^{μ} , where μ is the numerology.

where $P_{LRR}(i)$ is the probability of using a resource of i sub-channels after a latency reselection. $P_{SR}(i)$ is the probability of using a resource of i sub-channels after a size reselection without a simultaneous latency reselection. $P_{RC>0}(i)$ is the probability of using the same resource of i sub-channels when the *Reselection Counter* (RC) is higher than 0. $P_{RC=0}(i)$ is the probability of using a resource of i sub-channels when the RC is equal to 0. $P_{LRR}(i)$ is computed as a function of the probability of latency reselection (P_{LRR}) and the probability of generating a TB that needs i sub-channels ($P_G(i)$):

$$P_{LRR}(i) = P_{LRR} * P_G(i) \quad (11)$$

The probability $P_{RC>0}(i)$ is estimated as:

$$P_{RC>0}(i) = \overline{P_{RC}}(i) * P_{GNL}(j \leq i) \quad (12)$$

where $P_{GNL}(j \leq i)$ is the probability of generating a TB that needs $j \leq i$ sub-channels to be transmitted and does not produce a latency reselection. It can be expressed as:

$$P_{GNL}(j \leq i) = (1 - P_{LRR}) * \sum_{j=\min(i)}^i P_G(j) \quad (13)$$

$\overline{P_{RC}}(i)$ is the average of the probabilities of using the same resource of i sub-channels to transmit from 1 to $RC-1$ consecutive TBs. It is computed as:

$$\overline{P_{RC}}(i) = P_G(i) * \frac{\sum_{k=1}^{\overline{RC}-2} P_{RC}(i, k)}{\overline{RC} - 2} \quad (14)$$

where \overline{RC} is the average value of the RC and $P_{RC}(i, k)$ is the probability of using the same resource of i sub-channels to transmit k consecutive TBs. Eq. (14) limits the number of consecutive TBs to $\overline{RC}-2$ TBs so that the RC is always higher than 0. A vehicle uses the same resource of i sub-channels for k consecutive TBs if the TBs fit in the same resource of i sub-channels and the k TBs do not trigger any latency reselection. $P_{RC}(i, k)$ can then be expressed as:

$$P_{RC}(i, k) = (P_{GNL}(j \leq i))^k \quad (15)$$

When RC is equal to 0, a vehicle selects a new resource with probability $(1-P_k)$. $P_{RC=0}(i)$ is equal to:

$$P_{RC=0}(i) = P_{RCN}(i) + P_{RCK}(i) \quad (16)$$

where $P_{RCN}(i)$ and $P_{RCK}(i)$ are the probabilities of using a new resource of i sub-channels or keeping the same resource of i sub-channels when the RC is equal to 0. $P_{RCK}(i)$ is computed as:

$$P_{RCk}(i) = P_{kp}(i) * P_k * P_{GNL}(j \leq i) \quad (17)$$

where $P_{kp}(i)$ is the probability of using the same resource of i sub-channels during RC consecutive TBs. This can happen when the new selected resource needs i sub-channels to transmit the first TB, and the same resource is used for the following $RC-1$ TBs. $P_{kp}(i)$ is then equal to:

$$P_{kp}(i) = P_G(i) * P_{RC}(i, \overline{RC} - 1) \quad (18)$$

A vehicle selects a new resource of i sub-channels after RC is depleted if the first TB needs i sub-channels to be transmitted. The probability $P_{RCN}(i)$ of using a new resource of i sub-channels can be computed as:

$$P_{RCN}(i) = \overline{P_{kp}} * (1 - P_k) * P_G(i) \quad (19)$$

where $\overline{P_{kp}}$ is the probability of using the same resource of any number of sub-channels during RC consecutive TBs. $\overline{P_{kp}}$ is estimated as:

$$\overline{P_{kp}} = \sum_{j=\min(i)}^{\max(i)} P_G(j) * P_{kp}(j) \quad (20)$$

For brevity, we define the variable $P_X(i)$ as:

$$P_X(i) = P_{LRR}(i) + P_{RC>0}(i) + P_{RC=0}(i) \quad (21)$$

To compute $P_{RES}(i)$, we finally need to derive the probability $P_{SR}(i)$ that a vehicle uses i sub-channels after a size reselection without a simultaneous latency reselection:

$$P_{SR}(i) = P_{SR} * P_G(i) \quad (22)$$

Using eq. (7), we can express $P_{SR}(i)$ as:

$$P_{SR}(i) = P_{SRR} * (1 - P_{LRR}) * P_G(i) \quad (23)$$

Using eqs. (8), (9), (10), (21), (23) we can express P_{SRR} as:

$$P_{SRR} = \sum_{i=\min(i)}^{\max(i)} (P_G(j > i) * (P_X(i) + P_{SRR} * (1 - P_{LRR}) * P_G(i))) \quad (24)$$

We can then derive P_{SRR} as:

$$P_{SRR} = \frac{\sum_{i=\min(i)}^{\max(i)} (P_G(j > i) * P_X(i))}{1 - (1 - P_{LRR}) * \sum_{i=\min(i)}^{\max(i)} (P_G(j > i) * P_G(i))} \quad (25)$$

Using P_{SRR} and P_{LRR} , we can derive $P_{SR}(i)$ using eq. (23) and P_{SR} using eq. (7). P_{SR} can then be finally expressed as:

$$P_{SR} = \frac{(1 - P_{LRR}) * \sum_{i=\min(i)}^{\max(i)} (P_G(j > i) * P_X(i))}{1 - (1 - P_{LRR}) * \sum_{i=\min(i)}^{\max(i)} (P_G(j > i) * P_G(i))} \quad (26)$$

To compute P_{NUT} , we finally need to derive P_{URR} that represents the probability of an unutilized reservation because the TB is generated after the reserved resource. P_{URR} does not account for the cases where the TB generated after the reserved resource triggers a latency reselection or a size reselection. It can then be expressed as:

$$P_{URR} = P_{URS} * (1 - P_{SR}) \quad (27)$$

where P_{URS} is the probability of generating a TB after the reserved resource that does not produce a latency reselection and $(1 - P_{SR})$ represents the probability that a TB does not produce a size reselection. P_{URS} is defined as:

$$\begin{aligned} P_{URS} &= P((\tau > \bar{s}_r) \cap ((\tau + PDB) > \bar{s}_s)) \\ &= 1 - cdf_{\tau}(\max(\bar{s}_r, \bar{s}_s) - PDB) \end{aligned} \quad (28)$$

6.3.2 Two transmissions per TB

In the case that a vehicle does two transmissions per TB (the initial transmission and a retransmission), the resources reserved for a new TB are also not utilized if there is a latency reselection, a size reselection or an unutilized reservation because the TB is generated after the reserved resources. Compared to the scenario with one transmission per TB, the retransmission of a TB modifies when these three events happen and therefore their probabilities that we denote now P_{LRR}^R , P_{SR}^R and P_{URR}^R . The probability P_{NUT}^R of not utilizing a reservation when a vehicle does two transmissions per TB is the sum of P_{LRR}^R , P_{SR}^R and P_{URR}^R .

The process followed to compute P_{NUT}^R is similar to the one followed when each TB is transmitted only once. We define m_1 and m_2 as the slots where the new selected resources for the initial transmission and the retransmission are located respectively. We define \bar{s}_{m1} and \bar{s}_{m2} as the slots where the new selected resources for the initial transmission and retransmission are located on average, respectively, with respect to the slot n where a vehicle generates a new TB and selects new resources. The first selected resource is chosen randomly (Section 2.3) like for the scenario where each TB is transmitted only once; it is

hence equal on average to $\overline{s_m}$ (eq. (3), Section 6.3.1). Let's suppose that this first selected resource is located at slot m_α . The second selected resource must be located within the window $[m_\alpha-31, m_\alpha+31]$. We can then estimate $\overline{s_{m1}}$ and $\overline{s_{m2}}$ as:

$$\overline{s_{m1}} = \overline{m_1 - n} = \overline{s_m} - \frac{\overline{\Delta_{m1m2}}}{2} \quad (29)$$

$$\overline{s_{m2}} = \overline{m_2 - n} = \overline{s_m} + \frac{\overline{\Delta_{m1m2}}}{2} \quad (30)$$

where $\overline{\Delta_{m1m2}}$ is the average gap in number of slots between m_1 and m_2 . It is computed as:

$$\overline{\Delta_{m1m2}} = \frac{\max(\Delta_{m1m2}) + \min(\Delta_{m1m2})}{2} \quad (31)$$

where $\max(\Delta_{m1m2})$ and $\min(\Delta_{m1m2})$ are the maximum and minimum gap between m_1 and m_2 . The minimum gap is 1 slot because the initial transmission and the retransmission must be located in different slots. The maximum gap is 31 slots so that both transmissions are within a window of 32 slots (Section 2.3).

P_{LRR}^R can also be defined as the sum of the probability P_{LRR1}^R that the vehicle reserves a resource that does not fulfill the PDB of the new TB, and the probability P_{LRR2}^R that the vehicle generates a TB after the reserved resource and the following selected resource does not fulfill the PDB of the new TB. We define $\overline{s_{r1}}$ and $\overline{s_{r2}}$ as the slots where the resources reserved for the initial transmission and retransmission of a new TB are located on average, respectively. $\overline{s_{s2}}$ is the slot where the selected resource for the retransmission of the TB is located on average after the vehicle does not utilize the reserved resources. These resources are not utilized in the case of a latency reselection, and can be computed as:

$$\overline{s_{r1}} = \overline{s_{m1}} + RRI \quad (32)$$

$$\overline{s_{r2}} = \overline{s_{m2}} + RRI \quad (33)$$

$$\overline{s_{s2}} = \overline{s_{m2}} + 2 * RRI \quad (34)$$

To compute P_{LRR1}^R and P_{LRR2}^R , we should take into account that if the resource for the retransmission of the TB does not fulfill the PDB of the new TB, then the vehicle reselects the resources for the initial transmission and the retransmission [6]. P_{LRR1}^R can then be calculated as P_{LRR1} (eq. (5)) but replacing $\overline{s_r}$ by $\overline{s_{r2}}$. To compute P_{LRR2}^R , we should also note that if the vehicle generates the TB after the resource reserved for the initial transmission, then the vehicle uses the next selected resources available for the initial transmission and the retransmission [6] even if the resource reserved for the retransmission

is located after the generation of the TB. P_{LRR2}^R can then be calculated as P_{LRR2} (eq. (6)) but replacing \bar{s}_r by \bar{s}_{r1} and \bar{s}_s by \bar{s}_{s2} .

The probability P_{SR}^R of size reselection without a simultaneous latency reselection is estimated as:

$$P_{SR}^R = P_{SRR}^R * (1 - P_{LRR}^R) \quad (35)$$

P_{SRR}^R is computed as P_{SRR} (Section 6.3.1) but replacing P_{LRR} by P_{LRR}^R in eqs. (11) and (13). Similarly, P_{URR}^R is calculated as P_{URR} (Section 6.3.1) but replacing P_{SR} by P_{SR}^R in eq. (27), as well as \bar{s}_r by \bar{s}_{r1} and \bar{s}_s by \bar{s}_{s2} in eq. (28).

6.3.3 Numerical evaluation

The effectiveness of re-evaluation depends on the probability of not utilizing a reservation for a new TB. Figure 17 and Figure 18 plot this probability with respect to the average inter-packet arrival time (following the 3GPP guidelines [12]) for one and two transmissions per TB, respectively, when vehicles generate aperiodic traffic of variable size. The figure also plots the probabilities of latency reselection (LRR), size reselection (SR) and unutilized reservation because the vehicle has generated the TB after the reserved resource (URR). The results are derived for two strategies for the selection of the RRI (Section 3.2): average RRI (RRI is set equal to the average inter-packet arrival packet time) and minimum RRI (RRI is set equal to the minimum of the inter-packet arrival time). The PDB is set equal to the minimum inter-packet arrival time following the 3GPP guidelines [12].

The figures show that the probability of not utilizing a reservation for a new TB is higher than 0.5 in all evaluated scenarios when considering aperiodic traffic of variable size. In this case, a resource replacement triggered by re-evaluation is more likely to degrade the reliability than improve it because the detected collision might not happen and selecting new resources increases the probability of packet collision (Section 6.2). In line with the trends reported in Chapter 4, the figures show that the minimum RRI strategy reduces the probability of latency reselections at the cost of increasing the probability of unutilized reservations. On the other hand, the average RRI strategy reduces the probability of unutilized reservations but increases the probability of latency reselections. The probability of size reselection (without simultaneous latency reselections) is low in all scenarios.

Figure 17 shows that the probability of not utilizing a reservation for a new TB is maintained constant with respect to the average inter-packet arrival time for one transmission per TB. The variations in the generation of the traffic and PDB are compensated with the adaptation of the RRI to the inter-packet arrival time. This trend is not observed for two transmissions per TB (Figure 18) because the probability of latency reselections is higher than when there is one transmission per TB. This increase cannot be fully compensated with the adaptation of the RRI to the inter-packet arrival time, and the impact is higher with more stringent latency requirements, i.e. with low PDBs (corresponding to low inter-packet arrival times).

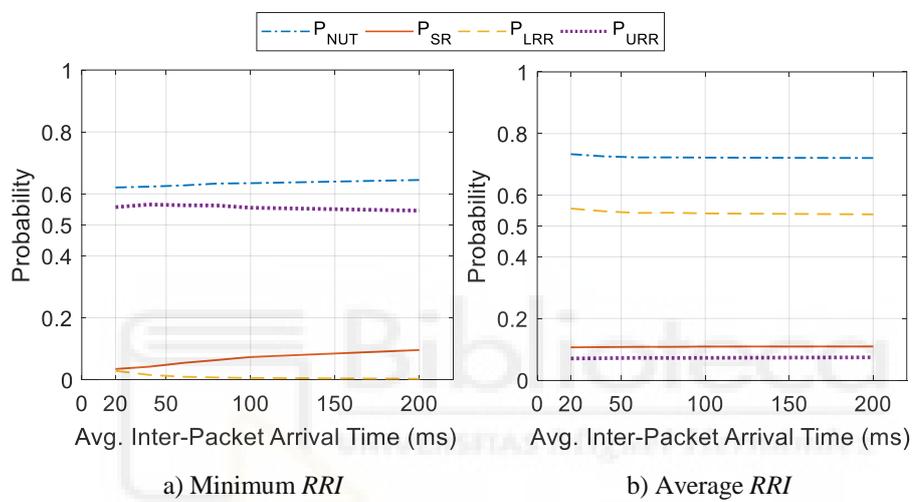


Figure 17. Probability of not utilizing a reservation for a new TB (P_{NUT}) for one transmission per TB [19].

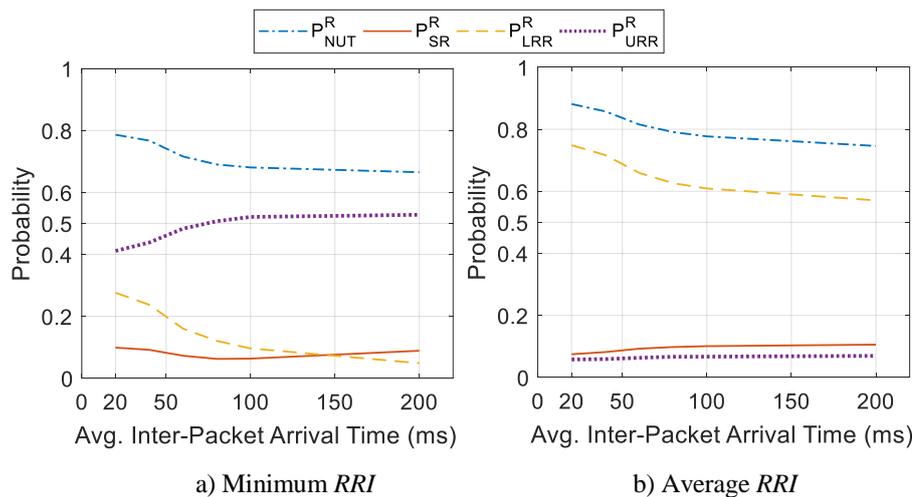


Figure 18. Probability of not utilizing a reservation for a new TB (P_{NUT}^R) for two transmissions per TB [19].

6.4 Selective re-evaluation mechanism

The previous section has demonstrated that the standard re-evaluation mechanism can produce more ineffective re-evaluations than effective ones when vehicles generate aperiodic traffic of variable size. In fact, the previous section has shown that, for all the evaluated scenarios, a resource replacement triggered by the standard re-evaluation mechanism is more likely to degrade the reliability than improve it (i.e. $P_{NUT} > 0.5$). To address this challenge, we propose a simple, yet effective, selective re-evaluation mechanism that modifies the standard re-evaluation mechanism to avoid ineffective re-evaluations.

The proposal only executes a resource replacement following a re-evaluation detection if the replacement is estimated to be effective, i.e. it is certain it will avoid a collision. Effective resource replacements are those generated by a re-evaluation detection triggered by a reservation for a retransmission of a TB. This is because a vehicle always utilizes a reservation for a retransmission since it is adapted in size and latency to the TB, and there is no risk of size or latency reselection. In this case, a replacement triggered by a reservation for a retransmission will always avoid a detected collision. Our proposal is therefore applicable only when NR V2X mode 2 is configured with retransmissions. However, we should remind that two transmissions per TB (i.e. one initial transmission and one retransmission of the TB) is the configuration recommended by SAE [14], and also the mode 2 configuration used in this study.

On the other hand, it is not possible to guarantee that a resource replacement following a re-evaluation detection triggered by a reservation for a new TB will be effective and will for sure avoid a detected collision (Section 6.3.3). The replacement might be ineffective if the new TB generates a size or latency reselection, and the colliding vehicle does not ultimately utilize the reservation announced for the new TB; these circumstances are more likely with aperiodic traffic of variable size. To avoid ineffective re-evaluations, the selective re-evaluation proposal does not execute the resource replacements following re-evaluation detections triggered by a reservation for a new TB. To identify effective resource replacements, a vehicle analyzes the information received in SCIs to identify whether a reservation by another vehicle that triggers a re-evaluation detection is for a retransmission of a TB or for the transmission of a new TB. We should note that a vehicle

informs in its SCI whether a reservation is for the retransmission of a TB or for the transmission of a new TB [36].

The proposed selective re-evaluation mechanism can also reduce the probability of undetectable collisions by not executing potentially ineffective resource replacements. This is because a resource replacement results in the selection of new resources, and it is not possible to detect collisions when vehicles are selecting new resources (no reservation has yet been announced). In addition, the vehicle that avoids the resource replacement maintains its selected resource and hence reduces the probability that other vehicles select the same resource as they can detect it is occupied. Finally, we should note that the selective re-evaluation proposal can also reduce the latency since resource replacements increase the probability of high latencies (Section 6.2).

6.5 Evaluation

This section evaluates and compares the performance of the standard and selective re-evaluation mechanisms. To do so, the following scenarios and configuration of parameters are employed: single traffic scenario (see Section 3.2), vehicles select 2 resources in step 2 (i.e. N is equal to 2) following the configuration recommended by SAE [14], the RRI is selected following the minimum RRI strategy (see Section 3.2), and vehicles execute the re-evaluation check in the mandatory slot according to the standard, i.e. T_3 slots before the selected resources (see Section 2.4).

Figure 19 compares the Packet Delivery Ratio (PDR) of the standard and selective re-evaluation schemes with aperiodic traffic of variable size. We should note that a TB is successfully decoded if at least 1 out of the N transmissions of a TB is successfully decoded. The comparison is done only for the transmissions that have experienced at least one re-evaluation detection (Figure 19-left, PDR–Re-evaluation) and for all the transmissions (Figure 19-right, PDR). Figure 19 reports results for the low, medium and high intensity scenarios and densities of 50 and 100 veh/km. Figure 19-left clearly shows that the selective re-evaluation mechanism improves the effectiveness of the re-evaluation process and hence the reliability under all the studied conditions. This is because the selective re-evaluation mechanism executes only the replacements that it is certain will result in a packet collision, i.e. those where the re-evaluation detection is triggered by a reservation for a retransmission. In contrast, the standard re-evaluation mechanism always

executes a resource replacement after a re-evaluation detection, and Figure 19-left shows that this does not increase the reliability due to the ineffective re-evaluations discussed in Section 6.2.

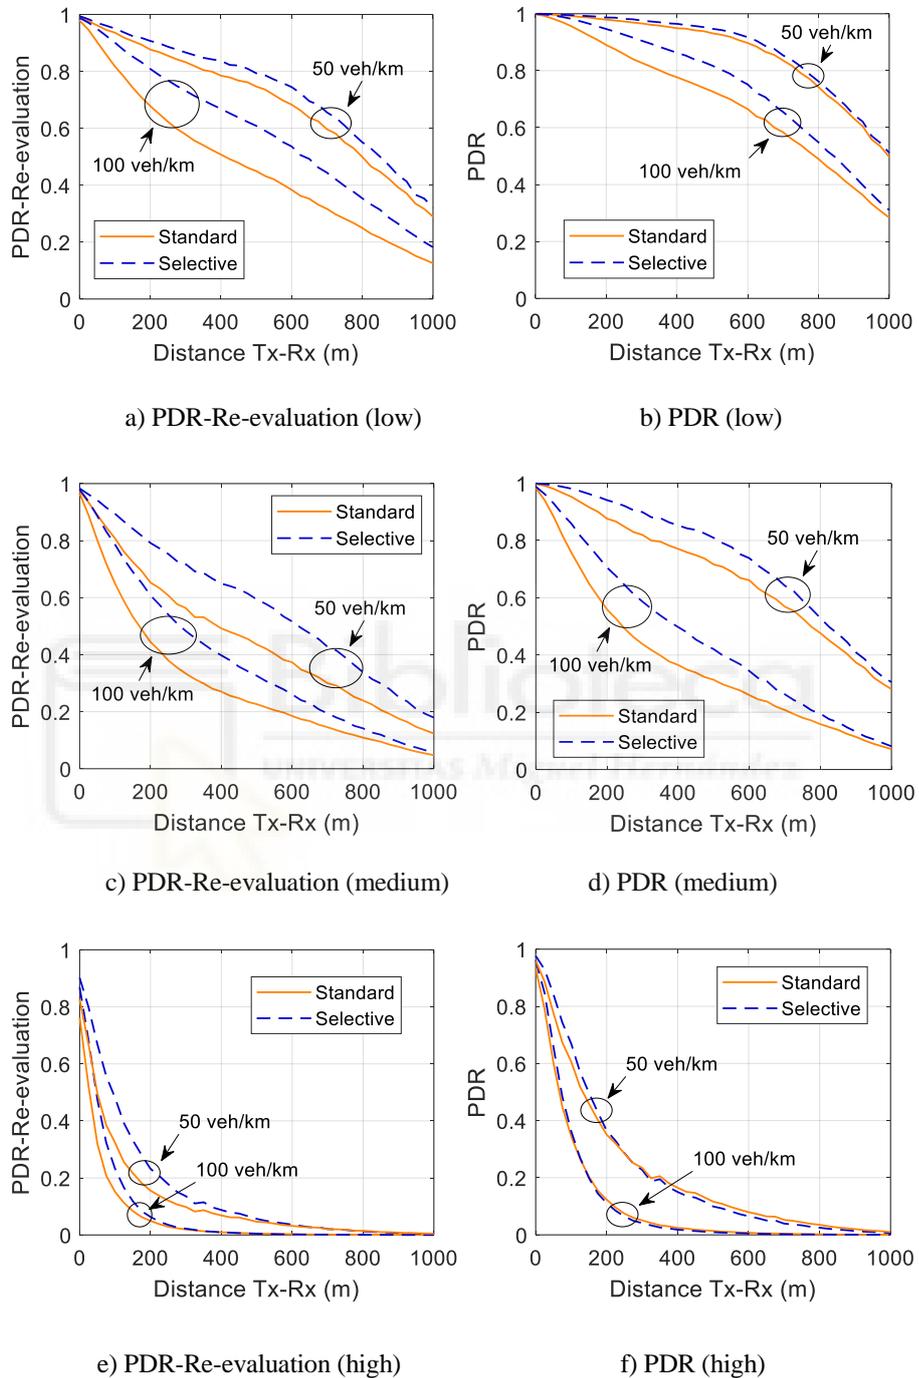


Figure 19. Reliability of the standard and selective re-evaluation mechanisms for aperiodic traffic of variable size in low, medium and high intensity scenarios with 50 and 100 veh/km [19].

The selective proposal significantly reduces the ratio of executed resource replacements (ReRR)¹² compared to the standard mechanism. Under the low intensity scenario, the standard mechanism experiences a ReRR of 40% and 52% with 50 and 100 veh/km, respectively. The selective proposal reduces these values to 17% and 25% respectively. Reduction levels above 50% are also observed under the medium and high intensity scenarios. The improvements observed in Figure 19-left with the selective proposal are primarily due to the fact that many of the detected collisions involving reservations for a new TB would not actually occur (Section 6.2) due to unutilized reservations. The selective proposal avoids the resource replacements triggered by these collisions compared to the standard mechanism. In fact, Figure 20¹³ shows that the ratio of unutilized reservations for new TBs is significantly higher than 50% with aperiodic traffic of variable size, which explains the PDR improvements obtained with the proposed selective re-evaluation mechanism compared to the standard one.

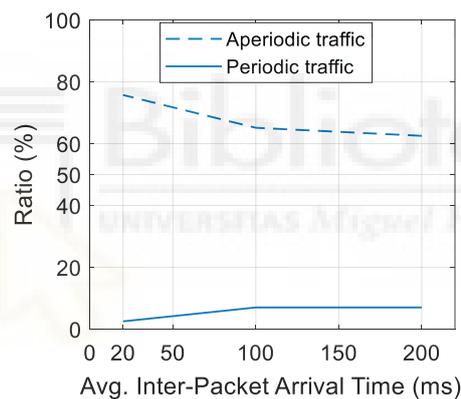


Figure 20. Ratio of Not Utilized reservations (*NUT*, Section 6.3.2). This is the ratio of reservations for a new TB that are not utilized to the total number of reservations for new TBs for 50 and 100 veh/km. Average inter-packet arrival times of 20, 100 and 200 ms correspond to the high, medium and low intensity scenarios respectively [19].

The selective proposal also helps in reducing undetectable packet collisions. When a vehicle maintains the same selected resource after a re-evaluation detection triggered by an unutilized reservation, it prevents other vehicles from transmitting in the same selected resource (Section 6.4). By reducing undetectable packet collisions, the selective re-

¹² ReRR is equal to the ReDR defined in Section 3.3 for the standard re-evaluation mechanism, as it executes a resource replacement after any re-evaluation detection. However, the selective re-evaluation mechanism only executes a resource replacement when the re-evaluation detection is triggered by a reservation for a retransmission of a TB. Therefore, the ReRR is not equal to the ReDR for the selective re-evaluation mechanism.

¹³ Figure 20 is obtained using simulations but the results are fully aligned with those obtained using the analytical models in Section 6.3 for all the scenarios.

evaluation mechanism improves not only the reliability of transmissions that have experienced a re-evaluation detection but also the reliability of transmissions that have not experienced any re-evaluation detection. This, together with the improvements obtained for the transmissions that have experienced at least one re-evaluation detection (Figure 19-left, PDR–Re-evaluation), explain the improvements in Figure 19-right. The PDR improvements obtained with the selective proposal (Figure 19-right) are particularly important in the low and medium intensity scenarios as the density and channel load increases. This is because the higher the load, the higher the risk that the standard re-evaluation mechanism encounters an undetectable packet collision after executing a resource replacement following a detected collision that was not going to actually happen.

The gains obtained with the selective proposal decrease for the high intensity scenario because the scenario is so loaded that it is not possible to avoid collisions anyway. In fact, the high loads experienced in this scenario require vehicles to execute up to 13 iterations of step 1 of the resource allocation algorithm in mode 2 to increase the RSRP threshold and reach the minimum percentage of available resources in the selection window (Section 2.3) with 100 veh/km. This is in contrast with only two iterations under the low intensity scenario. Increasing the RSRP threshold increases the probability of including reserved resources among the candidate resources, and therefore the probability of packet collisions (independently of the re-evaluation mechanism utilized).

Figure 21 compares the PDR of the standard and selective re-evaluation schemes with periodic traffic of variable size in the medium intensity scenario with 50 and 100 veh/km. The results show that the selective proposal does not improve the PDR-Re-evaluation compared to the standard re-evaluation mechanism (Figure 21(a)) under the presence of periodic traffic. This is because the ratio of not utilized reservations (*NUT*) is significantly lower than 50% with periodic traffic of variable size (Figure 20) as this traffic can only generate size reselections. In this case, the number of effective re-evaluations with the standard mechanism will be higher than the number of ineffective re-evaluations (Section 6.2), and the selective proposal is less effective for this traffic. In any case, Figure 21(b) shows that the selective proposal ultimately does not degrade the PDR compared to the standard re-evaluation mechanism since the ratio of TBs that experience one or more re-evaluation detections is below 5% in all scenarios with periodic traffic of variable size, and most of the re-evaluations are triggered by reservations for retransmissions of a TB.

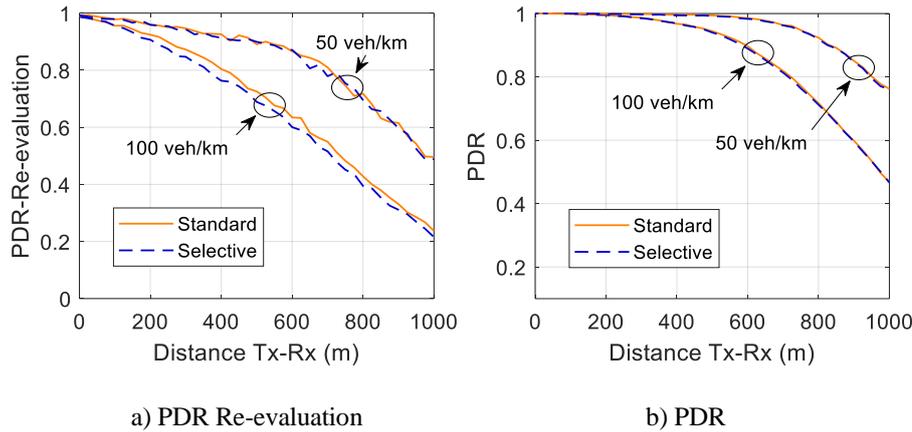


Figure 21. Reliability of the standard and selective re-evaluation mechanisms for periodic traffic of variable size in the medium intensity scenario with 50 and 100 veh/km. Similar trends are observed in the other intensity scenarios [19].

The selective re-evaluation mechanism improves the reliability of 5G NR V2X transmissions by avoiding the execution of the resource replacements that may not avoid a detected collision. Reducing the number of resource replacements can improve the latency as a resource replacement requires selecting new resources, and it is more likely that this new resource is in a slot after the initially selected one (Section 6.2). This is visible in Figure 22(a) that depicts a box plot of the Selected Resource Offset (SRO) for low, medium and high intensity scenarios and 100 veh/km when vehicles generate aperiodic traffic of variable size. This metric represents the time gap between the originally scheduled transmission time and the final transmission time for those transmissions that experience at least a re-evaluation detection. The top and bottom of the box represent the 75th and 25th percentiles, respectively, and the red horizontal line the median. The whiskers represent the 95th and 5th percentiles. Figure 22(a) clearly shows that the selective proposal significantly reduces the SRO since it does not execute uncertain resource replacements after a re-evaluation detection, and hence avoids the corresponding increase in SRO and latency. The higher reductions of SRO with the selective proposal are observed for the 95th percentile and the low intensity scenario. The highest SRO values occur when both transmissions of a TB experience a resource replacement. In this case, the new resources can be located in any slot of the selection window generated with re-evaluation, which increases the SRO¹⁴ (Section 6.2). The selective proposal reduces significantly the

¹⁴ The SRO can have negative values when a resource replacement advances the transmission, which is more likely to occur on resource replacements for retransmissions (Section 6.2).

percentage of TBs that are affected by two resource replacements for its original transmission and its retransmission compared to the standard mechanism: 5% compared to 27% in medium and low intensity scenarios with 100 veh/km. The low intensity scenario has higher PDB values, and this increases the selection window and hence the potential delay introduced by resource replacements. As the PDB decreases, the SRO differences between re-evaluation mechanisms decrease due to the shorter selection windows. This is particularly noticeable in the high intensity scenario that has a PDB of 10 ms compared to 100 ms in the low intensity scenario.

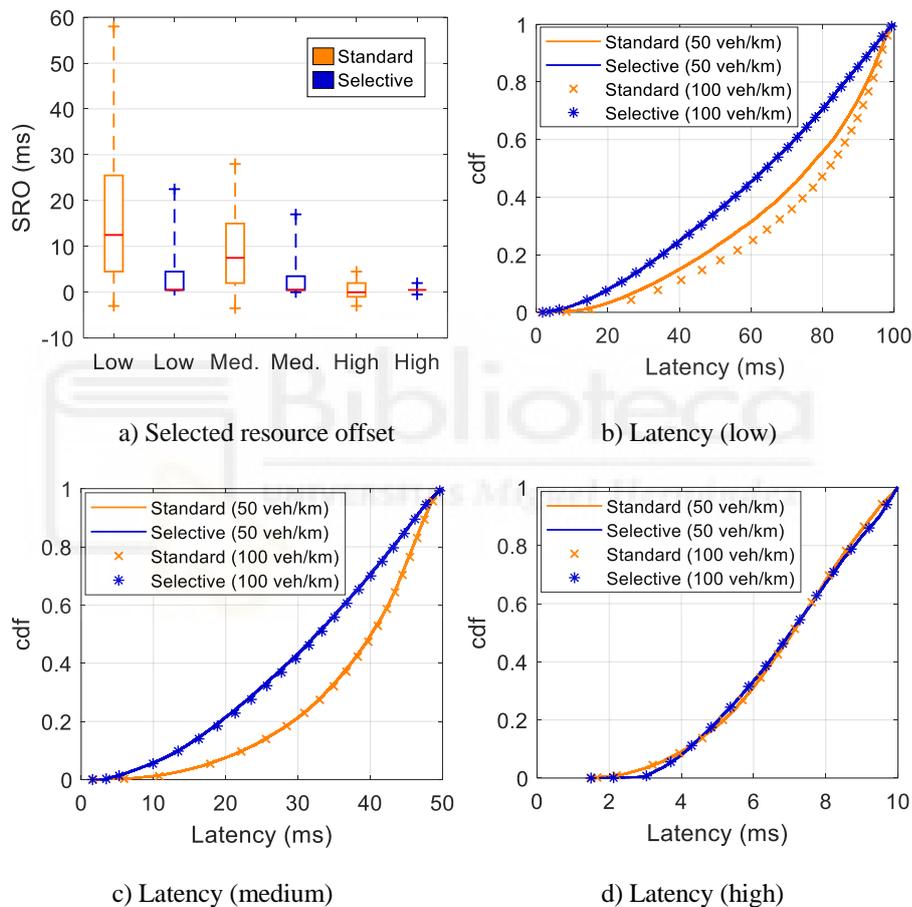


Figure 22. Impact of the re-evaluation mechanisms on the latency: SRO for 100 veh/km (a) and CDF of the latency for low intensity (b), medium intensity (c) and high intensity (d) scenarios with 50 and 100 veh/km [19].

The selective proposal significantly reduces the SRO by avoiding uncertain resource replacements, and this benefits the latency of 5G NR V2X transmissions. Figure 22(b)-(d) depict the cumulative distribution function (cdf) of the latency of transmissions that have undergone at least one re-evaluation detection in the low, medium, and high intensity scenarios with 50 and 100 veh/km. The latency is measured from the time a vehicle

generates a TB to the time it is received correctly (whether the initial transmission or the retransmission). We should note that the maximum latency (independently of the re-evaluation mechanism) is limited by the PDB of each scenario. In line with the trends observed in Figure 22(a), Figure 22(b)-(d) show that the selective proposal significantly reduces the latency under the low and medium intensity scenarios independently of the density. In the high intensity scenario, the standard and selective re-evaluation mechanisms exhibit similar SRO levels due to the smaller PDB and selection windows, and there are hence no noticeable differences in the latency.

6.6 Conclusions

This chapter presents and evaluates a novel selective re-evaluation mechanism that improves the performance of the standard 5G NR V2X mode 2 communications. The standard re-evaluation mechanism selects new resources when detecting a potential packet collision. However, many of these collisions ultimately do not happen when the traffic is aperiodic of variable size due to size and latency reselections as well as unutilized reservations. In addition, selecting new resources increases the probability of packet collisions. The proposed selective mechanism addresses these challenges by executing resource replacements only when the vehicle is certain that the collision detected with re-evaluation will occur, which is the case when a re-evaluation detection is triggered by a reservation for a retransmission of a TB. The selective re-evaluation proposal reduces the number of resource replacements, increases the reliability of NR V2X mode 2 and reduces its latency (in particular for larger PDB values) under the presence of aperiodic traffic of variable size.

7 V2X Rebroadcasting

7.1 Introduction

Connected and automated driving services generate V2X messages of varying sizes and generation times [12]. This traffic variability has been shown to significantly impact LTE V2X mode 4 [13] and NR V2X mode 2 (Chapter 4). This is because of the MAC challenges (Section 4.2) that both cellular V2X (C-V2X) technologies (i.e., LTE V2X mode 4 and NR V2X mode 2) experience when vehicles transmit aperiodic messages of variable size, which lead to a higher risk of packet collisions. NR V2X mode 2 introduces a re-evaluation mechanism to detect and prevent packet collisions [6]. However, Chapter 5 has shown that, while re-evaluation is effective with periodic traffic of fixed size, it loses effectiveness under aperiodic traffic of variable size. Chapter 6 has shown that the proposed selective re-evaluation mechanism improves the reliability of NR V2X mode 2 communications under aperiodic traffic of variable size when the MAC is configured with retransmissions. However, the analysis conducted in Chapter 4 has shown that NR V2X mode 2 also faces MAC challenges under aperiodic traffic of variable size without retransmissions (i.e. when vehicles only perform one transmission per TB) resulting in increased packet collisions. Recent studies in the literature propose mechanisms to reduce the risk of packet collisions in NR V2X mode 2 under aperiodic traffic of variable size. The study in [37] focuses on Decentralized Environmental Notification Messages (DENMs) that are generated when a new traffic event is detected. A DENM can be generated at any time, and is generally followed by one or more repetitions of the same DENM at predefined time intervals. The authors propose selecting new resources every time a new DENM is generated, and reserving the resources for the repetitions of the same DENM using the SPS scheme. The study in [38] focuses on the cooperative or collective perception service, and presents a solution that predicts the generation of Cooperative Perception Messages (CPMs) and adapts the allocation of resources to the predicted traffic. Each vehicle predicts the generation of CPMs by nearby vehicles based on their previous tracked positions and generated CPMs.

Existing proposals in the literature are bound to specific V2X messages, and this chapter extends the state-of-the-art with the first proposal that can be utilized to improve the

efficiency of the C-V2X MAC for any V2X message and traffic patterns. In particular, the proposal focuses on resolving the MAC inefficiencies generated by unutilized reservations under aperiodic traffic of variable size (see Chapter 4), and is applied in this study to the NR V2X mode 2 standard. Unutilized reservations occur when a vehicle reserves a radio resource, but does not finally use it because it has no packet to transmit at the time of the reservation. The vehicle cannot inform other vehicles about the resources it reserves for the following transmission, and this increases the risk of packet collisions. To reduce this risk, we propose *V2X Rebroadcasting*, a scheme enabling vehicles to rebroadcast packets from other vehicles in detected unutilized reservations. By doing so, the vehicle can utilize the rebroadcasted packet to reserve resources for its next transmission, and hence reduces the risk of packet collisions. In addition, the proposal increases the reliability of rebroadcasted packets as it can help resolve potential packet collisions and mitigate propagation losses experienced by the original transmission of the packet. We propose, evaluate and optimize different algorithms to select the packets to rebroadcast, and compare the performance achieved with *V2X Rebroadcasting* to the performance experienced with the standard NR V2X mode 2. The results demonstrate that *V2X Rebroadcasting* eliminates unutilized reservations, reduces the probability of packet collisions, improves the reliability of V2X communications, and reduces the packet inter-reception time. The benefits can diminish under very high channel load levels, but *V2X Rebroadcasting* still improves the reliability of NR V2X communications at short to medium distances where safety-critical V2X data is more necessary and relevant.

It should be noted that the *V2X Rebroadcasting* proposal can be applied in scenarios where NR V2X mode 2 is configured to select one candidate resource to perform one transmission per TB, as it is the case of this chapter. *V2X Rebroadcasting* then complements the selective re-evaluation mechanism proposal (see Chapter 6) that can only be applied when NR V2X mode 2 is configured with retransmissions. With one transmission per TB, only the SPS scheme is able to reserve resources since the DS scheme can only reserve resources for retransmissions (see Section 2.3). SPS and reserved resources were introduced for a higher stability in the operation of NR V2X mode 2, and this chapter focuses on SPS.

The rest of this chapter is organized as follows. Section 7.2 describes the impact of the RRI selection strategy on MAC challenges. Section 7.3 details the *V2X Rebroadcasting* proposal, including the various algorithms designed to select the packet to rebroadcast. Section 7.4 introduces some specific aspects of the simulation environment for this chapter.

Section 7.5 provides an evaluation of *V2X Rebroadcasting* and the different designed algorithms. Section 7.6 includes a performance comparison between the standard NR V2X mode 2 and *V2X Rebroadcasting* (utilizing the highest-performing algorithm evaluated in Section 7.5). Finally, Section 7.7 summarizes the main conclusions of this study. It should be noted that this chapter is related to a journal paper that we have elaborated and that will be submitted for publication just after the submission of this thesis [20].

7.2 Impact of RRI selection strategy on MAC challenges

Latency reselections and unutilized reservations depend on the selected *RRI*, the PDB and the time between TBs or inter-packet arrival time of the traffic (see Section 4.2). Latency reselections occur when $RRI > PDB$, while unutilized reservations occur when the *RRI* is smaller than the inter-packet arrival time. The PDB and inter-packet arrival time are traffic parameters that cannot be modified. However, vehicles can select the *RRI*. SAE proposes using an *RRI* equal to the minimum of the inter-packet arrival time (referred to as minimum *RRI* strategy) [14]. Chapter 4 demonstrated that this *RRI* selection strategy yields slightly better performance than configuring the *RRI* equal to the average of the inter-packet arrival time (referred to as average *RRI* strategy). This is because reducing the value of the *RRI* decreases the probability of latency reselections. Specifically, the minimum *RRI* strategy reduces the ratio of TBs experiencing latency reselections to a maximum of only 5%, and the total ratio of latency and size reselections is below 10% (see Chapter 4). Despite its higher performance, there is margin for improvement with the minimum *RRI* strategy since reducing the value of the *RRI* increases the probability of unutilized reservations under aperiodic traffic of variable size. NR V2X mode 2 can experience a ratio of ~55% of unutilized reservations relative to the total number of reservations (see Chapter 4). Unutilized reservations generate transmissions on selected resources that have not been reserved, and are therefore prone to generate packet or TB collisions (see Section 4.2.3).

7.3 V2X Rebroadcasting

We propose *V2X Rebroadcasting*, a novel MAC scheme designed to mitigate the packet collisions generated by unutilized reservations, which represent the dominant MAC challenge when NR V2X mode 2 is configured with the minimum *RRI* strategy proposed by SAE ([14], Section 4.2.3).

With *V2X Rebroadcasting*, vehicles use unutilized reservations to rebroadcast a TB or packet received from neighboring vehicles instead of leaving the reservation unutilized. Vehicles use these rebroadcasts to announce the resources they reserve for their next TB, which reduces the risk of packet collisions generated by unutilized reservations. A vehicle identifies a potential unutilized reservation in a reserved resource when it has not generated a new TB in the preceding slot to the reserved resource. In this case, the vehicle must select for rebroadcasting an eligible TB received from a neighboring vehicle. From the point of view of the resource allocation algorithm in NR V2X mode 2, the vehicle treats the TB selected for rebroadcasting as if it was a TB that the vehicle generated.

The operation of *V2X Rebroadcasting* is depicted in Figure 23. Vehicles V_A and V_B generate packets at slots s_{G1} and s_{G2} , respectively. Subsequently, both vehicles select new resources at s_{R1} and s_{R2} . V_B transmits a TB at s_{R2} and reserves the next resource at slot s_{R3} . At slot $s_{R3}-1$, V_B checks that it does not have a TB to be transmitted in s_{R3} , which would generate an unutilized reservation. V_B selects the packet received from V_A at s_{R1} and rebroadcasts it at slot s_{R3} . V_B uses this rebroadcast to reserve the resource at slot s_{R4} for its following TB (generated at slot s_{G3}). We should note that if V_B had not rebroadcasted V_A 's packet in s_{R3} , it would not have been able to notify other vehicles that it was reserving the resources at slot s_{R4} . In this case, there will be a risk of packet collisions with other vehicles that thought that the resources at s_{R4} were available since they did not receive a reservation from V_B . Rebroadcasting protects the transmissions from V_B against packet collisions. Additionally, it offers several benefits for the vehicle whose packet is rebroadcasted (V_A in Figure 23). Let's suppose that the original transmission from V_A occurred in a selected resource, and not a reserved resource. In this case, the transmission from V_A at s_{R1} will be prone to packet collisions, while it is protected against collisions when rebroadcasted by V_B in a reserved resource at slot s_{R3} . Rebroadcasting can also extend the communication range at which V_A 's packet can be received (e.g. V_C in Figure 24).

A key aspect of *V2X Rebroadcasting* is the selection of eligible TBs or packets for rebroadcasting. To this aim, we propose and evaluate several algorithms. For all proposals, we define as an eligible TB for rebroadcasting, a TB received from other vehicles in selected (i.e. non-reserved) resources. We apply this restriction because TBs transmitted in reserved resources are less prone to packet collisions. We should note that it is possible to identify TBs transmitted in selected resources because the transmitting vehicle uses the SCI to notify reservations. All proposals also leverage the TB or packet generation time to

compute the age of potential eligible TBs. The age of a TB is computed as the time elapsed from the generation of the TB to the execution of the algorithm to select a TB for rebroadcasting. We should note that the generation time is included in the *generationDeltaTime* data element of standard V2X messages such as Cooperative Perception Messages (CPMs) or Manoeuvre Coordination Messages (MCMs) [4][5].

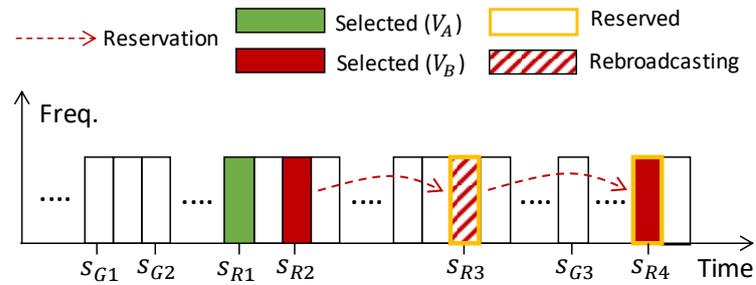


Figure 23. MAC operation of V2X Rebroadcasting [20].

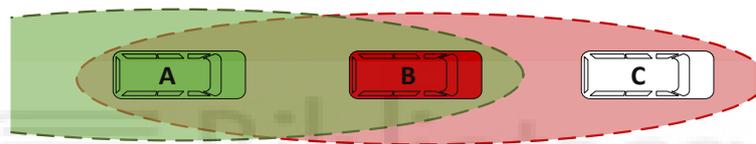


Figure 24. Illustration of V2X Rebroadcasting [20].

All algorithms share a common initial procedure to identify a first subset S_A of eligible TBs for rebroadcasting. TBs in S_A must satisfy two conditions. The first condition is that the eligible TBs must require a number of sub-channels for their transmission (NS_{Pk}) equal or smaller than the number of sub-channels of the unutilized reservation (NS_{Res}) where they would be rebroadcasted. This condition avoids a size reselection for rebroadcasting a TB, which would increase the probability of packet collisions as a new resource would have to be selected. The second condition is that the age of the TBs eligible for rebroadcasting must be smaller than 100ms, which is the maximum latency limit or PDB defined by 3GPP for V2X services [12]. This condition reduces the probability that a TB is rebroadcasted outside its latency limit. We should note that we consider the maximum latency or PDB of V2X services and not the specific PDB of each TB, since the PDB is not included in TBs as per the current 3GPP standards. After defining the first subset S_A of eligible TBs for rebroadcasting, the algorithms differ on how they select the TB that will be finally rebroadcasted.

7.3.1 Random

The *Random* algorithm randomly selects a TB from the subset S_A of TBs for rebroadcasting. A random selection reduces the risk that multiple vehicles select the same TB for rebroadcasting, which ultimately increases the number of TBs benefiting from rebroadcasting. On the other hand, a random selection can result in that a vehicle selects a TB for rebroadcasting that exceeds its latency limit.

7.3.2 Minimum Age

This algorithm selects for rebroadcasting the TB from the subset S_A with the minimum age. This approach reduces the risk of rebroadcasting a TB after its latency limit, but augments the probability that multiple vehicles select the same TB for rebroadcasting and reduces the number of eligible TBs that are rebroadcasted.

7.3.3 Low Age and Random (LAR)

LAR defines a *Threshold Age* and creates a second subset S_B of TBs that includes the TBs from S_A with an age equal or lower than *Threshold Age*. It then randomly selects from S_B the TB to be rebroadcasted. If the subset S_B is empty, *LAR* increments *Threshold Age* by Δ_{Age} , and follows the same process to create S_B . This process iterates until S_B is not empty. The *LAR* algorithm reduces the risk of rebroadcasting a TB later than its latency limit and the risk that multiple vehicles rebroadcast the same TB. The pseudocode for *LAR* is depicted in Algorithm I.

ALGORITHM I.

Input: Set of eligible TBs / Output: TB to be rebroadcasted

Execution: Slot before a potential unutilized reservation

1. **For** each TB in eligible TBs **do**
 2. Calculate *Age* and NS_{Pk} of the TB
 3. **If** $Age < 100\text{ms}$ && $NS_{Pk} \leq NS_{Res}$ **then**
 4. Add TB to subset S_A
 5. **End If**
 6. **End For**
 7. Set *Threshold Age* and Δ_{Age} equal to their initial values
 8. **Repeat**
 9. **For** each TB in S_A **do**
 10. **If** $Age \leq \text{Threshold Age}$ **then**
 11. Add TB to subset S_B
 12. **End If**
 13. **End For**
 14. **If** S_B is empty **then**
 15. $\text{Threshold Age} = \text{Threshold Age} + \Delta_{Age}$
-

-
16. **End If**
 17. **Until** S_B is not empty
 18. Randomly select a TB from subset S_B to be rebroadcasted
-

7.3.4 Low Age, Minimum Rebroadcasts and Random (LAMR)

The *LAMR* algorithm leverages the vehicle ID information included in V2X messages (e.g. the header of a CPM contains the *stationID* [4]) to further reduce the probability that multiple vehicles rebroadcast the same TB. A vehicle can detect if an eligible TB has been rebroadcasted by another vehicle when it receives a TB with the same vehicle ID and packet generation time as a TB in its stored subset of eligible TBs. Furthermore, it can estimate how many times an eligible TB has been rebroadcasted by other vehicles. *LAMR* leverages this information to prioritize rebroadcasting eligible TBs that have not been rebroadcasted by other vehicles or have been rebroadcasted fewer times, which increases the number of eligible TBs that are rebroadcasted. *NREB* denotes the number of times that a TB has been rebroadcasted by other vehicles.

ALGORITHM II.

Input: Set of eligible TBs / Output: TB to be rebroadcasted

Execution: Slot before a potential unutilized reservation

1. Lines 1-17 of Algorithm I to compute the subset S_B
 2. Set $counter = 0$
 3. **Repeat**
 4. **For** each TB in S_B **do**
 5. **If** $NREB == counter$ **then**
 6. Add TB to subset S_C
 7. **End If**
 8. **End For**
 9. **If** S_C is empty **then**
 10. $counter = counter + 1$
 11. **End If**
 12. **Until** S_C is not empty
 13. Randomly select a TB from subset S_C to be rebroadcasted
-

LAMR creates the subset S_B like *LAR*, but instead of selecting randomly a TB from S_B , it creates a new subset S_C with the TBs from S_B that have not been rebroadcasted by other vehicles. If S_C is empty, *LAMR* includes in S_C the TBs from S_B that have been rebroadcasted by other vehicles once, i.e. those with $NREB=1$. If necessary, this process is iterated augmenting $NREB$ until S_C is not empty. In each iteration, *LAMR* augments by one the number of times that a TB may have been rebroadcasted to be included in S_C . Finally, *LAMR* randomly selects the TB for rebroadcasting from S_C . The pseudocode for *LAMR* is shown in Algorithm II.

7.3.5 Low Age, Minimum Rebroadcasts and Power (LAMRP)

This algorithm creates the subset S_C like the *LAMR* algorithm, but instead of selecting randomly the TB from S_C to be rebroadcasted, it selects the one with the minimum received signal or RSRP with the objective to increase the range at which the rebroadcasted TB is received. We should remember that NR V2X mode 2 measures the RSRP of each received TB during the sensing process (see Section 2.3).

7.4 Scenario and Metrics

We evaluate the different *V2X Rebroadcasting* algorithms and compare their performance to the standard NR V2X mode 2 in the following sections. To do so, the following scenarios and configuration of parameters are employed: single traffic scenario (see Section 3.2), vehicles select 1 candidate resource in step 2 (i.e. N is equal to 1), the *RRI* is selected with the minimum *RRI* strategy following the configuration recommended by SAE [14] (see Section 7.2) and vehicles execute the re-evaluation check in the mandatory slot according to the standard, i.e. T_3 slots before the selected resources (see Section 2.4).

In this chapter, the PDR, PLR and PCR metrics (see Section 3.3) are represented as a function of the distance between transmitting and receiving vehicles (distance Tx-Rx). With *V2X Rebroadcasting*, the transmitting vehicle is the one that originally generates the TB. Additionally, this chapter uses the following evaluation metrics that have not been defined in Section 3.3 nor used in previous chapters:

- Packet within Delay Budget Ratio (PDBR): ratio of TBs rebroadcasted within their PDB to the total number of TBs rebroadcasted.
- Packet Inter-Reception (PIR) time: time elapsed between two TBs successfully received from the same transmitting vehicle. The PIR is measured within a distance Tx-Rx range $(0,D]$ as indicated by 3GPP guidelines [12].
- Available Resources Ratio (ARR): ratio of available candidate resources after the first iteration of step 1 of the resource allocation algorithm to the total number of candidate resources.
- Number of Rebroadcastings (NREB): number of times a TB is rebroadcasted using *V2X Rebroadcasting*.

7.5 Evaluation of V2X Rebroadcasting

This section analyzes first the most favorable configuration of the Low Age algorithms, and then compares the performance of the different *V2X Rebroadcasting* algorithms.

7.5.1 Configuration and selection of Low Age algorithms

We first analyze the most adequate configuration of the *Threshold Age* and Δ_{Age} values for the Low Age algorithms. To this aim, we consider the *LAMR* algorithm since similar trends were observed for *LAR* and *LAMRP*. Figure 25 represents the PDR of *LAMR* for various values of *Threshold Age* and Δ_{Age} under the medium intensity traffic scenario. The figure reveals that Δ_{Age} has a negligible impact on the PDR. This is because vehicles only had to increase *Threshold Age* by Δ_{Age} to find an eligible TB for rebroadcasting in less than 8% of the times that the *V2X Rebroadcasting* algorithm was executed. Figure 25 also shows that the highest PDR is achieved in this scenario with the highest values of *Threshold Age* analyzed, i.e. 50 ms and 100 ms. This is because augmenting *Threshold Age* increases the number of TBs eligible for rebroadcasting, reduces the probability that multiple vehicles rebroadcast the same TB, and augments the probability that each eligible TB is rebroadcasted at least once. This is visible in Table 7 that reports the probability that TBs are rebroadcasted once or more than once. The table shows that the highest values of *Threshold Age* increase the probability that TBs are rebroadcasted once, while the lowest value tends to select the same TBs more often when there is an opportunity to rebroadcast a TB. Reducing the number of TBs that benefit from rebroadcasting decreases the PDR. We should note that *Threshold Age* does not have a significant impact on the PDR of *LAMR* under the high intensity traffic scenario. In this scenario, the average inter-packet arrival time is equal to 20 ms (versus 100 ms in the medium intensity scenario), and vehicles transmit a large number of TBs in a short time. In this case, *LAMR* has a sufficient number of eligible TBs that have not been rebroadcasted yet with an age lower than 10 ms. This leads to *LAMR* with a *Threshold Age* of 10 ms achieving nearly the same probability of rebroadcasting TBs once than using a *Threshold Age* of 50 or 100 ms, which results in the same PDR for all *Threshold Age* values under the high intensity traffic scenario.

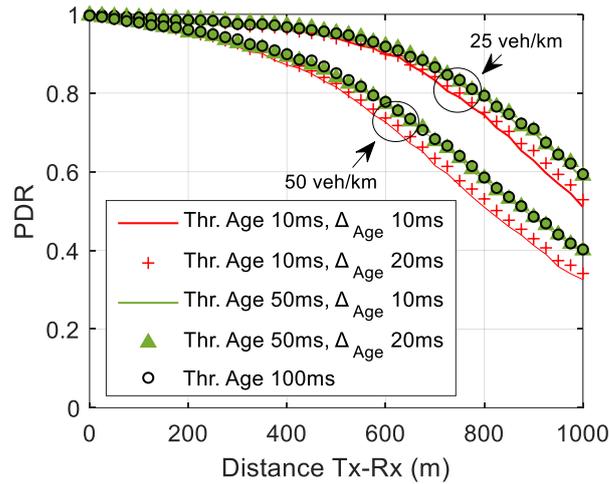


Figure 25. PDR of LAMR for different values of Threshold Age and Δ_{Age} . Medium intensity traffic [20].

Table 7. Probability that TBs are rebroadcasted once ($P(NREB=1)$) or more than once ($P(NREB > 1)$) for LAMR. Medium intensity traffic and 50 veh/km. Similar trends were observed for 25 veh/km.

Threshold Age	$P(NREB=1)$	$P(NREB>1)$
10ms	0.48	0.52
50ms	0.75	0.25
100ms	0.79	0.21

Table 8 shows that the PDBR decreases with higher values of *Threshold Age* for both traffic intensity scenarios. This is because increasing *Threshold Age* augments the age of eligible TBs, and increases the risk of rebroadcasting a TB outside its latency limit. We should note that the PDBR does not vary with the vehicle density because it only depends on the ratio of eligible TBs that can be rebroadcasted within their latency limit, and this ratio only depends on the PDB and the age limit established for a TB to be eligible for rebroadcasting. Δ_{Age} also has a negligible impact on the PDBR.

Table 8. Impact of Threshold Age on the PDBR for LAMR

Traffic	Medium intensity			High intensity		
Threshold Age	10ms	50ms	100ms	10ms	50ms	100ms
PDBR	100%	99%	47%	96%	18%	10%

Following the trends observed for the PDR and PDBR, we set the value of *Threshold Age* equal to the PDB of each scenario, i.e. equal to 50 ms and 10 ms for the medium and high intensity traffic scenarios, respectively. This configuration ensures PDBR values higher than 95% in all scenarios while maximizing the PDR. Δ_{Age} has a negligible impact in both the PDR and PDBR, and we set Δ_{Age} equal to 20ms for the rest of the chapter.

Figure 26 compares the PDR experienced by the Low Age algorithms considering their most favourable configuration under the medium intensity traffic scenario. The figure shows that *LAMR* achieves the highest PDR. *LAMR* outperforms *LAR* because *LAMR* considers the number of times that eligible TBs have been rebroadcasted by other vehicles, and prioritizes rebroadcasting TBs that have not been yet rebroadcasted or that have been rebroadcasted fewer times. This approach increases the number of TBs that are rebroadcasted, and augments the PDR. This is visible in Table 9 that shows that *LAMR* tends to rebroadcast more TBs once, while *LAR* increases the probability that fewer TBs are rebroadcasted more than once.

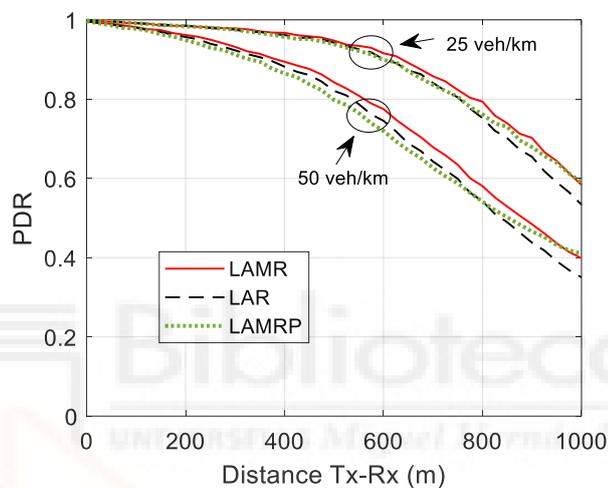


Figure 26. PDR of Low Age algorithms for the medium intensity traffic scenario. PDB = 50ms, Threshold Age = PDB, $\Delta_{Age} = 20\text{ms}$ [20].

Table 9. Probability that TBs are rebroadcasted once ($P(\text{NREB}=1)$) or more than once ($P(\text{NREB} > 1)$) for Low Age algorithms. Medium intensity traffic and 50 veh/km. Similar trends were observed for high intensity traffic and 25 veh/km.

Algorithm	$P(\text{NREB}=1)$	$P(\text{NREB} > 1)$
LAMR	0.75	0.25
LAMRP	0.75	0.25
LAR	0.54	0.46

We should note that all algorithms have the same number of opportunities to rebroadcast TBs, and they all achieve an URR (see Section 3.3) equal to 0. *LAMRP* also increases the number of TBs that are rebroadcasted once, but favours the selection of TBs with lower RSRP to counteract propagation losses¹⁵. This benefits the PDR at larger distances but reduces the number of TBs from vehicles at shorter distances that are rebroadcasted. This

¹⁵There are more eligible TBs than rebroadcasting opportunities in all the scenarios.

impacts the capacity to resolve packet collisions at shorter distances, which are the most critical conditions between two vehicles with a potential safety risk. Similar trends have been observed for the high intensity traffic scenario. In terms of PDBR, all Low Age algorithms perform equally well, achieving a PDBR higher than 98% for the medium intensity traffic scenario, and higher than 95% for the high intensity traffic scenario. Based on the observed trends in PDR and PDBR, we select *LAMR* as the most suitable Low Age algorithm for *V2X Rebroadcasting*.

7.5.2 Evaluation of V2X Rebroadcasting algorithms

This section compares the performance obtained with *LAMR*, *Random* and *Minimum Age* in order to identify the best *V2X Rebroadcasting* algorithm to compare against the standard 5G NR V2X mode 2 operation in Section 7.6.

Figure 27 shows that *LAMR* outperforms *Random* and *Minimum Age* in terms of PDR. Results are shown for the medium intensity traffic scenario but similar trends were observed for the high intensity traffic scenario. All proposals have the same opportunities to rebroadcast TBs, and they all experience an URR equal to zero, i.e. there are no unutilized reservations. However, *LAMR* augments the number of eligible TBs that are rebroadcasted, which reduces the PLR (Figure 27) and improves the PDR since more TBs benefit from rebroadcasting. This is visible in Table 10 that reports the probability that TBs are rebroadcasted once or more than once. The table shows that *LAMR* increases the probability that TBs are rebroadcasted once compared to *Random* and *Minimum Age*. This results in more TBs benefiting from rebroadcasting, thereby reducing collisions and mitigating propagation losses for more TBs. On the other hand, *Minimum Age* tends to select the same TBs more often when there is a rebroadcasting opportunity, which reduces the gains from rebroadcasting as it benefits fewer TBs.

Table 11 shows that *LAMR* experiences almost the same PDBR as *Minimum Age* and outperforms *Random*. *LAMR* and *Minimum Age* achieve PDBR values higher than 98% and 95% for medium and high intensity traffic scenarios, respectively. These results show that both algorithms are able to rebroadcast most of the TBs within their latency limit or PDB. On the other hand, *Random* experiences a significantly lower PDBR, and its PDBR decreases with the traffic intensity. This is because the PDB is lower for the high intensity scenario (10 ms) than the medium one (50 ms), and *Random* selects randomly an eligible

TB within the subset S_A that includes all TBs with an age smaller than 100 ms. The results show that *LAMR* achieves a PDBR nearly as good as *Minimum Age* while significantly improving the PDR. We compare then *LAMR* against the standard 5G NR V2X mode 2 operation in the next section.

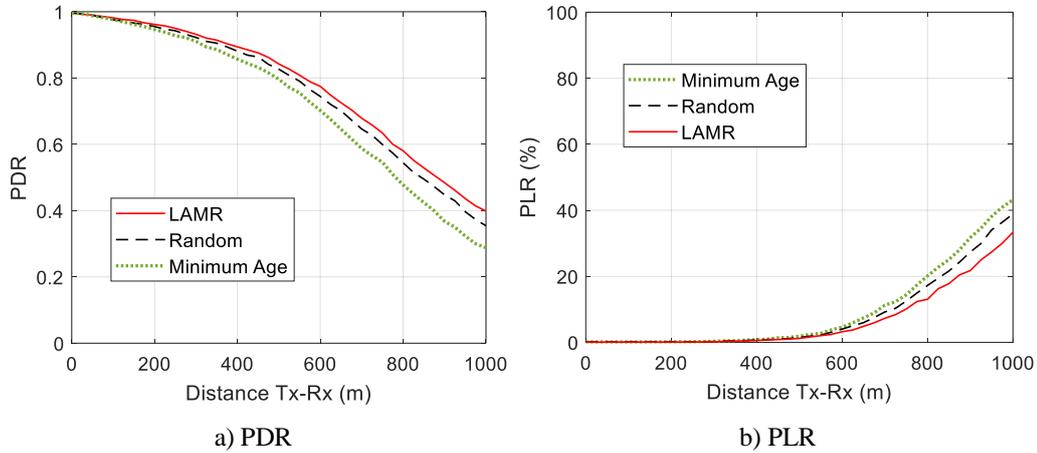


Figure 27. PDR and PLR of Random, Minimum Age and LAMR algorithms under the medium intensity traffic scenario and a density of 50 veh/km. Similar trends were observed for high traffic intensity and 25 veh/km [20].

Table 10. Probability that TBs are rebroadcasted once ($P(\text{NREB}=1)$) or more than once ($P(\text{NREB}>1)$) for Random, Minimum Age and LAMR. Medium intensity traffic and 50 veh/km. Similar trends were observed for high intensity traffic and 25 veh/km.

Algorithm	$P(\text{NREB}=1)$	$P(\text{NREB}>1)$
LAMR	0.75	0.25
Random	0.60	0.40
Minimum Age	0.27	0.73

Table 11. PDBR for LAMR (Threshold Age = PDB), Random and Minimum Age algorithms.

Traffic	Medium intensity			High intensity		
Algorithm	LAMR	Minimum Age	Random	LAMR	Minimum Age	Random
PDBR	99%	100%	43%	96%	99%	8%

7.6 Comparison of V2X Rebroadcasting and the standard 5G NR V2X mode 2

Figure 28 compares the PDR achieved with the standard NR V2X mode 2 resource allocation algorithm and with *V2X Rebroadcasting* when using the *LAMR* algorithm. The figure plots the PDR achieved under the medium and high intensity traffic scenarios and the two vehicle densities. Figure 28 shows that *V2X Rebroadcasting* outperforms the standard under the medium intensity traffic scenario, as well as for short to medium distances when the traffic intensity increases. The trends observed for the PDR in Figure

28 are explained by the trends of the PCR and PLR depicted in Figure 29. In general, the PLR augments with the distance due to propagation errors. However, *V2X Rebroadcasting* improves the PLR for all the scenarios. This is because rebroadcasted TBs can be received with better signal level at a given distance, and they can reach larger distances as illustrated with the example in Figure 24.

Figure 29 shows that *V2X Rebroadcasting* reduces the PCR¹⁶ at short distances across all scenarios, but increases it at long distances under the high intensity traffic scenarios. The crossover point between the PCR curves (Figure 29(b)) matches the crossover point between the PDR curves (Figure 28(b), 25 veh/km). *V2X Rebroadcasting* decreases the PCR because it eliminates all unutilized reservations (i.e. URR=0¹⁷) compared to the standard NR V2X mode 2 that exhibits a URR of ~55% for all scenarios. In addition, *V2X Rebroadcasting* maintains the sum of the ratio of size and latency reselections below 10% like the standard NR V2X mode 2 (see Chapter 4). Eliminating unutilized reservations decreases the percentage of TBs transmitted in selected resources, and augments the percentage of TBs transmitted in reserved resources. The use of selected resources increases the probability of packet collisions (and hence augments the PCR), since nearby vehicles are not aware that the transmitting vehicle was planning to transmit a TB over a selected resource. This risk decreases with *V2X Rebroadcasting* as vehicles do not have to transmit a TB using selected resources after a detected unutilized reservation, and transmit instead using reserved resources. In this case, the transmitting vehicle announces the reservation of resources using the SCI of its previous TB, and nearby vehicles discard then the reserved resources when selecting resources to transmit their own TBs.

¹⁶ We should note that the PCR first increases with the distance, and from a certain distance it starts decreasing. This distance is the distance from which PLR becomes the dominant source of packet errors. By definition, PCR excludes propagation errors.

¹⁷ URR is always equal to 0 for *V2X Rebroadcasting* because there are always more eligible TBs than rebroadcasting opportunities. In fact, the ratio of eligible TBs to rebroadcasting opportunities is higher than 1.3 in all the scenarios.

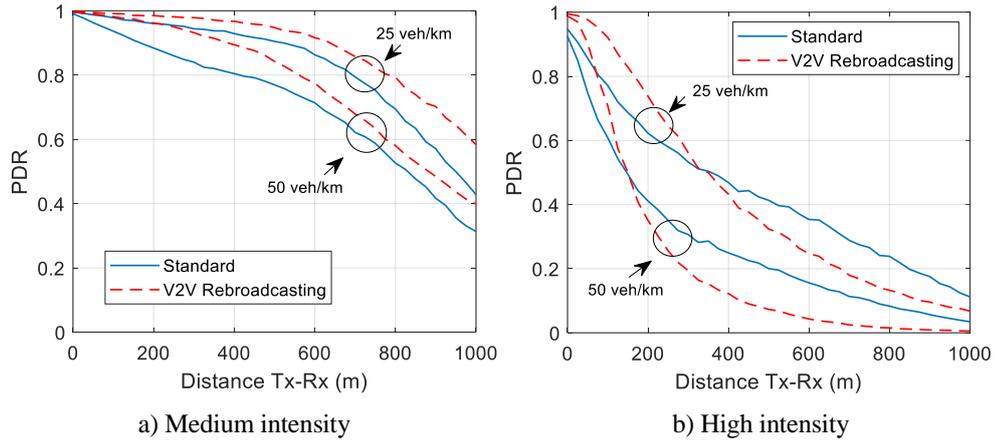


Figure 28. PDR achieved by the standard NR V2X mode 2 and V2X Rebroadcasting for (a) medium intensity and (b) high intensity traffic scenario [20].

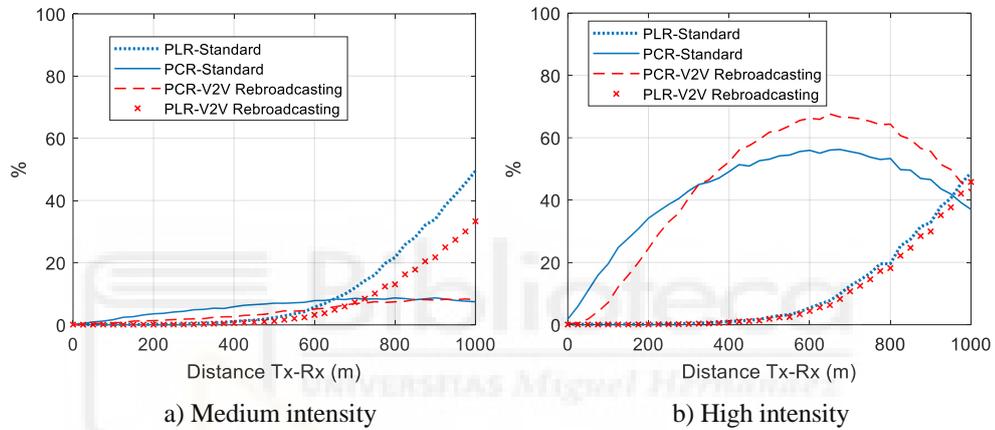


Figure 29. PCR and PLR of V2X Rebroadcasting and standard NR V2X mode 2 for (a) medium intensity and (b) high intensity traffic scenarios with 25 veh/km. Similar trends were observed for a density of 50 veh/km [20].

Despite eliminating unutilized reservations, Figure 29(b) shows that *V2X Rebroadcasting* increases the PCR from medium to long distances under the high intensity traffic scenario when the channel load increases to very high levels. This occurs when the gains obtained by avoiding unutilized reservations are offset by the reduction of the ARR (Available Resources Ratio) with *V2X Rebroadcasting* as a result of the lower unutilized reservations and the higher number of transmissions in reserved resources. When the ARR decreases, vehicles have less candidate resources to choose from when executing the mode 2 resource allocation algorithm. If the ARR decreases below 20%, the algorithm must execute a new or multiple iterations of step 1 until ARR is above 20%. In each iteration, the value of the RSRP threshold is increased by 3dB (Section 2.3). This increases the likelihood for a vehicle selecting new resources to reuse resources reserved by a vehicle that is at long distances (i.e. with a low RSRP), leading to an increase in packet collisions at long

distances. This effect emerges in the high intensity traffic scenario since the number of transmissions and reservations is higher. For example, the ARR is above 20% in the medium intensity traffic scenario when the density is 25 veh/km; in particular, the ARR is equal to 72% and 37% with the standard NR V2X mode 2 and with *V2X Rebroadcasting*, respectively. The ARR decreases to 30% and 9% with the standard NR V2X mode 2 and *V2X Rebroadcasting*, respectively, in the high intensity traffic scenario with a density of 25 veh/km. Since the ARR is below 20% with *V2X Rebroadcasting*, vehicles have to reuse reserved resources used by vehicles at long distances (i.e. with a low RSRP), and the PCR augments at long distances. The lower the ARR is, the higher the number of iterations of step 1 that must be executed¹⁸. At each iteration, the threshold RSRP increases and more reserved resources by vehicles at large distances become candidate resources, which increases the likelihood of packet collisions at large distances. This explains why *V2X Rebroadcasting* augments the PCR at medium to large distances (Figure 29(b)) and decreases the PDR for these distances (Figure 28(b)) compared to the standard NR V2X mode 2. We should though note that 200 m is the most common communication range required by 3GPP for advanced V2X services to be supported by NR V2X [11] (e.g. cooperative perception and cooperative manoeuvre with high levels of automation), and *V2X Rebroadcasting* improves the PDR below 200 m. We should note that 3GPP relaxes the reliability and latency requirements for larger ranges (e.g. 500m or 1000m) [11].

Rebroadcasting a TB increases the probability to correctly receive a TB that was lost at short distances due to a packet collision. This improves the Packet Inter-Reception (PIR) time, or time between two TBs successfully received from the same transmitting vehicle. Figure 30 plots the PIR measured within the most common communication range for advanced V2X services following [11], i.e. in the range $(0, D=200]$ m. The PIR is represented using box plots, where the top and bottom of the box represent the 75th and 25th percentiles, respectively. The red horizontal line within each box denotes the median, and the whiskers outside the box represent the 95th and 5th percentiles, respectively. The median of the PIR remains constant for both standard and *V2X Rebroadcasting*, and equal to $2 \cdot RRI$ ms across all scenarios. It is noteworthy that most PIR values are multiples of the *RRI* as the *RRI* represents the period between resources in NR V2X mode 2. The median is equal

¹⁸ With *V2X Rebroadcasting*, the threshold RSRP was increased, on average, 7.9 times (compared to 1.6 times with the standard NR V2X mode 2) during the execution of step 1 of the resource allocation algorithm under the high intensity traffic scenario with a density of 25 veh/km.

to $2 \cdot RRI$ because the average inter-packet arrival time is equal to $2 \cdot RRI$, and consecutive packets are generated and transmitted with an average period of $2 \cdot RRI$ ms. In addition, the PIR is measured in the range (0,200]m where there is a high probability of correct reception of TBs for both the standard NR V2X mode 2 and *V2X Rebroadcasting*. This is why both schemes experience the same median for the PIR.

The improvement of *V2X Rebroadcasting* in PCR at short to medium distances reduces the 75th and 95th percentiles of the PIR, i.e., the higher PIR values. This is important as the higher PIR values represent instances where a vehicle goes longer without receiving a packet from another vehicle, which increases road safety risks. For instance, *V2X Rebroadcasting* reduces the 95th percentile compared to the standard NR V2X mode 2 by 40 ms for medium intensity and 50 veh/km, and by 17 ms for high intensity and 25 veh/km. *V2X Rebroadcasting* reduces the 75th percentile, for instance, by 22 ms for medium intensity and 50 veh/km. We should note that *V2X Rebroadcasting* improves the high PIR percentiles despite the PDR trends discussed as it can better handle the increase in packet collisions at shorter distances when the traffic load augments.

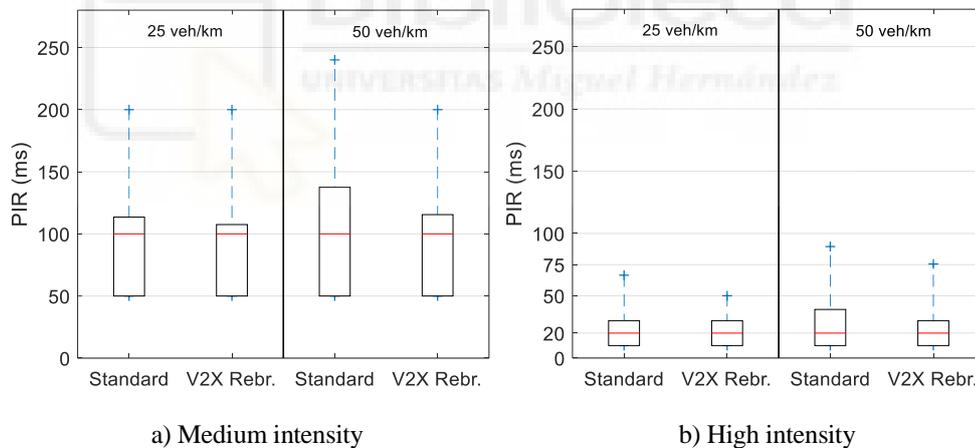
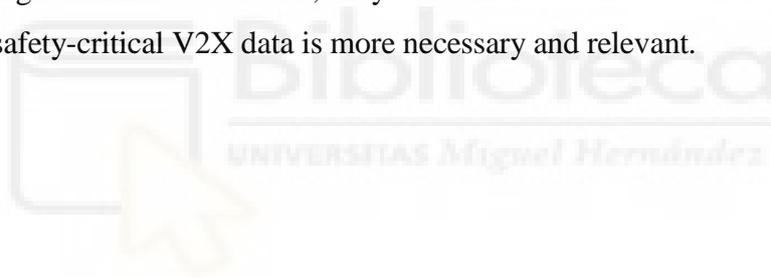


Figure 30. PIR of standard NR V2X mode 2 and V2X Rebroadcasting in the (a) medium and (b) high intensity traffic scenarios [20].

7.7 Conclusions

This chapter proposes *V2X Rebroadcasting*, a scheme designed to address the inefficiencies generated by unutilized reservations in NR V2X communications when transmitting aperiodic traffic of variable size characteristic of connected and automated driving services. Unutilized reservations increase transmissions on selected resources, which are prone to packet collisions as transmissions occur on resources that were not

previously reserved by the transmitting vehicle. With *V2X Rebroadcasting*, vehicles rebroadcast packets from other vehicles in detected unutilized reservations. By rebroadcasting a packet, the vehicle can announce its next transmission on a reserved resource, which prevents other vehicles from selecting the reserved resource and, hence, reduces the risk of packet collisions. In addition, rebroadcasting a packet can help resolve potential packet collisions and mitigate propagation losses experienced by the original transmission of the packet. The study evaluates different algorithms to select the packets to rebroadcast, and reveals that selecting packets that have not been previously rebroadcasted and with low age augments the packet delivery ratio and the probability of rebroadcasting packets within their latency limit. We compare the performance of *V2X Rebroadcasting* against the standard NR V2X mode 2, and demonstrate that *V2X Rebroadcasting* eliminates unutilized reservations, reduces the probability of packet collisions, improves the reliability of V2X communications, and reduces the packet inter-reception time. The study shows that, even if the benefits of *V2X Rebroadcasting* can diminish under high channel load levels, they are still maintained at short to medium distances where safety-critical V2X data is more necessary and relevant.



8 Conclusions and future work

The 3GPP 5G NR V2X (or NR V2X) standard supports direct or sidelink (SL) Vehicle-to-Vehicle (V2V) and Vehicle-to-Infrastructure (V2I) communications, and complements LTE V2X that was designed for basic awareness V2X services. NR V2X has been designed to support advanced V2X services for connected and automated driving with stringent requirements such as cooperative perception and driving, among others. To this aim, the standard defines two 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 [17]. This thesis has focused on NR V2X mode 2 since it ensures that V2X service provisioning is not limited by cellular coverage.

The initial focus of this thesis was an in-depth study of the 3GPP 5G NR V2X standard documents and the related literature. This study has been useful for two main purposes: 1) for writing and publishing the first journal paper of this thesis which is an in-depth and comprehensive tutorial on the 5G NR V2X standard, and 2) for developing a standard-compliant NR V2X mode 2 simulator that adheres to 3GPP standards and follows 3GPP evaluation guidelines. This simulator has been essential for implementing and evaluating the standard NR V2X mode 2 as well as the proposed enhancements in this thesis.

The first study conducted in this thesis has focused on analyzing the efficiency of NR V2X mode 2 to support advanced V2X services. These services are expected to generate aperiodic V2X traffic of variable size according to 3GPP. A key mandatory feature introduced in the MAC of NR V2X mode 2 is the re-evaluation mechanism that is designed with the objective of detecting and avoiding packet collisions. Previous studies had evaluated the performance of NR V2X mode 2 under different traffic patterns. However, these studies did not implement the re-evaluation mechanism [16]. In addition, these studies only considered periodic or aperiodic V2X traffic of fixed size. However, the 3GPP evaluation methodology guidelines for NR V2X recommend traffic generation models for advanced V2X services that also include traffic of variable size in line with the message patterns characteristic of Day 2 or Day 3 V2X services such as cooperative perception or

maneuver coordination [17]. The first study conducted in this thesis is then the first evaluation of a fully standard compliant implementation of NR V2X mode 2 using the re-evaluation mechanism under periodic and aperiodic traffic of fixed and variable packet size, in accordance with the 3GPP traffic model recommendations. The evaluation has been conducted in scenarios where vehicles perform one transmission per TB and execute the re-evaluation check only in the mandatory slot indicated by the standard. The conducted study has identified and quantified the MAC challenges that NR V2X mode 2 can face under aperiodic traffic of variable size (i.e. latency and size reselections, unutilized reservations and unused sub-channels). The configuration of NR V2X mode 2 has considered two strategies that set the *RRI* to the minimum or to the average inter-packet arrival times. Both strategies result in multiple trade-offs in terms of different MAC challenges, but none reduces the resource management instability of NR V2X mode 2 observed when transmitting aperiodic traffic with variable packet size. These instabilities have shown to reduce the PDR and increase packet collisions with respect to scenarios where vehicles generate TBs periodically. The obtained results presented similar trends to those observed with LTE V2X mode 4 when transmitting aperiodic traffic, despite the new and mandatory re-evaluation mechanism of NR V2X mode 2 that identifies and should, in principle, avoid possible packet collisions [16].

The first study conducted in this thesis provided an initial system-level evaluation of NR V2X mode 2 but does not explain why re-evaluation is not effective in certain scenarios. This thesis then focused on understanding the low effectiveness of the re-evaluation mechanism, as a crucial step for optimizing NR V2X mode 2. For example, this thesis deemed it necessary to determine if these observations hold when retransmissions are used. In this context, a new study was conducted to provide an in-depth analysis and evaluation of the impact of the re-evaluation mechanism on the operation and performance of NR V2X mode 2. This study analyzed for the first time when and why re-evaluation is effective or not in detecting and avoiding packet collisions. The conducted analysis considered vehicles transmitting periodic or aperiodic packets of fixed or variable size, following 3GPP guidelines. This study showed that the effectiveness of re-evaluation to avoid collisions depends on the data traffic patterns and mode 2 configurations. In particular, the study has shown that re-evaluation is effective in detecting collisions when vehicles transmit periodic traffic of fixed size. However, the results obtained have shown that the impact on the performance of NR V2X mode 2 is small since the number of packet

collisions detected by re-evaluation is low under periodic traffic of fixed size. This study has also shown that the effectiveness of re-evaluation can decrease under the presence of aperiodic traffic of variable size because traffic variability increases the probability of selecting new resources, and re-evaluation cannot detect collisions on new selected resources. This is particularly the case when there are no retransmissions. With retransmissions, this thesis has demonstrated that re-evaluation is more effective in detecting packet collisions, even with aperiodic traffic of variable size. However, the obtained results have shown that the impact of re-evaluation on the performance with retransmissions is low since, without re-evaluation, a TB is correctly received if just one of the two transmissions is correctly received [17].

The 3GPP standard allows different strategies for when and how often to perform re-evaluation checks. The first study conducted in this thesis considers the strategy where vehicles execute re-evaluation checks only in the mandatory slot according to the standard (this thesis named this the one-slot strategy). This thesis has also carried out a study to evaluate the impact of different re-evaluation check strategies on the performance of NR V2X mode 2 under aperiodic traffic of variable size. In particular, this thesis has evaluated two different re-evaluation check strategies proposed by the 3GPP standard for NR V2X mode 2 (one-slot and all-slots). The conducted study has shown that the all-slots strategy can significantly reduce the transmission latency at the expense of a high computational cost compared to the one-slot strategy. This study has also included an alternative re-evaluation check strategy (first-slot), not considered in the 3GPP standard. The proposed strategy has shown to reduce the latency to values similar to the all-slots strategy with a low computational cost similar to the one of one-slot strategy. The study has also shown that these benefits are achieved without degrading the reliability of the NR V2X mode 2 transmissions compared to the standardized re-evaluation check strategies [18].

The study about the impact of re-evaluation conducted in this thesis has shown that re-evaluation is not fully effective in avoiding packet collisions generated by aperiodic traffic of variable size because many of the detected collisions ultimately do not happen (due to size or latency reselections and unutilized reservations), and selecting new resources with the resource replacement increases the risk of collisions undetectable by re-evaluation. To address this inefficiency, this thesis has proposed a selective re-evaluation mechanism that only selects new resources when the vehicle is certain that a detected collision is going to occur, which is the case when a re-evaluation detection is triggered by a reservation for a

retransmission of a TB. This thesis' proposal builds from a thorough analysis and characterization of the challenges faced by the standard re-evaluation mechanism to avoid packet collisions when dealing with aperiodic packets of variable size. The proposed selective re-evaluation mechanism has shown its potential to reduce the number of resource replacements, increase the reliability of NR V2X mode 2 and reduce its latency (in particular for larger PDB values) under the presence of aperiodic traffic of variable size.

The proposed selective re-evaluation mechanism has demonstrated improvements in terms of reliability for NR V2X mode 2 communications under aperiodic traffic of variable size when the MAC is configured with retransmissions. However, the first study conducted in this thesis has shown that NR V2X mode 2 also faces MAC challenges under aperiodic traffic of variable size without retransmissions (i.e. when vehicles only perform one transmission per TB) resulting in increased packet collisions. On the other hand, recent studies in the literature have proposed mechanisms to reduce the risk of packet collisions in NR V2X mode 2 under aperiodic traffic of variable size. However, existing proposals in the literature are bound to specific V2X messages, and this thesis has extended the state-of-the-art with the *V2X Rebroadcasting* proposal. *V2X Rebroadcasting* is the first proposal that can be utilized to improve the efficiency of the NR V2X MAC for any V2X message and traffic patterns. Moreover, *V2X Rebroadcasting* can be applied in scenarios where NR V2X mode 2 is configured to select one candidate resource to perform one transmission per TB. *V2X Rebroadcasting* then complements the selective re-evaluation mechanism proposal that can only be applied when NR V2X mode 2 is configured with retransmissions. *V2X Rebroadcasting* has been proposed to resolve the MAC inefficiencies generated by unutilized reservations under aperiodic traffic of variable size. Unutilized reservations increase transmissions on selected resources, which are prone to packet collisions as transmissions occur on resources that were not previously reserved by the transmitting vehicle. With the proposed *V2X Rebroadcasting*, vehicles rebroadcast packets from other vehicles in detected unutilized reservations. By rebroadcasting a packet, the vehicle can announce the reservation for its next transmission and therefore the vehicle can perform the next transmission on a reserved resource. This prevents other vehicles from selecting the reserved resource and, hence, reduces the risk of packet collisions. In addition, rebroadcasting a packet can help resolve potential packet collisions and mitigate propagation losses experienced by the original transmission of the packet. The study conducted in this thesis has evaluated different algorithms to select the packets to

rebroadcast. The study has revealed that selecting packets that have not been previously rebroadcasted and with low age augments the packet delivery ratio and the probability of rebroadcasting packets within their latency limit. The conducted study has also compared the performance of *V2X Rebroadcasting* against the standard NR V2X mode 2, and it has demonstrated that *V2X Rebroadcasting* eliminates unutilized reservations, reduces the probability of packet collisions, improves the reliability of V2X communications, and reduces the packet inter-reception time. The study has also demonstrated that, even if the benefits of *V2X Rebroadcasting* can diminish under high channel load levels, they are still maintained at short to medium distances where safety-critical V2X data is more necessary and relevant.

With respect to future work, two potential research lines are identified:

Artificial intelligence techniques can be employed to enhance the performance of NR V2X mode 2 communications [38]. Vehicles can use these techniques, for instance, to predict both their own traffic patterns and those generated by other vehicles. By leveraging these predictions, vehicles can adapt resource allocation algorithms to improve NR V2X mode 2 communications. Within the context of this thesis, artificial intelligence techniques could be implemented and evaluated to check if they can further enhance the proposed improvements for NR V2X mode 2. For example, vehicles could predict the traffic generated by other vehicles to estimate which reservations for a new TB will be utilized. With this information, a modified (i.e. AI-based) selective re-evaluation mechanism could also execute a resource replacement after the re-evaluation detections triggered by reservations for a new TB that the vehicle predicts will be used. Additionally, a new algorithm based on artificial intelligence could be designed for *V2X Rebroadcasting* to select the eligible TB to be rebroadcasted. This algorithm would predict and select the eligible TB that, when rebroadcasted, would provide the maximum reliability enhancement to NR V2X mode 2.

The proposals presented in this thesis, as well as those suggested for future work involving artificial intelligence, could be implemented and experimentally validated through field tests to demonstrate their functionality and technical feasibility. Given the limited availability of NR V2X chipsets [39][40] and the constraints to modify existing ones to fit the proposed designs, these implementations would require open-source solutions. A possible option would be to use fully reprogrammable software-defined radio (SDR)

devices in combination with the OpenAirInterface (OAI) [41] open-source software platform. This platform integrates the 3GPP protocol stack and enables real-time emulation of mobile devices, base stations, and core network nodes.



9 Conclusiones y trabajo futuro

El estándar 3GPP 5G NR V2X (o NR V2X) soporta comunicaciones directas o sidelink (SL) de Vehículo a Vehículo (V2V) y de Vehículo a Infraestructura (V2I), y complementa LTE V2X que fue diseñado para servicios V2X de seguridad vial básicos. NR V2X ha sido diseñado para soportar servicios V2X avanzados para conducción conectada y automatizada con requisitos estrictos, tales como percepción y conducción cooperativas, entre otros. Con este fin, el estándar define dos modos de operación: modo 1 y modo 2. En el modo 1, la infraestructura celular gestiona y selecciona los recursos de comunicación para cada comunicación SL, mientras que en el modo 2, los vehículos seleccionan y gestionan autónomamente los recursos radio sin el soporte de la infraestructura celular. Esta tesis se ha centrado en NR V2X modo 2, ya que asegura que la provisión del servicio V2X no esté limitada por la cobertura celular.

El trabajo inicial de esta tesis fue un estudio profundo de los documentos del estándar 3GPP 5G NR V2X y la literatura relacionada. Este estudio ha sido útil para dos propósitos principales: 1) para escribir y publicar el primer artículo de esta tesis, el cual es un tutorial profundo y comprensivo sobre el estándar 5G NR V2X, y 2) para desarrollar un simulador de NR V2X modo 2 conforme a los estándares del 3GPP, el cual también sigue la metodología de evaluación indicada por el 3GPP. Este simulador ha sido útil para implementar y evaluar el estándar NR V2X modo 2, así como las mejoras propuestas en esta tesis.

El primer estudio realizado en esta tesis se ha centrado en analizar la eficiencia de NR V2X modo 2 para soportar servicios V2X avanzados. Se espera que estos servicios generen tráfico V2X aperiódico de tamaño variable según el 3GPP. Una característica clave introducida en la MAC de NR V2X modo 2 es el mecanismo de re-evaluación o *re-evaluation mechanism* diseñado con el objetivo de detectar y evitar colisiones de paquetes. Estudios previos habían evaluado el rendimiento de NR V2X modo 2 bajo diferentes patrones de tráfico. Sin embargo, estos estudios no implementaron el mecanismo de re-evaluación. Además, estos estudios solo consideraron tráfico periódico o aperiódico de tamaño fijo. Sin embargo, las directrices de la metodología de evaluación del 3GPP para NR V2X recomiendan modelos de generación de tráfico para servicios V2X avanzados

que también incluyan tráfico de tamaño variable en línea con los patrones de mensajes característicos de los servicios V2X de Día 2 o Día 3, tales como percepción cooperativa o coordinación de maniobras. El primer estudio realizado en esta tesis es entonces la primera evaluación de una implementación completamente conforme al estándar de NR V2X modo 2, utilizando el mecanismo de re-evaluación bajo tráfico periódico y aperiódico de tamaño fijo y variable, de acuerdo con las recomendaciones del modelo de tráfico del 3GPP. La evaluación se realiza en escenarios donde los vehículos realizan una transmisión por TB y ejecutan el chequeo de re-evaluación o *re-evaluation check* solo en el slot obligatorio indicado por el estándar. El estudio realizado identifica y cuantifica los desafíos a nivel MAC que NR V2X modo 2 puede experimentar bajo tráfico aperiódico de tamaño variable (es decir, reselecciones de latencia y tamaño, reservas no utilizadas y subcanales no utilizados). La configuración de NR V2X modo 2 ha considerado dos estrategias que establecen el RRI al mínimo o al promedio de los tiempos de llegada entre paquetes. Ambas estrategias resultan en múltiples compensaciones en términos de diferentes desafíos de la MAC, pero ninguna reduce la inestabilidad de gestión de recursos de NR V2X modo 2 observada al transmitir tráfico aperiódico de tamaño variable. Estas inestabilidades reducen la PDR y aumentan las colisiones de paquetes con respecto a escenarios donde los vehículos generan TBs periódicamente. Los resultados obtenidos presentan tendencias similares a las observadas con LTE V2X modo 4 al transmitir tráfico aperiódico, a pesar del nuevo y obligatorio mecanismo de re-evaluación de NR V2X modo 2 que, en principio, debería identificar y evitar posibles colisiones de paquetes.

El primer estudio realizado en esta tesis proporciona una evaluación inicial a nivel de sistema de NR V2X modo 2, pero no explica por qué el mecanismo de re-evaluación no es efectivo en ciertos escenarios. Entender esto es crucial para optimizar NR V2X modo 2. También es necesario determinar si estas observaciones se mantienen cuando se utilizan retransmisiones. En este contexto, se realizó un nuevo estudio para proporcionar un análisis profundo y una evaluación del impacto del mecanismo de re-evaluación en el funcionamiento y rendimiento de NR V2X modo 2. Este estudio es el primero que analiza cuándo y por qué el mecanismo de re-evaluación es efectivo o no para detectar y evitar colisiones de paquetes. El análisis considera vehículos que transmiten paquetes periódicos o aperiódicos de tamaño fijo o variable, siguiendo las directrices del 3GPP. Este estudio muestra que la efectividad del mecanismo de re-evaluación para evitar colisiones depende de los patrones de tráfico de datos y las configuraciones de NR V2X modo 2. En particular,

el estudio muestra que el mecanismo de re-evaluación es efectivo para detectar colisiones cuando los vehículos transmiten tráfico periódico de tamaño fijo. Sin embargo, el impacto en el rendimiento de NR V2X modo 2 es pequeño ya que el número de colisiones de paquetes detectadas por el mecanismo de re-evaluación es bajo con tráfico periódico de tamaño fijo. La efectividad del mecanismo de re-evaluación puede disminuir en presencia de tráfico aperiódico de tamaño variable porque la variabilidad del tráfico aumenta la probabilidad de seleccionar nuevos recursos, y el mecanismo de re-evaluación no puede detectar colisiones en nuevos recursos seleccionados. Esto es particularmente el caso cuando no hay retransmisiones. Con retransmisiones, el mecanismo de re-evaluación es más efectivo para detectar colisiones de paquetes, incluso con tráfico aperiódico de tamaño variable. Sin embargo, el impacto del mecanismo de re-evaluación en el rendimiento con retransmisiones es bajo ya que, sin mecanismo de re-evaluación, un TB se recibe correctamente si solo una de las dos transmisiones se recibe correctamente.

El estándar 3GPP permite diferentes estrategias para cuándo y con qué frecuencia realizar chequeos de re-evaluación o *re-evaluation checks*. El primer estudio realizado en esta tesis considera la estrategia en la que los vehículos ejecutan chequeos de re-evaluación solo en el slot obligatorio según el estándar (estrategia de *one-slot*). Esta tesis ha incluido un estudio para evaluar el impacto de diferentes estrategias de chequeo de re-evaluación en el rendimiento de NR V2X modo 2 bajo tráfico aperiódico de tamaño variable. En particular, evalúa dos estrategias diferentes de chequeo de re-evaluación propuestas por el estándar 3GPP para NR V2X modo 2 (*one-slot* y *all-slots*). El estudio realizado ha mostrado que la estrategia de *all-slots* puede reducir significativamente la latencia de transmisión a expensas de un alto costo computacional en comparación con la estrategia de *one-slot*. Este estudio presenta una estrategia alternativa de chequeo de re-evaluación (*first-slot*) que reduce la latencia a valores similares a la estrategia de *all-slots* con un costo computacional bajo similar al de la estrategia de *one-slot*. Estos beneficios se logran sin degradar la fiabilidad de las transmisiones de NR V2X modo 2 en comparación con las estrategias de chequeo de re-evaluación estandarizadas.

El estudio sobre el impacto del mecanismo de re-evaluación realizado en esta tesis ha mostrado que el mecanismo de re-evaluación no es completamente efectivo para evitar colisiones de paquetes generadas por tráfico aperiódico de tamaño variable porque muchas de las colisiones detectadas finalmente no ocurren (debido a reselecciones de tamaño o latencia y reservas no utilizadas), y la selección de nuevos recursos con el reemplazo de

recursos o *resource replacement* aumenta el riesgo de colisiones no detectables por el mecanismo de re-evaluación. Para abordar esta ineficiencia, esta tesis propone un mecanismo de re-evaluación selectivo (*selective re-evaluation mechanism*) que solo selecciona nuevos recursos cuando el vehículo está seguro de que una colisión detectada va a ocurrir, que es el caso cuando una detección de re-evaluación o *re-evaluation detection* es producida por una reserva para una retransmisión de un TB. La propuesta se basa en un análisis y caracterización exhaustivos de los desafíos experimentados por el mecanismo de re-evaluación estándar para evitar colisiones de paquetes al tratar con paquetes aperiódicos de tamaño variable. La propuesta de re-evaluación selectiva reduce el número de reemplazos de recursos, aumenta la fiabilidad de NR V2X modo 2 y reduce su latencia (en particular para valores de PDB más grandes) en presencia de tráfico aperiódico de tamaño variable.

El mecanismo de re-evaluación selectivo propuesto mejora la fiabilidad de las comunicaciones de NR V2X modo 2 bajo tráfico aperiódico de tamaño variable cuando la MAC está configurada con retransmisiones. Sin embargo, el primer estudio realizado en esta tesis ha mostrado que NR V2X modo 2 también experimenta desafíos MAC bajo tráfico aperiódico de tamaño variable sin retransmisiones (es decir, cuando los vehículos solo realizan una transmisión por TB) resultando en un aumento de colisiones de paquetes. Por otro lado, estudios recientes en la literatura han propuesto mecanismos para reducir el riesgo de colisiones de paquetes en NR V2X modo 2 bajo tráfico aperiódico de tamaño variable. Sin embargo, las propuestas existentes en la literatura están limitadas a mensajes V2X específicos, y esta tesis extiende el estado del arte con la propuesta de *V2X Rebroadcasting*. *V2X Rebroadcasting* es la primera propuesta que puede ser utilizada para mejorar la eficiencia de la MAC de NR V2X para cualquier mensaje V2X y patrones de tráfico. Además, *V2X Rebroadcasting* puede ser aplicada en escenarios donde NR V2X modo 2 está configurado para seleccionar un recurso para realizar una transmisión por TB. *V2X Rebroadcasting* complementa entonces la propuesta del mecanismo de re-evaluación selectivo que solo puede ser aplicada cuando NR V2X modo 2 está configurado con retransmisiones. *V2X Rebroadcasting* se centra en resolver las ineficiencias MAC generadas por reservas no utilizadas bajo tráfico aperiódico de tamaño variable. Las reservas no utilizadas aumentan las transmisiones en recursos seleccionados, los cuales son propensos a colisiones de paquetes ya que las transmisiones ocurren en recursos que no fueron previamente reservados por el vehículo transmisor. Con *V2X Rebroadcasting*, los

vehículos *rebroadcastean* o retransmiten paquetes recibidos de otros vehículos en reservas no utilizadas detectadas. Al retransmitir un paquete de otro vehículo, el vehículo puede anunciar su próxima transmisión en un recurso reservado, lo cual evita que otros vehículos seleccionen el recurso reservado y, por lo tanto, reduce el riesgo de colisiones de paquetes. Además, la retransmisión de un paquete de otro vehículo puede ayudar a resolver colisiones de paquetes potenciales y mitigar las pérdidas de propagación experimentadas por la transmisión original del paquete. El estudio evalúa diferentes algoritmos para seleccionar los paquetes de otros vehículos a retransmitir, y revela que seleccionar paquetes que no han sido previamente *rebroadcasteados* y con poca antigüedad (o poco tiempo transcurrido desde su generación) aumenta la fiabilidad y la probabilidad de retransmitir paquetes de otros vehículos dentro de su límite de latencia. Comparamos el rendimiento de *V2X Rebroadcasting* con el estándar NR V2X modo 2, y demostramos que *V2X Rebroadcasting* elimina todas las reservas no utilizadas, reduce la probabilidad de colisiones de paquetes, mejora la fiabilidad de las comunicaciones V2X y reduce el tiempo de recepción entre paquetes. El estudio muestra que, aunque los beneficios de *V2X Rebroadcasting* pueden disminuir bajo altos niveles de carga del canal, éstos se mantienen a distancias cortas y medias donde los datos V2X de seguridad vial son más necesarios y relevantes.

Con respecto a trabajos futuros, se identifican dos líneas potenciales de investigación:

Las técnicas de inteligencia artificial pueden ser empleadas para mejorar el rendimiento de las comunicaciones de NR V2X modo 2 [38]. Los vehículos pueden utilizar estas técnicas, por ejemplo, para predecir tanto sus propios patrones de tráfico como los generados por otros vehículos. Aprovechando estas predicciones, los vehículos pueden adaptar los algoritmos de asignación de recursos para mejorar las comunicaciones de NR V2X modo 2. Dentro del contexto de esta tesis, las técnicas de inteligencia artificial podrían implementarse y evaluarse para comprobar si pueden mejorar aún más las mejoras propuestas para NR V2X modo 2. Por ejemplo, los vehículos podrían predecir el tráfico generado por otros vehículos para estimar qué reservas para un nuevo TB serán utilizadas. Con esta información, un mecanismo de re-evaluación selectiva modificado podría también ejecutar un reemplazo de recursos o *resource replacement* después de las detecciones de re-evaluación o *re-evaluation detections* producidas por reservas para un nuevo TB que el vehículo predice que serán utilizadas. Adicionalmente, podría diseñarse un nuevo algoritmo basado en inteligencia artificial para *V2X Rebroadcasting* para seleccionar el TB elegible de otro vehículo para ser retransmitido. Este algoritmo predeciría

y seleccionaría el TB elegible que, cuando se retransmita, proporcionaría la máxima mejora en la fiabilidad a NR V2X modo 2.

Las propuestas presentadas en esta tesis, así como las sugeridas para trabajos futuros que involucran inteligencia artificial, podrían ser implementadas y validadas experimentalmente mediante pruebas de campo para demostrar su funcionalidad y viabilidad técnica. Debido a la disponibilidad limitada de chipsets NR V2X [39][40] y a las restricciones de modificar los existentes para que se ajusten a las propuestas, estas implementaciones requerirían soluciones de código abierto. Una alternativa posible sería utilizar dispositivos de radio definida por software (SDR) completamente reprogramables en combinación con la plataforma de software de código abierto OpenAirInterface (OAI) [41]. Esta plataforma integra la pila de protocolos del 3GPP y permite la emulación en tiempo real de dispositivos móviles, estaciones base y nodos de la red troncal.



Bibliography

- [1] K. Jo, J. Kim, D. Kim, C. Jang, and M. Sunwoo, "Development of autonomous car Part I: Distributed system architecture and development process", *IEEE Transactions on Industrial Electronics*, vol. 61, no. 12, pp. 7131–7140, December 2014.
- [2] F. Rosique, P. J. Navarro, C. Fernández and A. Padilla, "A Systematic Review of Perception System and Simulators for Autonomous Vehicles Research", *Sensors*, no. 19, 648, January 2019.
- [3] J. Van Brummelen, M. O'Brien, D. Gruyer and H. Najjaran, "Autonomous vehicle perception: The technology of today and tomorrow", *Transportation Research Part C*, vol. 89, pp. 384-406, April 2018.
- [4] ETSI ITS, "Intelligent Transport System (ITS); Vehicular Communications; Basic Set of Applications; Analysis of the Collective -Perception Service (CPS) ", ETSI TR 103 562 V2.1.1, December 2019.
- [5] ETSI ITS, "Intelligent Transport System (ITS); Vehicular Communications; Basic Set of Applications; Manoeuvre Coordination Service (MCS) ", ETSI TS 103 561 V0.0.6, December 2023.
- [6] 3GPP TS 38.321, "NR; Medium Access Control (MAC) protocol specification," Release 16 V16.7.0, Jan. 2022.
- [7] 3GPP, "TS 38.214 NR; Physical layer procedure for data (v16.8.0, Release 16)," 3GPP, Tech. Spec., Jan. 2022.
- [8] 3GPP, "TS 36.213 Evolved Universal Terrestrial Radio Access (E-UTRA); Physical layer procedures (v15.7.0, Release 15)," 3GPP, Tech. Spec., Sept. 2019.
- [9] 3GPP, "TS 36.321 Evolved Universal Terrestrial Radio Access (E-UTRA); Medium Access Control (MAC) protocol specification (v15.7.0, Release 15)," 3GPP, Tech. Spec., Sept. 2019.
- [10] ETSI ITS, "Intelligent Transport Systems; Vehicular Communications; Basic Set of Applications; Part 2: Specification of Cooperative Awareness Basic Service", EN 302 637-2 V1.4.1, January 2019.
- [11] 3GPP, "TS 22.186 Enhancement of 3GPP support for V2X scenarios (v17.0.0, Release 17)," 3GPP, Tech. Spec., March 2022.
- [12] 3GPP TR 37.885, "Study on evaluation methodology of new Vehicle-to-Everything (V2X) use cases for LTE and NR," Release 15 V15.3.0, June 2019.
- [13] 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," in *IEEE Access*, vol. 8, pp. 121526-121548, 2020.

- [14] SAE International, "On-Board System Requirements for LTE-V2X; V2V Safety Communications", SAE J3161/1, March 2022.
- [15] 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.
- [16] A. Molina-Galan, B. Coll-Perales, L. Lusvarghi, J. Gozalvez and M. L. Merani, "How does 5G NR V2X Mode 2 Handle Aperiodic Packets and Variable Packet Sizes?," *2022 IEEE 23rd HPSR*, Taicang, Jiangsu, China, 2022, pp. 183-188.
- [17] A. Molina-Galan, L. Lusvarghi, B. Coll-Perales, J. Gozalvez and M. L. Merani, "On the Impact of Re-Evaluation in 5G NR V2X Mode 2," in *IEEE Transactions on Vehicular Technology*, vol. 73, no. 2, pp. 2669-2683, Feb. 2024.
- [18] A. Molina-Galan, B. Coll-Perales and J. Gozalvez, "Re-Evaluation Strategies for 5G NR V2X Communications," *2022 IEEE 96th Vehicular Technology Conference (VTC2022-Fall)*, London, United Kingdom, 2022, pp. 1-5
- [19] A. Molina-Galan, J. Gozalvez and B. Coll-Perales, "A Selective Re-Evaluation Mechanism for 5G NR V2X Mode 2 Communications," in *IEEE Transactions on Vehicular Technology* (submitted).
- [20] A. Molina-Galan, J. Gozalvez and B. Coll-Perales, "V2X Rebroadcasting for C-V2X Communications".
- [21] 3GPP, "TS 38.215 NR; Physical layer measurements (v16.5.0, Release 16)," 3GPP, Tech. Spec., March 2022.
- [22] ns-3 official website, <https://www.nsnam.org/>. Accessed: May 2024.
- [23] R1-1900852, "Link level evaluations on sidelink for NR V2X," Huawei, HiSilicon, 3GPP TSG RAN WG1 Ad-Hoc Meeting 1901, Taipei, Jan. 2019.
- [24] R1-1903180, "Link level evaluations of NR PSCCH," Ericsson, 3GPP TSG RAN WG1 Meeting #96, Athens, Greece, March 2019.
- [25] 5GAA, "V2X Functional and Performance Test Report; Test Procedures and Results," 5GAA Report, Apr. 2019.
- [26] 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," in *IEEE Access*, vol. 9, pp. 145648-145661, 2021.
- [27] 3GPP, "TS 38.331 NR; Radio Resource Control (RRC); Protocol specification (v16.7.0, Release 16)," 3GPP, Tech. Spec., December 2021.
- [28] L. Lusvarghi and M. L. Merani, "On the Coexistence of Aperiodic and Periodic Traffic in Cellular Vehicle-to-Everything," in *IEEE Access*, vol. 8, pp. 207076-207088, 2020.
- [29] L. Lusvarghi and M. L. Merani, "Machine Learning for Disseminating Cooperative Awareness Messages in Cellular V2V Communications," in *IEEE Transactions on Vehicular Technology*, vol. 71, no. 7, pp. 7890-7903, July 2022.

- [30] Z. Ali, S. Lagén, L. Giupponi and R. Rouil, "3GPP NR V2X Mode 2: Overview, Models and System-Level Evaluation," in *IEEE Access*, vol. 9, pp. 89554-89579, 2021.
- [31] C. Campolo, V. Todisco, A. Molinaro, A. Berthet, S. Bartoletti and A. Bazzi, "Improving Resource Allocation for beyond 5G V2X Sidelink Connectivity," *2021 55th Asilomar Conference on Signals, Systems, and Computers*, Pacific Grove, CA, USA, 2021, pp. 55-60.
- [32] 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," *2022 International Conference on Artificial Intelligence in Information and Communication (ICAIIIC)*, Jeju Island, Korea, Republic of, 2022, pp. 171-175.
- [33] G. Thandavarayan, M. Sepulcre and J. Gozalvez, "Analysis of Message Generation Rules for Collective Perception in Connected and Automated Driving," *2019 IEEE Intelligent Vehicles Symposium (IV)*, Paris, France, 2019, pp. 134-139.
- [34] A. Correa *et al.*, "Infrastructure Support for Cooperative Maneuvers in Connected and Automated Driving," *2019 IEEE Intelligent Vehicles Symposium (IV)*, Paris, France, 2019, pp. 20-25.
- [35] G. Grimmett and D. Welsh, "Probability: An introduction," Oxford, U.K.: Oxford Science, Nov. 2014.
- [36] 3GPP, "TS 38.212 NR; Multiplexing and channel coding (v16.8.0, Release 16)," 3GPP, Tech. Spec., Dec. 2021.
- [37] S. Daw, et al., "LAMP: A latency-aware MAC protocol for joint scheduling of CAM and DENM traffic over 5G-NR sidelink," *Elsevier's Computer Communications*, vol. 217, pp. 41-56, Jan. 2024.
- [38] A. Hegde, R. Song and A. Festag, "Radio Resource Allocation in 5G-NR V2X: A Multi-Agent Actor-Critic Based Approach," *IEEE Access*, vol. 11, pp. 87225-87244, August 2023.
- [39] Autotalks, "TEKTON3," [Online]. Available: <https://autotalks.com/products/tekton3/>. [Accessed: Jun. 2024].
- [40] Autotalks, "SECTON3," [Online]. Available: <https://autotalks.com/products/secton3/>. [Accessed: Jun. 2024].
- [41] OpenAirInterface, "OpenAirInterface – 5G software alliance for democratizing wireless innovation," [Online]. Available: <https://openairinterface.org/>. [Accessed: Jun. 2024].

Annex A.1 Publication



A Tutorial on 5G NR V2X Communications

Mario H. Castañeda Garcia, Alejandro Molina-Galan, Mate Boban, Javier Gozalvez, Baldomero Coll-Perales, Taylan Şahin and Apostolos Kousaridas

Abstract— The Third Generation Partnership Project (3GPP) has recently published its Release 16 that includes the first Vehicle-to-Everything (V2X) standard based on the 5G New Radio (NR) air interface. 5G NR V2X introduces advanced functionalities on top of the 5G NR air interface to support connected and automated driving use cases with stringent requirements. This paper presents an in-depth tutorial of the 3GPP Release 16 5G NR V2X standard for V2X communications, with a particular focus on the sidelink, since it is the most significant part of 5G NR V2X. The main part of the paper is an in-depth treatment of the key aspects of 5G NR V2X: the physical layer, the resource allocation, the quality of service management, the enhancements introduced to the Uu interface and the mobility management for V2N (Vehicle to Network) communications, as well as the co-existence mechanisms between 5G NR V2X and LTE V2X. We also review the use cases, the system architecture, and describe the evaluation methodology and simulation assumptions for 5G NR V2X. Finally, we provide an outlook on possible 5G NR V2X enhancements, including those identified within Release 17.

Index Terms—5G NR V2X, 5G V2X, 3GPP, Release 16, sidelink, vehicle-to-everything, V2X, 5G, New Radio, 5G NR, LTE V2X, Cellular V2X, C-V2X, connected and automated vehicles, CAV, connected and automated driving.

I. INTRODUCTION

IN 1999, the US Federal Communications Commission (FCC) allocated 75 MHz of spectrum in the 5.9 GHz band for Intelligent Transportation Services (ITS). The allocation triggered significant research activity around the world to develop and deploy V2X communications over the past two decades (e.g., the CAMP consortium in the US [1], the Car 2 Car Communication Consortium in Europe [2], and countless research projects). V2X communications includes Vehicle-to-vehicle (V2V), Vehicle-to-Network (V2N) or Vehicle-to-Infrastructure (V2I), Vehicle-to-Road Side Unit (V2R), and Vehicle-to-Pedestrian (V2P). The research resulted in a first set of radio standards for V2X completed by 2010. These standards are based on the IEEE 802.11p technology and are referred to as Dedicated Short Range Communications (DSRC) [3]. The development of radio standards was followed by the definition of higher layer standards, message formats, protocols, and applications (e.g., [4] in Europe and [5] in the US).

Manuscript received June 17, 2020.

M. H. Castañeda Garcia, M. Boban, and A. Kousaridas are with Huawei Technologies Duesseldorf GmbH, Munich Research Center, Riesstrasse 25, 80992 Munich, Germany, e-mail: {mario.castaneda, mate.boban, apostolos.kousaridas}@huawei.com.

T. Şahin is with the Telecommunication Networks Group, Technische Universität Berlin, 10587 Berlin, Germany, e-mail: taylan.sahin.1@campus.tu-berlin.de. At the time of submission, T. Şahin was also with Huawei Technologies Duesseldorf GmbH, Munich Research Center.

3GPP Release 12 (Rel. 12) was the first standard to introduce direct Device-to-Device (D2D) communications for proximity services (ProSe) using cellular technologies [6]. This work was used by 3GPP to develop LTE V2X, the first cellular V2X (C-V2X) standards based on the 4G Long Term Evolution (LTE) air interface. LTE V2X was developed under Release 14 (Rel. 14) [7] and was further enhanced in Release 15 (Rel. 15). It is only under Release 16 (Rel. 16) that 3GPP has developed a new cellular V2X standard based on the 5G NR (New Radio) air interface. The precursor to the technical work on Rel. 16 NR V2X was the study item (SI) approved under Rel. 15. This SI developed the evaluation methodology and assumptions for LTE and NR V2X [8] that were necessary to evaluate and compare the various proposals to be included in the 5G NR V2X standard. Next, 3GPP approved a SI and a work item (WI)¹ to develop the first set of 5G NR V2X standards in Rel. 16. Specifically, the SI on radio interface technologies [9] ran until March 2019 followed by a WI [10] that officially concluded in December 2019. This WI resulted in the first set of 5G NR V2X specifications included in the 3GPP technical specifications (TS). Fig. 1 summarizes the timeline of the development of cellular V2X standards under 3GPP with a focus on Radio Access Network (RAN) developments.

The 5G NR standard was developed under Rel. 15 but it did not include sidelink (SL) aspects. SL refers to direct communication between terminal nodes or User Equipments (UEs) without the data going through the network. In NR V2X, UEs are vehicles, Road Side Units (RSUs), or mobile devices that are carried by pedestrians. Rel. 16 is the first to introduce V2X communications, including SL communications, based on the 5G NR air interface. This makes Rel. 16 NR V2X SL the first 5G V2X standard available, and a basis for future enhancements and extensions for V2X and non-V2X SL applications. As noted in [10], the NR V2X SL has been developed to complement and not replace LTE V2X SL communications. The goal of NR V2X SL is to support enhanced V2X (eV2X) use cases related to connected and automated driving [11]. Some of these use cases have requirements that cannot be satisfied by the LTE V2X standard [12].

This paper presents an exhaustive overview of the first

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

¹In 3GPP, SI refers to the more exploratory work that investigates the feasibility and benefit that certain technologies can bring. On the other hand, WI refers to specification work that often builds on top of the work performed in a SI and results in a set of technical specifications (i.e., a standard).

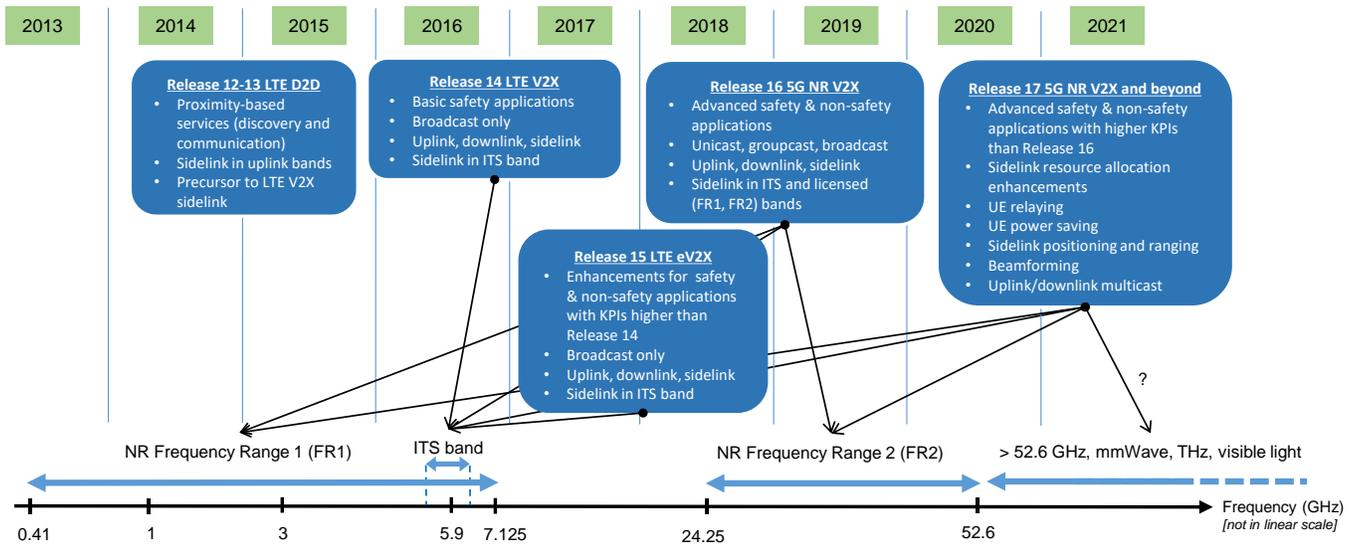


Fig. 1. Progress of 3GPP work on V2X with a focus on RAN.

standard for 5G NR V2X communications that 3GPP developed in Rel. 16. The main focus of the standard and of this paper is on NR V2X SL communications as the basis to support ubiquitous V2X communications. The paper provides a comprehensive and reference tutorial that introduces the major 3GPP standard developments essential to understand how NR V2X communications operate. To this aim, Section II briefly reviews first the LTE V2X standard designed under Rel. 14 and Rel. 15. This review helps understand and highlight the differences and novelties introduced by NR V2X under Rel. 16. One key difference between LTE V2X and NR V2X is the use cases that each technology should support. To that end, Section III presents the use cases that 5G NR V2X supports following the work done at the 3GPP and the 5G Automotive Association (5GAA) [13]. Section III also includes a review of the main use case requirements and key performance indicators (KPIs). Section IV presents a high level overview of the 5G NR system architecture for V2X communication for SL and uplink/downlink communications. Section V provides an in-depth summary of the NR V2X SL physical layer including its

structure, the physical sidelink channels, as well as physical layer sidelink procedures such as HARQ feedback, synchronization, and power control. Section VI describes the resource allocation in NR V2X SL for its two communication modes (mode 1 and mode 2) as well as procedures for supporting congestion control. Section VII describes the framework and mechanisms defined under Rel. 16 to manage quality of service (QoS) in NR V2X communications. Section VIII presents the major enhancements introduced in Rel. 16 to the Uu interface and to the mobility management in order to support V2N communications. NR V2X has been designed to complement LTE V2X. To that end, Section IX explains how the two technologies can co-exist. The publication of Rel. 16 will trigger significant efforts to evaluate the performance and capabilities of NR V2X. Section X presents the evaluation methodology defined in 3GPP that includes new channel models and assumptions for system and link level simulations. Finally, Section XI presents future possible enhancements for NR V2X communications including those already identified as study and work items under Release 17 (Rel. 17).

TABLE I
LIST OF ACRONYMS

Acronym	Definition	Acronym	Definition
3GPP	Third Generation Partnership Project	CDM	Code Division Multiplexing
5GAA	5G Automotive Association	CG	Configured Grant
5GC	5G Core	CHO	Conditional Handover
5GS	5G System	CP	Cyclic Prefix
5QI	5G New Radio Standardized Quality of Service Identifier	CQI	Channel Quality Indicator
ACK	Acknowledgement	CR	Channel Occupancy Ratio
AF	Application Function	CRC	Cyclic Redundancy Check
AGC	Automatic Gain Control	CRlimit	Maximum Channel Occupancy Ratio
AMF	Access and Mobility Management Function	CSI-RS	Channel State Information-Reference Signal
AS layer	Access Stratum layer	C-V2X	Cellular V2X
AS	Application Server	D2D	Device-to-Device (Communications)
BS	Base Station	DAPS	Dual Active Protocol Stack
BSM	Basic Safety Message	DCC	Decentralized Congestion Control
BWP	Bandwidth Part	DCI	Downlink Control Information
CAM	Cooperative Awareness Message	DENM	Decentralized Environmental Notification Message
CBR	Channel Busy Ratio	DFN	Direct Frame Number
CDL	Cluster Delay Line	DG	Dynamic Grant

TABLE I
LIST OF ACRONYMS

Acronym	Definition	Acronym	Definition
DL	Downlink	RA	Random Access
DMRS	Demodulation Reference Signal	RAN	Radio Access Network
DSRC	Dedicated Short Range Communications	RAT	Radio Access Technology
eNB	Evolved Node B (LTE Base Station)	RB	Resource Block
EPC	Evolved Packet Core	RF	Radio Frequency
EPS	Evolved Packet System	RI	Rank Indicator
ETSI	European Telecommunications Standardization Institute	RLC	Radio Link Control
eV2X	Enhanced Vehicle-to-Everything	RLF	Radio Link Failure
EVM	Error Vector Magnitude	RP	Resource Pool
FCC	Federal Communications Commission	RRC	Radio Resource Control
FDM	Frequency Division Multiplexing	RRI	Resource Reservation Interval
FEC	Forward Error Correction	RSRP	Reference Signal Received Power
FR1	Frequency Range 1	RSRQ	Reference Signal Received Quality
FR2	Frequency Range 2	RSSI	Received Signal Strength Indicator
GLOSA	Green Light Optimal Speed Advisory	RS-SINR	Reference Signal-Signal to Noise and Interference Ratio
gNB	Next generation Node B (NR Base Station)	RSU	Road Side Unit
gNB-CU	gNB-Control Unit	RV	Redundancy Version
gNB-DU	gNB-Distributed Unit	RX UE	Receiving User Equipment
GNSS	Global Navigation Satellite System	SA	3GPP Services and System Aspects
HARQ	Hybrid Automatic Repeat Request	SAE	Society of Automotive Engineers
HO	Handover	SC-FDMA	Single-Carrier Frequency-Division Multiple Access
HPLMN	Home Public Land Mobile Network	SCI	Sidelink Control Information
ISD	Inter Site Distance	SC-PTM	Single-Cell Point-To-Multipoint
ITS	Intelligent Transportation Systems	SCS	Subcarrier Spacing
KPI	Key Performance Indicator	SDAP	Service Data Adaptation Protocol
LDPC	Low-Density Parity-Check	SF	Subframe
LTE	Long-Term Evolution	SFN	System Frame Number
MAC	Medium Access Control	SI	Study Item
MAC CE	MAC Control Element	SIB	System Information Block
MBMS	Multimedia Broadcast Multicast Services	SID	Study Item Description
MBS	Multicast Broadcast Services	SL	Sidelink
MBSFN	Multimedia Broadcast Single Frequency Network	SLR	Service Level Requirement
MCS	Modulation and Coding Scheme	SLRB	Sidelink Radio Bearer
NACK	Negative Acknowledgement	SLSS	Sidelink Synchronization Signal
NAS	Non-Access Stratum	SMF	Session Management Function
NDI	New Data Indicator	SPS	Semi-Persistent Scheduling
NEF	Network Exposure Function	S-PSS	Sidelink Primary Synchronization Signal
NF	Network Functions	SR	Scheduling Request
NG-RAN	Next-Generation Radio Access Network	SSB	Synchronization Signal Block
NR	New Radio	S-SSB	Sidelink Synchronization Signal Block
NRF	Network Repository Function	S-SSS	Sidelink Secondary Synchronization Signal
NWDAF	Network Data Analytics Function	TAI	Tracking Area Identifier
OAM	Operations, Administration and Maintenance	TB	Transport Block
OFDM	Orthogonal Frequency Division Multiplexing	TDD	Time Division Duplex
P2P	Pedestrian-to-Pedestrian	TDM	Time Division Multiplexing
PBCH	Physical Broadcast Channel	TFRP	Time-Frequency Resource Pattern
PCF	Policy Control Function	TPC	Transmit power control
PDB	Packet Delay Budget	TS	Technical Specification
PDCCH	Physical Downlink Control Channel	TTI	Transmission Time Interval
PDCP	Packet Data Convergence Protocol	TX UE	Transmitting User Equipment
PDSCH	Physical Downlink Shared Channel	UDM	Unified Data Management
PFI	PC5 QoS Flow ID	UDR	Unified Data Repository
PMI	Precoding Matrix Indicator	UICC	Universal Integrated Circuit Card
PPPP	Proximity Service Per-Packet Priority	UL	Uplink
PPPR	Proximity Service Per-Packet Reliability	UPF	User Plane Functions
PQI	PC5 5QI	URLLC	Ultra-Reliable Low Latency Communication
PRB	Physical Resource Block	UTC	Coordinated Universal Time
ProSe	Proximity Service	V2I	Vehicle-to-Infrastructure
PSBCH	Physical Sidelink Broadcast Channel	V2N	Vehicle-to-Network
PSCCH	Physical Sidelink Control Channel	V2P	Vehicle-to-Pedestrian
PSFCH	Physical Sidelink Feedback Channel	V2R	Vehicle-to-RSU
PSSCH	Physical Sidelink Shared Channel	V2V	Vehicle-to-Vehicle
PT-RS	Phase Tracking-Reference Signal	V2X	Vehicle-to-Everything
PUCCH	Physical Uplink Control Channel	VPLMN	Visited Public Land Mobile Network
PUSCH	Physical Uplink Shared Channel	VRU	Vulnerable Road User
QoS	Quality of Service	WI	Work Item
R2R	RSU(Road Side Unit)-to-RSU		

II. RELEASE 14/15: LTE V2X

3GPP defined in Rel. 14 (and later refined in Rel. 15) the LTE V2X standard for V2X communications using the LTE air interface [14]. LTE V2X is expected to operate on the 5.9 GHz band reserved in certain markets (e.g., United States, Europe, China) for ITS services. For SL communications, vehicles utilize the so-called PC5 interface, whereas they utilize the Uu interface for V2N. LTE V2X has been designed to support basic cooperative active traffic safety, traffic management, and telematics applications [15]. LTE V2X supports similar services as those supported by DSRC or its European counterpart ITS-G5. These services rely on the broadcast transmission of small awareness messages such as CAMs (Cooperative Awareness Messages) in ITS-G5 [4] or BSMs (Basic Safety Messages) in DSRC [3] to regularly provide basic information such as the location, direction, speed, and acceleration of the transmitting vehicle. LTE V2X defines new PHY (Physical) and MAC (Medium Access Control) layers for V2X and reuses the upper V2X layers and protocols specified by ETSI (European Telecommunications Standardization Institute), IEEE (Institute of Electrical and Electronic Engineers), and SAE (Society of Automotive Engineers).

A. Physical Layer

LTE V2X uses SC-FDMA (Single-Carrier Frequency-Division Multiple Access) and supports 10 MHz and 20 MHz channels. The channel is divided into 180 kHz Resource Blocks (RBs) that correspond to 12 subcarriers of 15 kHz each. In the time domain, the channel is organized into 1 ms subframes. Fig. 2 illustrates the channelization in LTE V2X. Each subframe has 14 OFDM symbols with normal cyclic prefix. Nine of these symbols are used to transmit data and four of them (3rd, 6th, 9th, and 12th) are used to transmit demodulation reference signals (DMRSs) for channel estimation and combating the Doppler effect at high speeds [17]. The last symbol is used as a guard symbol for timing adjustments and for allowing vehicles to switch between transmission and reception across subframes.

RBs are grouped into sub-channels (Fig. 2). A sub-channel can include RBs only within the same subframe. The number of RBs per sub-channel can vary and is (pre-)configured. (Pre-)configuration refers to a configuration that is: 1) defined by the network and signaled to the UE by the cellular base station (eNB or gNB) when a UE is in network coverage; or 2) predefined in the UE when the UE is out of network coverage. Sub-channels are used to transmit data and control information. The data is organized in Transport Blocks (TBs) (Fig. 2) that are carried in the Physical Sidelink Shared Channel (PSSCH). A TB contains a full packet (e.g., a CAM or a BSM). A TB can occupy one or several sub-channels depending on the size of the packet, the number of RBs per sub-channel, and the utilized Modulation and Coding Scheme (MCS). TBs can be transmitted using QPSK, 16-QAM or 64QAM (introduced in Rel. 15) modulations and turbo coding.

Each TB has an associated Sidelink Control Information (SCI) message (Fig. 2) that is carried in the Physical Sidelink Control Channel (PSCCH) [18]. It is also referred to as Scheduling Assignment (SA). An SCI occupies 2 RBs and

includes information such as: an indication of the RBs occupied by the associated TB; the MCS used for the TB; the priority of the message that is being transmitted; an indication of whether it is a first transmission or a blind retransmission of the TB; and the resource reservation interval. A blind retransmission refers to a scheduled retransmission or repetition of the TB (i.e., not based on feedback from the receiver). The resource reservation interval specifies when the vehicle will utilize the reserved sub-channel(s) to transmit its next TB. The SCI includes critical information for the correct reception of the TB. A TB cannot be decoded properly if the associated SCI is not received correctly. A TB and its associated SCI must be transmitted always in the same subframe (Fig. 2). The TB and its associated SCI can be transmitted in adjacent sub-channels. Alternatively, RBs can be divided into two pools. One pool is dedicated to transmit SCIs and the other one to transmit TBs.

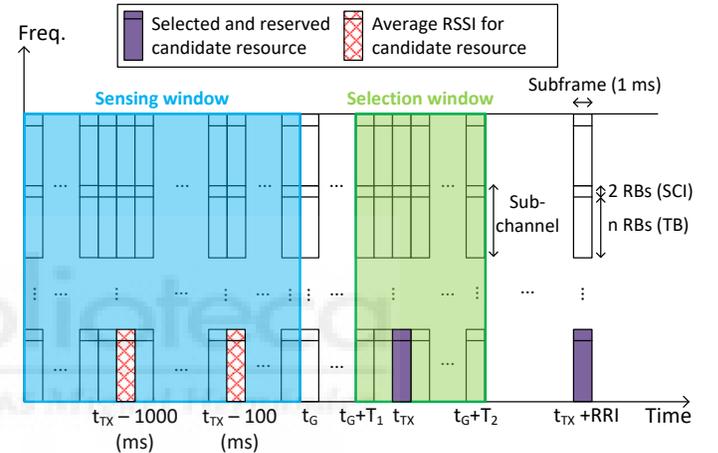


Fig. 2. LTE V2X channelization and illustration of mode 4 sensing-based SPS scheduling (illustrative example when $T=100$ ms).

B. Resource Allocation

LTE V2X defines two new resource allocation modes (mode 3 and mode 4) for V2X SL communications [16]. In mode 3, the cellular infrastructure (eNB) manages the V2X SL communications. This includes selecting and configuring the communication resources (sub-channels). Mode 4 can operate without cellular infrastructure support. In this case, vehicles autonomously select, manage and configure the sub-channels. Vehicles utilizing mode 3 need to be in network coverage, while vehicles using mode 4 can operate out of network coverage.

The standard does not specify an algorithm for the selection of sub-channels in mode 3. Instead, it defines two scheduling approaches [14]: dynamic scheduling and Semi-Persistent Scheduling (SPS). With dynamic scheduling, vehicles must request sub-channels from the eNB for each TB. With SPS scheduling, the eNB reserves sub-channels so that a vehicle can transmit several TBs. The eNB can configure the periodicity of the reserved sub-channels [23] using the DCI (Downlink Control Information) transmitted over the Physical Downlink Control Channel (PDCCH). Mode 3 can outperform mode 4 since the scheduling of transmissions is centralized at the eNB ([20], [21]). However, it requires operating in network coverage

and introduces cellular uplink (UL) and downlink (DL) signaling overhead. Mode 3 can also encounter challenges at the cell boundaries, in particular when different operators serve neighboring vehicles. 3GPP has defined an architecture to support multi-operator scenarios [22].

Under mode 4, vehicles autonomously select their sub-channels using the sensing-based SPS scheduling scheme specified in Rel. 14/15. A vehicle uses the selected sub-channels for the transmission of its following *Reselection Counter* consecutive TBs. The vehicle announces the reservation of the selected sub-channels for the transmission of the next TB using the *Resource Reservation Interval (RRI)* included in the SCI. This is illustrated in Fig. 2 where a vehicle selects sub-channel(s) at subframe t_{TX} , and informs neighboring vehicles that it reserves them for its following transmission at subframe $t_{TX} + RRI$. This is done to prevent other vehicles from utilizing the same sub-channels at the same time. The *RRI* can be equal to 0 ms, 20 ms, 50 ms, 100 ms or any multiple of 100 ms up to a maximum value of 1000 ms. A vehicle sets the *RRI* equal to 0 ms to announce neighboring vehicles that it is not reserving the same sub-channels for the next TB. A vehicle can only select *RRI*s values higher than 0 ms from a (pre-)configured list of permitted *RRI* values. This list can contain up to 16 values although currently 3GPP standards only define 12 possible *RRI*s values higher than 0 ms for mode 4 [19]. *Reselection Counter* is randomly chosen between 5 and 15 for a selected *RRI* higher than or equal to 100 ms. It varies between 10 and 30 for a selected *RRI* equal to 50 ms and between 25 and 75 for a *RRI* equal to 20 ms, respectively. *Reselection Counter* is decremented by one after transmitting a TB. A vehicle must select new sub-channels with probability $(1-P)$ when *Reselection Counter* is equal to zero, where $P \in [0, 0.8]$. The standard does not specify a fixed value for P . New sub-channels must also be selected if a new TB does not fit in the previously reserved sub-channels or if the current reservation cannot satisfy the latency deadline of a new TB. The later occurs if the time until the next reserved sub-channels is higher than the latency deadline of the new TB. Vehicles select new sub-channel(s) using the sensing-based SPS scheduling scheme that operates as follows ([23], [24]):

- 1) The vehicle first identifies candidate resources within a Selection Window (SW). A candidate resource is a group of adjacent sub-channels within the same subframe where the new SCI+TB to be transmitted fits (see Fig. 2). SW is a time window that includes the subframes in the range $[t_G+T_1, t_G+T_2]$, where t_G is the subframe at which the vehicle wants to select a new candidate resource. T_1 is the processing time (in subframes) required by a UE to identify and select candidate resources for transmission. T_1 is left to UE implementation but it must be equal to or smaller than 4 subframes. The value of T_2 (in number of subframes) is left to UE implementation but must be included within the range $T_{2min} \leq T_2 \leq 100$ if T_{2min} has been (pre-)configured. If T_{2min} has not been (pre-)configured, then the value of T_2 (in number of subframes) must be in the range $20 \leq T_2 \leq 100$. T_{2min} depends on the priority of the transmission and its value (in number of subframes) must be within the range

$10 \leq T_{2min} \leq 20$ [19]. Additionally, T_2 must be set so that t_G+T_2 fulfils the maximum latency requirement for transmitting the TB. The maximum latency is equal to 100 ms, 50 ms and 20 ms when vehicles transmit 10 packets per second (pps), 20 pps and 50 pps, respectively.

- 2) The vehicle then identifies candidate resources from SW that it should exclude. To do so, it senses the transmissions from other vehicles during the last 1000 subframes before t_G (*sensing window* in Fig. 2). If the vehicle was transmitting in any previous subframe t_N in the sensing window, it excludes all candidate resources within any subframe $t_N+q \cdot RRI_i$ located in the selection window where q is an integer value defined in [23]. RRI_i is equal to any value of the *RRI* (in number of subframes in the above expression) included in the (pre-)configured list of permitted *RRI*s. Note that if the list of permitted *RRI*s only includes values higher than or equal to 100 ms, then q can take only the value 1. These resources are excluded because the vehicle could not sense the transmissions from other vehicles at the same subframe it was transmitting a TB. We should note that other vehicles could be transmitting with any of the permitted values of RRI_i . Additionally, the vehicle should exclude all candidate resources within any subframe t_F of the selection window if it estimates that any of its following j transmissions at $t_F+j \cdot RRI_{TX}$ (with $1 \leq j \leq 10 \cdot ReselectionCounter-1$ and RRI_{TX} being the *RRI* selected by the vehicle) can coincide with any subframe $t_N+q \cdot RRI_i$. Again, if the list of permitted *RRI*s only includes values higher than or equal to 100 ms, then q can take only the value 1. The vehicle also excludes the candidate resources that it estimates will be used by other vehicles. A candidate resource is estimated to be used by other vehicles, and hence excluded, if the two following conditions are met: 1) the vehicle estimates with an SCI received from another vehicle that this other vehicle will utilize this candidate resource in the current SW or at the same time the vehicle will need it to transmit any of its following $10 \cdot ReselectionCounter-1$ TBs; 2) the vehicle excludes a candidate resource if its average Reference Signal Received Power (RSRP) measured over the TB associated to the corresponding SCI is higher than a threshold that can be configured. The vehicle checks then if the number of remaining available candidate resources is equal or higher than 20% of all candidate resources within the SW. If this is not the case, the RSRP threshold is increased by 3 dB and the process to identify available candidate resources is repeated iteratively until the number of available candidate resources is at least equal to 20% of the candidate resources within SW. The vehicle creates a list L_1 with the identified available candidate resources.
- 3) The vehicle creates a second list L_2 with the candidate resources from L_1 experiencing the lowest average RSSI (Received Signal Strength Indicator) in the sensing window. The total number of candidate resources in L_2 must be greater than or equal to 20% of all candidate resources in the SW. The RSSI of a candidate resource located at subframe t_{TX} is computed by averaging the RSSI

over all the RBs of the candidate resource. The average RSSI is computed over the previous $t_{TX}-T*j$ subframes located in the sensing window, where j is a positive integer (note that T is in number of subframes in the above expression). If the vehicle has selected an $RRI \geq 100$ ms, then $T=100$ ms. When the RRI is set equal to 50 ms then $T=50$ ms and when $RRI=20$ ms then $T=20$ ms.

- 4) The vehicle randomly chooses one of the candidate resources from L_2 to reduce the probability that multiple vehicles select the same candidate resource or candidate resources that partially overlap. The selected candidate resource is used by the vehicle to transmit its new TB, and the vehicle maintains the selection for its next *ReselectionCounter*-1 transmissions.

The operation of LTE V2X mode 4 depends on parameters (e.g., P and RSRP threshold) that can be (pre-)configured. An analysis of their configuration can be found in [20] and [25].

Since LTE V2X mode 4 utilizes a sensing-based scheduling mechanism, it is prone to packet collisions as the network load and channel congestion increase [26]. 3GPP defines two metrics to characterize the channel congestion: Channel Busy Ratio (CBR) and Channel occupancy Ratio (CR). The CBR is defined as the ratio of sub-channels that experience an RSSI higher than a (pre-)configured threshold to the total number of sub-channels in the previous 100 subframes. The CR quantifies the channel occupancy generated by the transmitting vehicle. It is estimated in subframe n as the ratio between the number of sub-channels utilized by the transmitting vehicle in subframes $[n-a, n-1]$ and selected by the vehicle for its remaining *Reselection Counter* transmissions in subframes $[n, n+b]$ and the total number of sub-channels within $[n-a, n+b]$. a and b must satisfy $a+b+1 = 1000$ with $a \geq 500$. The standard defines up to sixteen CBR ranges. For each range, the standard specifies a CR_{Limit} that cannot be surpassed by a transmitting vehicle and that can take different values as a function of the priority of the transmission. When a vehicle wants to transmit a TB, it measures the CBR and maps it to one of the ranges to get the CR_{Limit} . The vehicle also estimates its CR and if it is higher than the CR_{Limit} , it adjusts its transmission parameters. For example, the vehicle could drop certain packets (while maintaining the reserved sub-channels) or reduce the transmission power to decrease the CBR measured by other vehicles and maintain its CR. A vehicle could also augment the MCS to utilize less sub-channels to transmit a TB and hence reduce its CR. 3GPP does not specify a particular congestion control mechanism. However, first studies have analyzed the performance achieved with the Decentralized Congestion Control (DCC) technique defined by ETSI [27] or the distributed congestion control mechanism specified in the SAE J2945/1 standard [28].

III. 5G NR V2X USE CASES

5G NR V2X has been designed to complement LTE V2X. LTE V2X supports basic active safety and traffic management use cases while 5G NR V2X supports advanced use cases and higher automation levels. The 5G NR V2X use cases have been specified by 3GPP Services and System Aspects (SA) Working

Group 1 (SA1) and have been further elaborated by the 5GAA Working Group 1 (Use Cases and Technical Requirements). Both 3GPP and 5GAA organize use cases in groups, and a use case can be a member of more than one group. This section provides an overview of the use cases supported by 5G NR V2X along with their main KPIs. Readers are referred to 3GPP [11] and 5GAA [13] documents or papers on the topic [29] for their comprehensive description.

A. 3GPP use case groups

3GPP technical report (TR) 22.886 [11] and TS 22.186 [12] present a comprehensive description of the NR V2X use cases and requirements, respectively. For each use case, 3GPP further distinguishes different degrees of automation following the SAE automation levels [30] ranging from 0 (no automation) to 5 (full automation). Typically, the higher automation levels of a use case, the more stringent the NR V2X QoS requirements are. The use cases are divided in the following four groups [11]:

- 1) *Vehicles Platooning*: This group includes use cases for the dynamic formation and management of groups of vehicles in platoons. Vehicles in a platoon exchange data periodically to ensure the correct functioning of the platoon. The inter-vehicle distance between vehicles in a platoon may depend on the available QoS.
- 2) *Advanced Driving*: This group includes use cases enabling semi-automated or fully-automated driving. In this group, vehicles share data obtained from their local sensors with surrounding vehicles in proximity. In addition, vehicles share their driving intention in order to coordinate their trajectories or maneuvers, thus increasing safety and improving traffic efficiency.
- 3) *Extended Sensors*: This group enables the exchange of sensor data – either raw or processed – collected through local sensors between vehicles, RSUs, devices of pedestrians, and V2X application servers. The objective is to improve the perception of the environment beyond the perception capabilities of the vehicles' own sensors.
- 4) *Remote Driving*: This group enables a remote (tele-operated) driver or a V2X application to operate a vehicle. The main use cases are for passengers who cannot drive themselves, for vehicles located in hazardous environments (e.g., construction areas or locations with adverse weather conditions), and for complex situations which automated vehicles are unable to drive safely.

B. 5GAA use case groups

5GAA combines the use cases and groups defined by 3GPP in Rel. 14 [31] and Rel. 15 [11] with new use cases in Rel. 16, and defines the following groups [13]:

- 1) *Safety*: This group includes use cases that provide safety for vehicles and other traffic participants. It includes basic safety use cases such as emergency braking, collision warning, etc. [31], along with more advanced use cases requiring higher automation levels (e.g., intersection management).
- 2) *Vehicle operations management*: This group comprises commercial use cases aimed at improved operation of

TABLE II
REQUIREMENT RANGES FOR THE 3GPP USE CASE GROUPS [12]

Use case group	Payload (Bytes)	Tx rate (Message/ Sec)	Max end-to-end latency (ms)	Reliability (%)	Data rate (Mbps)	Required communication range (meters)
Vehicles Platooning	50-6000	2-50	10-25	90-99.99	<= 65	80-350
Advanced Driving	SL: 300-12000 UL: 450	SL: 10-100 UL: 50	10-100	90-99.999	SL: 10-50 UL: 0.25-10 DL: 50	360-700
Extended Sensors	1600	10	3-100	90-99.999	10-1000	50-1000
Remote Driving	16000-41700	33-200	5	99.999	UL: 25 DL: 1	1000+

Note 1: If not specified otherwise, the requirement applies to all link types (SL, DL, and UL).
Note 2: In case of the Remote Driving use case group, [12] does not specify the values for Payload, Tx rate, and Required communication range. For completeness, we include these missing values based on [127].

vehicles for diverse types of users (e.g., vehicle owners/drivers, transport/delivery companies, etc.). Examples include sensors monitoring, software updates, remote support, etc.

- 3) *Convenience*: This group includes a diverse set of use cases that provide value and convenience to either the driver or the fleet management company operating the vehicle. Examples include infotainment, assisted navigation, and smart parking.
- 4) *Autonomous driving*: This group includes the advanced driving, remote driving, and extended sensors groups defined by 3GPP.
- 5) *Platooning*: This group is the same as the 3GPP vehicles platooning use case group.
- 6) *Traffic efficiency and environmental friendliness*: This group includes use cases that provide enhanced value to infrastructure or city providers in areas where the vehicles will be operating. Examples include Green Light Optimal Speed Advisory (GLOSA), traffic jam information, routing advice, etc.
- 7) *Society and community*: This group includes use cases that are of value and interest to the society and public. Examples in this group are Vulnerable Road User (VRU) protection, emergency vehicle approaching, emergency answering points, etc.

C. Use case requirements

1) 3GPP requirements

3GPP WG SA1 defined in Rel. 14 [32] and Rel. 15 [12] the key requirements for V2X services as follows:

- 1) *Payload* refers to the amount of data required by a specific service and generated by the application.
- 2) *Tx rate* is the number of messages per unit time that the transmitter generates and the receiver is expected to receive subject to other relevant requirements (e.g., payload size, latency, communication range, etc.).
- 3) *Maximum end-to-end latency* is the maximum allowed time between the generation of a message at the transmitter's application and the reception of the message at the receiver's application.
- 4) *Reliability* is defined as the probability that a transmitted message is correctly received within a specified maximum

end-to-end latency subject to other relevant requirements (e.g., payload size, communication range, etc.).

- 5) *Data rate* represents the total amount of data that needs to be received by the receiver per unit time. It is related directly to the payload and the *Tx rate* and is measured in bits per second (bps). It is also subject to other requirements (e.g., latency, reliability etc.).
- 6) *Required communication range* specifies the minimum distance between a transmitter and its intended receiver allowing communication with a targeted payload size, maximum latency, reliability, and data rate.

Table II summarizes the range of values for the above requirements as identified by 3GPP [12] for the four 3GPP use case groups. Each range is defined by the minimum and maximum requirements based on [12]. Use cases typically do not need to meet all of the most stringent requirements simultaneously. For example, platooning might require relatively low latency and high reliability, but the required data rate is moderate (exchange of maneuver information among the platoon members). We note that higher degrees of automation generally lead to more stringent requirements.

2) 5GAA requirements

For each use case, 5GAA defines multiple possible [13] use case scenarios, where the scenarios differ in terms of road configuration, actors involved, service flows, etc. 5GAA complements largely network-centric requirements laid out by 3GPP with the concept of Service Level Requirements (SLRs). SLRs focus more on the automotive-centric requirements (e.g., service level reliability, interoperability, positioning, etc.), and consider system-level aspects (e.g., vehicle density) that need to be supported by the network for a specific use case scenario. SLRs are established for each user story defined within a use case. A user story describes, from the user's perspective, a use case scenario that results in a specific SLR. Key SLRs are defined in Table III.

5GAA has so far not performed an exhaustive analysis of service level requirements for all use cases. Rather, it provided in [13] some examples of use cases, scenarios, and user stories for each use case group. Example of 5GAA use cases and user stories descriptions can be found in [13].

TABLE III
SERVICE LEVEL REQUIREMENTS DEFINED BY 5GAA [13]

Service level Requirement (SLR)	SLR Unit
Range: same as 3GPP definition of required communication range.	[m]
Information requested/generated: Information that a service needs or creates.	Quality of information / Information needs
Service Level Latency: same as 3GPP definition of maximum end-to-end latency.	[ms]
Service Level Reliability: probability that at least one of a set of messages carrying the same information from the transmitter (i.e., including retransmissions on any of the layers) is correctly received within a specified maximum latency, subject to other relevant requirements.	[%]
Velocity: maximum speed of the user (vehicle, pedestrian, etc.) that is required by the use case story.	[m/s]
Vehicle Density: maximum density of vehicles in an area required by the use case story.	[vehicle/km ²]
Positioning: maximum allowed error in estimating the location of a user (vehicle, pedestrian, etc.) required by the use case story.	[m]
Interoperability/ Regulatory/ Standardization Required (yes/no): requirements by the use case story in terms of: 1) interoperability between users; 2) regulatory action (e.g., in terms of spectrum, safety, etc.); and 3) standardization of the use case story for its functioning.	Not applicable

IV. OVERVIEW OF 5G SYSTEM ARCHITECTURE FOR V2X COMMUNICATION OVER PC5 AND UU

The 5G system architecture supports two operation modes for V2X communication, namely V2X communication over the PC5 reference point or interface and V2X communication over the Uu reference point or interface. The PC5 interface supports

SL V2X communications for NR and LTE. V2X communications over Uu for UL and DL transmissions are possible under NR Non-Standalone (NSA) and Standalone (SA) deployments. In Rel. 16, V2X communication over Uu is only supported for unicast communications. However, Rel. 17 includes an ongoing SI [33] to discuss enhancements for multicast and broadcast transmissions in 5G. It should be noted that V2X messages over LTE-Uu can be broadcast in DL via Multimedia Broadcast Multicast Services (MBMS) [34]. Since NR V2X includes both the SL (over the PC5 interface) and UL/DL (over the Uu interface), when referring specifically to either the SL or the UL/DL, in the rest of the paper we name them explicitly (e.g., NR V2X SL).

Fig. 3 shows the high-level view of the 5G System architecture for V2X communication over PC5 and Uu interface. More specifically, Fig. 3 illustrates the roaming architecture with local breakout. Home Public Land Mobile Network (HPLMN) refers to the network that a user is a subscriber to while visited PLMN (VPLMN) is the network to which the UE roams to when leaving the HPLMN. When a UE roams, services provided by the HPLMN are used to provide V2X service related parameters to the VPLMN. The local breakout is a deployment option where the Session Management Function (SMF) for establishment, modification or release of a session and all User Plane Functions (UPFs) that are involved in a Protocol Data Unit (PDU) session (i.e., a logical connection between the UE and network) are under control of the VPLMN. Local breakout is critical to reduce the latency. V2X communication over PC5 supports roaming and inter-PLMN operations. In case of inter-PLMN V2X communication over PC5, the PC5 parameters must be

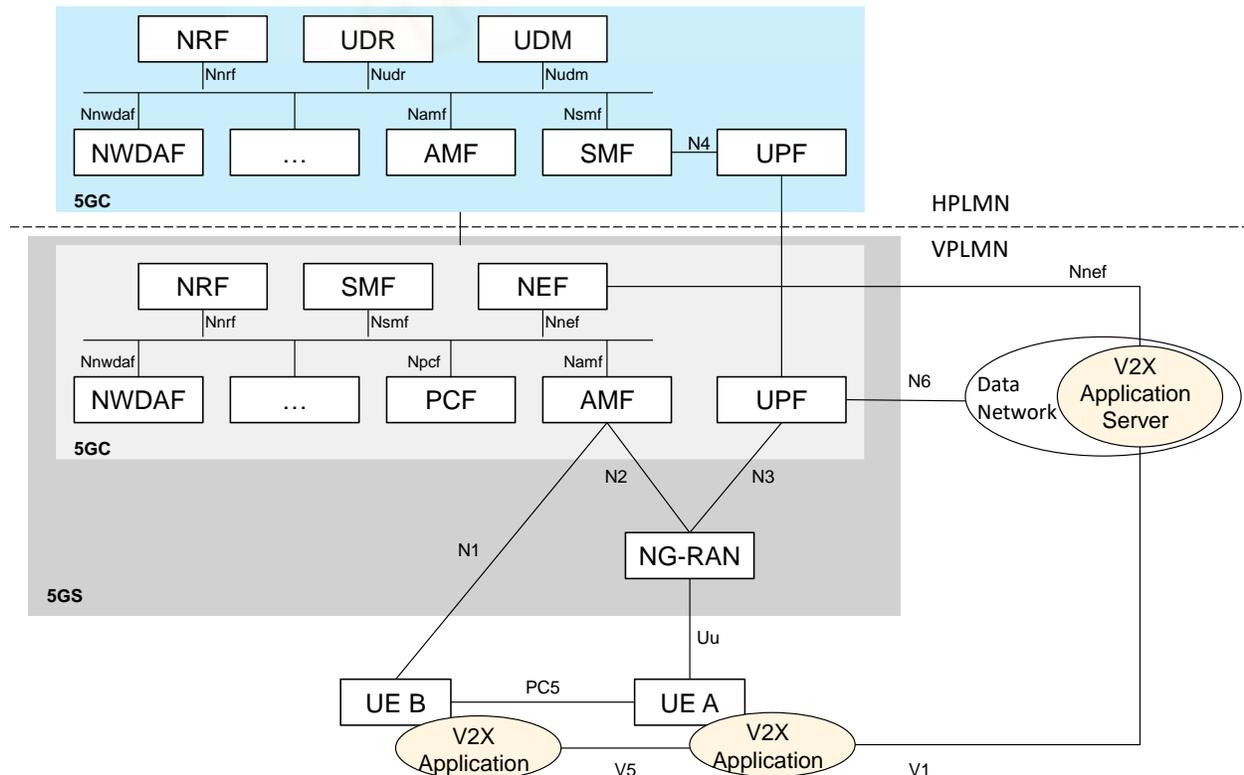


Fig. 3. 5G System architecture for V2X communication over PC5 and Uu – Local breakout scenario (Roaming case).

configured consistently among the UEs within a certain geographical area.

5G System (5GS) consists of the Next-Generation Radio Access Network (NG-RAN) and the 5G Core network (5GC) domains. The 5GC consists of several Network Functions (NF) such as the Access and Mobility Management Function (AMF), Policy Control Function (PCF), Network Data Analytics Function (NWDAF), Network Repository Function (NRF), Network Exposure Function (NEF), Unified Data Repository (UDR), Unified Data Management (UDM), UPF and SMF among others. It should be noted that in Fig. 3 the NFs within the 5GC Control Plane (that includes all NFs except UPF) use service-based interfaces for their interactions (e.g., Npcf, Nsmf, Nnef, Nnwdaf, etc). In service-based representation, NFs within the Control Plane enable other authorized network functions to access their services. Fig. 3 also shows the application layer interfaces between V2X applications in the UEs (interface V5) and the interface between V2X applications in the UE and in the V2X Application Server (AS) (interface V1). The application layer interfaces can be used to exchange application layer information and configuration parameters useful for the UE to configure its V2X communication. The 5GS architecture specified in [35] does not introduce new NFs to support V2X communication. Instead, existing NFs have been extended with V2X related functionalities. Next, we present a non-exhaustive list of these extensions at the 5GS and the UE.

The UE can report the V2X capability and V2X PC5 capability information to the 5GC (as well as receive V2X parameters from the 5GC) over the N1 reference point and/or from a V2X AS over the V1 reference point (see Fig. 3). The UE includes procedures for V2X communication over PC5 and for the configuration of parameters (e.g., destination Layer-2 IDs, radio resource parameters, V2X AS address information, mapping between V2X service types and V2X frequencies). A UE can receive the parameters for its V2X communications over PC5 and Uu from different sources. In this case, the UE shall consider them in the following priority order: a) provided/updated by the PCF through the AMF using Non-access stratum (NAS) signaling; b) provided/updated by the V2X AS via the V1 reference point; c) (pre-)configured in the Universal integrated circuit card (UICC); and d) pre-configured in the UE. The PCF can provide the UE with authorization and policy parameters for V2X communication over PC5 and Uu, for example, radio parameters when the UE is not served or under coverage by LTE or 5G. The PCF also provides the AMF with necessary parameters that are part of the UE context to configure and manage V2X communications. This includes, for example, parameters related to PC5 QoS flows and profiles. These parameters are retrieved from a UDR. The NRF helps other NFs to discover and select the appropriate PCF taking into account the V2X capabilities.

The 5GS provides NEF services to enable communication between NFs in the PLMN and the V2X AS. In addition, the NEF stores the V2X service parameters in the UDR. The V2X AS includes AF (Application Function) functionalities such as: request 5GC services (e.g., notifications about expected changes of provided QoS, QoS Sustainability Analytics),

provision the UE and/or the 5GC with parameters for V2X communications over PC5 and/or Uu (e.g., mapping of V2X service types to V2X frequencies with geographical areas).

The interworking between 5GS V2X and 4G Evolved Packet System (EPS) V2X does not require new interfaces at the architecture level and does not impact existing network function entities in the 4G Evolved Packet Core (EPC) network entities and the 5GC entities. When the UE is in 5GS or EPS, the UE shall use the valid V2X policy and parameters provisioned by the PCF in 5GC or by the V2X Control Function (CF) in EPC for V2X communication. The V2X related parameters for EPS are defined in TS 23.285 [34]. They can be provided by the PCF or by the V2X CF while the V2X policy and parameters for 5GS are provided by the PCF only.

V. PHYSICAL LAYER DESIGN FOR NR V2X SIDELINK

The physical layer structure for the NR V2X sidelink is based on the Rel. 15 NR Uu design. In addition, the physical layer procedures for the NR V2X SL reuse some of the concepts of Rel. 14 LTE V2X, with the introduction of additional procedures for providing physical layer support for unicast and groupcast transmissions.

A. Physical Layer Structures for NR V2X sidelink

1) Numerology

Rel. 16 NR V2X sidelink can operate at the same frequencies as Rel. 15 NR Uu [10], [36], i.e., at frequencies within the two following frequency ranges [37], [38]:

- Frequency range 1 (FR1): 410 MHz – 7.125 GHz.
- Frequency range 2 (FR2): 24.25 GHz – 52.6 GHz.

Although both frequency ranges are supported in NR V2X sidelink, the design of NR V2X sidelink has been based mainly on FR1 [10]. For NR V2X sidelink, no specific optimization is performed for FR2 [10], except for addressing phase noise which is more prominent at higher frequencies [39].

Transmissions in NR V2X SL use the orthogonal frequency division multiplexing (OFDM) waveform with a cyclic prefix (CP). The sidelink frame structure is organized in radio frames (also referred simply as frames), each with a duration of 10 ms. A radio frame is divided into 10 subframes, each with a duration of 1 ms. The number of slots per subframe and the subcarrier spacing (SCS) for the OFDM waveform can be flexible for NR V2X. To support diverse requirements and different operating frequencies in FR1 and FR2, a scalable OFDM numerology is considered for NR V2X based on Rel. 15 NR Uu. Each OFDM numerology is defined by an SCS and a CP. NR V2X supports multiples of 15 kHz (i.e., the SCS in LTE V2X) for the SCS of the OFDM waveform [36]. As shown in Table IV, different OFDM numerologies can be obtained with a scalable SCS given by $2^\mu \times 15$ kHz, where μ is an SCS configuration factor. For NR V2X, the SCS configuration factor can be $\mu = 0, 1, 2, 3$ such that the SCS can be equal to 15 kHz, 30 kHz, 60 kHz or 120 KHz. In FR1, 15 kHz, 30 kHz and 60 kHz are supported for the SCS, while 60 kHz and 120 kHz are supported for the SCS in FR2. Supporting higher SCS improves robustness of the OFDM waveform against frequency impairments caused by Doppler effects, carrier frequency offsets and hardware phase

TABLE IV
SUPPORTED NUMEROLOGIES IN NR V2X SIDELINK

μ	SCS ($2^\mu \times 15$ kHz)	Frequency Range	Cyclic Prefix	Symbols per slot	Slots per subframe (2^μ)	Slot duration ($2^{-\mu}$ ms)	Maximum Carrier Bandwidth
0	15 kHz	FR1	Normal	14	1	1 ms	50 MHz
1	30 kHz	FR1	Normal	14	2	0.5 ms	100 MHz
2	60 kHz	FR1, FR2	Normal	14	4	0.25 ms	200 MHz
			Extended	12			
3	120 kHz	FR2	Normal	14	8	0.125 ms	400 MHz

noise [41], which are more prominent in FR2.

In NR V2X SL, the number of slots in a subframe is equal to 2^μ , i.e., it is determined by the SCS configuration factor μ . As the subframe duration is 1 ms, the slot duration is given by $2^{-\mu}$ ms. Thus, a larger SCS results in a shorter slot duration. As in Rel. 15 NR Uu, two different CP lengths are supported in NR V2X sidelink: a first CP referred to as normal CP and a longer CP referred to as extended CP. A normal CP can be used for all supported SCS while the extended CP is supported only for an SCS of 60 kHz [36]. With the normal CP, increasing the SCS from 15 kHz to 60 kHz leads to a shorter CP (similar to how the slot duration is reduced by a factor of 4). However, this reduced normal CP with an SCS of 60 kHz may not be enough to accommodate for the delay spread in certain deployments and channels. For this purpose, Rel. 15 NR Uu introduced the extended CP for an SCS of 60 KHz, which has a similar duration as the normal CP for an SCS of 15 KHz [41]. Depending on whether a normal CP or an extended CP is used, each slot consists of 14 or 12 OFDM symbols, respectively. The duration of the normal or extended CP as well as the resulting OFDM symbol duration scale with $2^{-\mu}$, i.e., decreasing with increasing SCS. The exact expressions for the CP duration² and the resulting OFDM symbol duration is given in [36]. In the following, we refer to OFDM symbols simply as symbols.

As in NR Uu, the maximum bandwidth in NR V2X SL depends on the SCS ([37],[38]), see Table IV. Only one numerology (i.e., one combination of SCS and CP) can be used in a carrier at a time in NR V2X SL. As a general design guideline, the selection of the numerology depends on aspects like the carrier frequency, radio channel conditions (e.g., due to mobility), requirements (e.g., latency) and hardware features (e.g., complexity). For example, a large SCS may be preferred for low latency applications since a larger SCS leads to a shorter slot duration.

The smallest unit of time for scheduling SL transmissions in NR V2X is a slot. This is in contrast to Rel. 15 NR Uu which supports mini-slot scheduling, where a data transmission can be scheduled on only some of the OFDM symbols within a slot³ [42]. NR V2X does not support mini-slot scheduling in the SL.

2) Sidelink Bandwidth Parts

In NR Uu, the maximum carrier bandwidth is 200 MHz for FR1 and 400 MHz in FR2. Although gNBs can support such

wide bandwidths, this may not be the case for all UEs, in particular low-end UEs. Furthermore, supporting a very large bandwidth may also imply higher power consumption at the UE, both from the radio frequency (RF) and baseband signal processing perspectives. To support UEs that cannot handle large bandwidths (e.g., due to processing limitations or high power consumption), Rel. 15 NR Uu introduced the concept of bandwidth part (BWP) [43]. A BWP consists of a contiguous portion of bandwidth within the carrier bandwidth where a single numerology is employed. By defining a small BWP, the computational complexity and power consumption of a UE can be reduced. As each BWP can have a different bandwidth and numerology, BWPs enable a more flexible and efficient use of the resources by dividing the carrier bandwidth for multiplexing transmissions with different configurations and requirements.

The concept of BWP has also been adopted for the NR V2X sidelink ([44], [45]) where a SL BWP occupies a contiguous portion of bandwidth within a carrier as depicted in Fig. 4. In a carrier, only one SL BWP⁴ is (pre-)configured for all UEs [46]. Sidelink UE transmissions and receptions are contained within the SL BWP and employ the same numerology. Thus, all physical channels, reference signals and synchronization signals in NR V2X sidelink are transmitted within the SL BWP. This also means that in the sidelink a UE is not expected to receive or transmit in a carrier with more than one numerology [45]. The SL BWP is divided into common RBs. A common RB consists of 12 consecutive subcarriers with the same SCS, where the SCS is given by the numerology of the SL BWP.

3) Resource Pools

In NR V2X, only certain slots are (pre-)configured to accommodate SL transmissions. Thus, the available sidelink resources consist of slots allocated for sidelink (time resources) and common RBs within a SL BWP (frequency resources). In NR V2X, a subset of the available SL resources is (pre-)configured to be used by several UEs for their SL transmissions. This subset of available SL resources is referred to as a resource pool [45] and is illustrated in Fig. 4. The common resource blocks within a resource pool are referred to as physical resource blocks (PRBs)⁵. A resource pool consists of contiguous PRBs and contiguous or non-contiguous slots that have been (pre-)configured for SL transmissions. A resource pool must be defined within the SL BWP [47].

²For the extended CP, all symbols have the same CP duration. For the normal CP, the first symbol every 0.5 ms has a slightly longer CP than the CP in the rest of the symbols [40].

³Mini-slot scheduling can be beneficial for latency critical transmissions and shorter data transmissions, e.g., ultra-reliable low latency communication (URLLC).

⁴This is in contrast to NR Uu where up to four BWPs can be configured for the downlink and uplink of a UE (although only one BWP is active for the downlink and one BWP is active for the uplink) [36].

⁵The PRBs are indexed within a resource pool while the common RBs are indexed within the SL BWP.

Therefore, a single numerology is used within a resource pool. If a UE has an active UL BWP, the SL BWP must use the same numerology as the UL BWP if they are both included in the same carrier [48]. Otherwise, the SL BWP is deactivated [48].

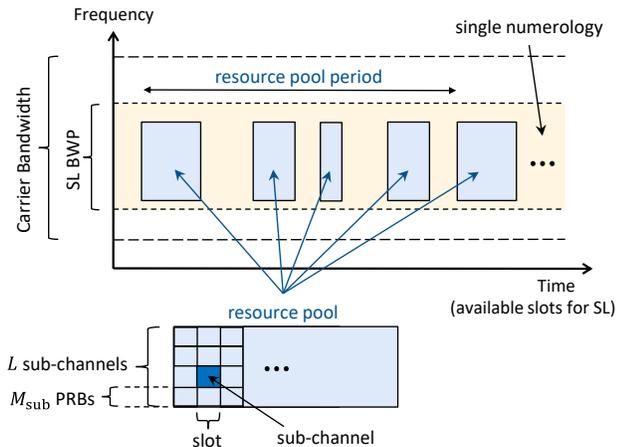


Fig. 4. SL bandwidth part and resource pool for NR V2X sidelink.

In the frequency domain, a resource pool is divided into a (pre-)configured number L of contiguous sub-channels, where a sub-channel consists of a group of consecutive PRBs in a slot [42]. The number M_{sub} of PRBs in a sub-channel corresponds to the sub-channel size, which is (pre-)configured within a resource pool. In NR V2X SL, the sub-channel size M_{sub} can be equal to 10, 12, 15, 20, 25, 50, 75, or 100 PRBs. A sub-channel represents the smallest unit for a sidelink data transmission or reception. A sidelink transmission can use one or multiple sub-channels. In the time domain, the slots that are part of a resource pool are (pre-)configured and occur with a periodicity of 10240 ms [46]. The slots that are part of a resource pool can be (pre-)configured with a bitmap [49]. The length of the bitmap can be equal to 10, 11, 12, ..., 160 [46].

At each slot of a resource pool, only a subset of consecutive symbols are (pre-)configured for the sidelink [42], i.e., out of the 14 or 12 symbols per slot for a normal or extended CP, respectively. The subset of SL symbols per slot is indicated with a starting symbol and a number of consecutive symbols, where these two parameters are (pre-)configured per resource pool [55]. The number of consecutive SL symbols⁶ can vary between 7 and 14 symbols [42], e.g., depending on the physical channels which are carried within a slot (e.g., see Fig. 5 and Fig. 6).

A resource pool (RP) can be shared by several UEs for their SL transmissions. An RP can be used for all transmission types (i.e., unicast, groupcast, and broadcast). A UE can be (pre-)configured with multiple RPs for transmission (transmit RPs) and with multiple RPs for reception (receive RPs) [45]. A UE can then receive data on resource pools used for SL transmissions by other UEs, while the UE can still transmit on the SL using its transmit resource pools. For the case when UEs in network coverage do not have a stable network connection, exceptional transmit RPs are configured for the UEs [45]. These

situations include when a UE is in a transition from idle to connected mode, when a UE experiences a link failure or a handover, or when a UE is changing between different configured transmit RPs [45]. The use of exceptional transmit RPs in such situations aids in improving service continuity [50].

B. Physical Channels and Signals in NR V2X sidelink

The physical channels specified in NR V2X SL [36]:

- 1) *Physical Sidelink Control Channel (PSCCH)*: carries control information in the sidelink.
- 2) *Physical Sidelink Shared Channel (PSSCH)*: carries data payload in the sidelink and additional control information.
- 3) *Physical Sidelink Broadcast Channel (PSBCH)*: carries information for supporting synchronization in the sidelink. PSBCH is sent within a sidelink synchronization signal block (S-SSB).
- 4) *Physical Sidelink Feedback Channel (PSFCH)*: carries feedback related to the successful or failed reception of a sidelink transmission.

Furthermore, the following signals (carried in or associated with the previous channels) are specified in NR V2X [36]:

- 1) *Demodulation reference signal (DMRS)*: used by a receiver for decoding the associated sidelink physical channel, i.e., PSCCH, PSSCH, PSBCH. The DMRS is sent within the associated sidelink physical channel.
- 2) *Sidelink primary synchronization signal (S-PSS) and sidelink secondary synchronization signal (S-SSS)*: used by a receiver to synchronize to the transmitter of these signals. S-PSS and S-SSS are sent within an S-SSB.
- 3) *Sidelink Channel state information reference signal (SL CSI-RS)*: used for measuring channel state information (CSI) at the receiver that is then fed back to the transmitter. The transmitter can adjust its transmission based on the fed back CSI. SL CSI-RS is sent within the PSSCH region of the slot.
- 4) *Sidelink Phase-tracking reference signal (SL PT-RS)*: used for mitigating the effect of phase noise (in particular at higher frequencies) resulting from imperfections of the oscillator. SL PT-RS is sent within the PSSCH region of the slot.

Compared to LTE V2X that supports only broadcast transmissions in the sidelink, NR V2X provides physical layer support for unicast, groupcast, and broadcast transmissions in the SL. In the following, we refer to a UE performing a transmission as a TX UE while the intended recipients of a transmission are denoted as RX UEs. In NR V2X SL, data is organized into TBs [42] and each TB is associated with a SCI. A TB is carried in a PSSCH. The SCI indicates the resources used by the PSSCH that carries the associated TB, as well as further information required for decoding the TB. A PSCCH is sent with a PSSCH. The SCI in NR V2X is transmitted in two stages [45] compared to a single one for LTE V2X. The 1st-stage SCI in NR V2X is carried on the PSCCH while the 2nd-stage SCI is carried on the corresponding PSSCH. As discussed in Section V.B.2), the introduction of the 2nd-stage SCI enables

⁶Differently from mini-slot scheduling from NR Uu, the number of symbols configured for a SL transmission in NR V2X is not changed dynamically.

a flexible SCI design to support unicast, groupcast, and broadcast transmissions in NR V2X, in contrast to LTE V2X where only broadcast is supported. Splitting the SCI in two stages (1st-stage SCI and 2nd-stage SCI) allows other UEs which are not RX UEs of a transmission to decode only the 1st-stage SCI for channel sensing purposes, i.e., for determining the resources reserved by other transmissions. On the other hand, the 2nd-stage SCI provides additional control information which is required for the RX UE(s) of a transmission. UEs performing channel sensing are referred in the following as sensing UEs.

NR V2X introduces the use of feedback in unicast and groupcast communications in order to increase the reliability of SL transmissions. The feedback consists of hybrid automatic repeat request (HARQ) feedback. Two options for sending the SL HARQ feedback are possible for groupcast communication: option 1 and option 2. In option 1, only RX UEs at a specified distance to the TX UE (smaller than the required communication range) should send HARQ feedback. The required communication range (see Section III.C.1)) depends on the service and represents the minimum distance for which the QoS parameters of a service need to be fulfilled [51]. In option 2, all RX UEs send HARQ feedback. The RX UEs send the HARQ feedback on PSFCH in response to a unicast or groupcast transmission carried in a PSSCH. Further details about the HARQ procedure are provided in Section V.C.1). NR V2X also supports CSI reporting in unicast communications. To this aim, a TX UE transmits CSI-RS so that a RX UE can measure the CSI and report it back to the TX UE via CSI reporting carried within a PSSCH.

In NR V2X, a UE can also transmit information for supporting synchronization in the sidelink. If it does, the UE serves as a synchronization reference and is referred to as a SyncRef UE [52]. The synchronization information in NR V2X SL is carried on the S-SSB that consists of the PSBCH, S-PSSS and S-SSS. Nearby UEs that may be out of network or GNSS coverage can receive S-SSB transmissions from a SyncRef UE and synchronize to it. Thus, nearby UEs can then have the same SL timing reference and establish SL communication to/from the SyncRef UE and among nearby UEs. It should be noted that a SyncRef UE is not always a TX UE, i.e., that transmits data.

PSCCH is sent on the same slot as the associated PSSCH, and such a slot is referred to as a PSCCH/PSSCH slot. While PSCCH and PSSCH can be sent in every slot of a resource pool, resources for PSFCH are allocated periodically, with a transmission on PSFCH performed in response to a PSSCH received a few slots before. Furthermore, S-SSBs are transmitted periodically within a SL BWP on slots which are not part of a resource pool. The following sub-sections present in more detail the physical sidelink channels, the S-SSB, the associated signals and the two stage SCI design. A summary is presented afterwards in Table VI.

1) Physical Sidelink Control Channel (PSCCH)

The PSCCH carries the 1st-stage SCI that contains control information associated with a PSSCH and the 2nd-stage SCI. For

this purpose, SCI format 1-A is used [53]. The 1st-stage SCI indicates the frequency resources (e.g., sub-channels) of the PSSCH carrying the current (re)transmission of a TB, as well as the resource reservation for up to two further retransmissions of the TB. The 1st-stage SCI also informs about the resource reservation period if the UE reserves resources semi-persistently for PSSCH (see Section VI.B). In addition, the 1st-stage SCI includes the priority of the associated PSSCH, as well as the format and size of the 2nd-stage SCI [53]. The 1st-stage SCI also indicates the MCS of the data payload carried in the associated PSSCH. The MCS is determined with an MCS index within an MCS table. In NR V2X SL, a specific MCS table is used by default and additionally up to two further MCS tables can be (pre-)configured per resource pool [54]. The MCS tables are selected among the three MCS tables that are supported for the physical downlink shared channel (PDSCH) in Rel. 15 NR Uu [42]. To indicate the MCS of a PSSCH, the 1st-stage SCI indicates the MCS index as well as the MCS table, if one or two MCS tables are (pre-)configured within the resource pool. For supporting different channel conditions, the DMRS associated with a PSSCH in NR V2X can be carried on different symbols within a PSSCH slot, i.e., with different time patterns. Within a resource pool, multiple time patterns can be (pre-)configured for the PSSCH DMRS, and the 1st-stage SCI indicates which time pattern is used for the associated PSSCH. The 1st-stage SCI also provides the number of ports of the PSSCH DMRS, which can be equal to one or two. This represents the number of layers (i.e., number of data streams) supported in the PSSCH. Thus, by exploiting multiple transmit and receive antennas up to two streams of data can be sent within a PSSCH in NR V2X SL.

The PSCCH is multiplexed in non-overlapping resources with the associated PSSCH in the same slot. The PSCCH is transmitted from the second SL symbol in the slot and starting from the lowest PRB within the sub-channel(s) occupied by the associated PSSCH. The number of symbols for the PSCCH is (pre-)configured per resource pool and can be equal to 2 or 3 symbols. In the frequency domain, the PSCCH occupies a (pre-)configurable number M_{PSCCH} of PRBs per resource pool that can be equal to 10, 12, 15, 20 or 25 PRBs [55]. However, as the PSCCH is to be contained within one sub-channel [56], the number M_{PSCCH} of PRBs for the PSCCH is limited by the number M_{sub} of PRBs in a sub-channel⁷, i.e., $M_{\text{PSCCH}} < M_{\text{sub}}$. The possible number of symbols and PRBs for PSCCH allow different allocations of PSCCH in the time and frequency domain. For instance, for a large sub-channel size (e.g., $M_{\text{sub}} = 75$ PRBs), the PSCCH can occupy a large number of PRBs (i.e., $M_{\text{PSCCH}} = 25$ PRBs), and hence 2 symbols for the PSCCH may suffice. On the other hand, for a smaller sub-channel size (e.g., $M_{\text{sub}} = 15$ PRBs), the PSCCH may only be able to occupy $M_{\text{PSCCH}} = 10$ or 12 PRBs and thus, the PSCCH may need to use 3 symbols. Fig. 5 shows examples of a PSCCH/PSSCH slot with a PSCCH of 2 or 3 symbols.

A cyclic redundancy check (CRC) of 24 parity bits is appended to the 1st-stage SCI payload [53] to support error

⁷In fact, the sub-channel size M_{sub} should take into account the number M_{PSCCH} of PRBs that a PSCCH may need to occupy, e.g., depending on the

service [57]. For example, if 2 symbols are preferred for PSCCH, e.g., for faster PSCCH decoding, this may impose a minimum number of PRBs for PSCCH.

detection. The 1st-stage SCI payload with the appended CRC is encoded with the same Polar code as used for the physical downlink control channel (PDCCH) in Rel. 15 NR Uu [53]. The coded bits are then modulated with QPSK and mapped onto the resources for PSCCH [36].

As the number of symbols and number of PRBs for PSCCH are both (pre-)configured, the size of the 1st-stage SCI is fixed within a resource pool. Since the resource pool can be used for any transmission type [44], the payload size of the 1st-stage SCI is the same for unicast, groupcast or broadcast transmissions within a resource pool. With L sub-channels within a resource pool, there are L possible locations for a PSCCH in a slot, starting from the second SL symbol in a slot and from the lowest PRB in each sub-channel. To receive a 1st-stage SCI in NR V2X, a UE needs to check the L possible PSSCH locations at each slot within a resource pool. This is in contrast to LTE V2X where a UE needs to check different pairs of PRBs to determine if they carry a PSCCH [45].

For demodulating the PSCCH, DMRS is transmitted within the PSCCH. The PSCCH DMRS follows the design of the DMRS associated with PDCCH in Rel. 15 NR Uu [44]. The PSCCH DMRS reuses the same pseudo-random sequence used for the PDCCH DMRS in Rel. 15 NR Uu [36], with the PSCCH DMRS sequence initialization based on a (pre-)configured value per resource pool. In addition, every PSCCH symbol contains PSCCH DMRS. The pattern of PSCCH DMRS in the frequency domain also reuses the DMRS frequency pattern employed for the PDCCH in Rel. 15 NR Uu [44].

2) Physical Sidelink Shared Channel (PSSCH)

The PSSCH carries the 2nd-stage SCI and the data payload consisting of a TB [45]. The 2nd-stage SCI carries information used for decoding PSSCH and for supporting HARQ feedback and CSI reporting [53]. The 2nd-stage SCI indicates the Layer 1 source ID and destination ID of a transmission that represent identifiers (in the physical layer) of the TX UE and intended recipients (RX UEs) of the TB. The Layer 1 source ID allows an RX UE to know the identity of the TX UE. This is used for determining the PSFCH for HARQ feedback as explained in Section V.C.2). The 2nd-stage SCI also carries a one-bit new data indicator (NDI) that is used to specify whether the TB sent in the PSSCH corresponds to the transmission of new data or a retransmission. Furthermore, a HARQ process ID is also included in the 2nd-stage SCI in order to identify a TB. For instance, if the NDI notifies of a retransmission in a PSSCH, an RX UE can determine the TB for which the retransmission corresponds to using the HARQ process ID. The 2nd-stage SCI also informs about the redundancy version (RV) that depends on the index of a retransmission⁸. The 2nd-stage SCI also indicates whether HARQ feedback is enabled/disabled for the PSSCH [46]. In NR V2X, two formats are supported for the 2nd-stage SCI: SCI format 2-A and SCI format 2-B [53].

SCI format 2-A is used when there is no HARQ feedback or

for supporting unicast HARQ feedback or groupcast HARQ feedback (option 1 or option 2). Thus, the 2nd-stage SCI with SCI format 2-A indicates the cast type among broadcast, unicast, and groupcast with HARQ feedback option 1 or option 2 [46]. To request CSI feedback from an RX UE, a one bit CSI request is sent in the 2nd-stage SCI with SCI format 2-A.

On the other hand, SCI format 2-B is used when there is no HARQ feedback or for supporting groupcast HARQ feedback option 1. For this purpose, the 2nd-stage SCI with SCI format 2-B also includes the required communication range and the TX UE's zone ID [46]. The zone ID is indicated with 12 bits so a given area can be divided into 2^{12} squared regions of equal size⁹. The TX UE's zone ID indicates the zone in which the TX UE is located, i.e., it provides an indication of the TX UE's location [56]. The required communication range is represented with 4 bits using a set of 16 (pre-)configured values that can be selected out of a defined set of possible values¹⁰.

After decoding the 1st-stage SCI in PSCCH, an RX UE has the necessary information to decode the 2nd-stage SCI carried in PSSCH. Thus, no blind decoding of the 2nd-stage SCI is needed [44]. The 2nd-stage SCI is decoded using PSSCH DMRS. To support error detection, a 24 bits CRC is appended to the 2nd-stage SCI [46]. Like the 1st-stage SCI, the 2nd-stage SCI and the appended CRC is encoded with the same Polar coding as the NR PDCCH. The coded bits are modulated with QPSK [53].

The two-stage SCI in NR V2X reduces the complexity of the SCI decoding, not only for RX UEs, but also for sensing UEs which only need to decode the 1st-stage SCI to know which resources are reserved by a TX UE. This key advantage of the two-stage SCI stems from its design features [57]: (i) the 1st-stage SCI having a fixed size (independent of transmission type); (ii) the 1st-stage SCI being carried in PSCCH on a known possible location within a sub-channel; (iii) the 1st-stage SCI indicating the resources of the 2nd-stage SCI carried in PSSCH; and (iv) the 2nd-stage SCI having a varying payload size (depending on transmission type).

The PSSCH also carries the data payload, i.e., the TB. For error detection, A CRC of 24 parity bits is appended to the TB [53]. The TB is encoded using the Low-Density Parity-Check (LDPC) coding employed for PDSCH in Rel. 15 NR Uu [44]. If the TB size is larger than the maximum code block size that can be handled by the LDPC encoder, the TB (including the appended CRC) is divided into code blocks [53]. Each code block is encoded into codewords using the LDPC coding. The codewords of the TB are modulated using QPSK, 16-QAM, 64-QAM or 256-QAM [36]. The MCS is selected from the default MCS table or from the up to two further MCS tables which can be (pre-)configured per resource pool (See Section V.B.1)).

Before being mapped onto a PSSCH, the coded 2nd-stage SCI and coded TB are multiplexed according to the procedure described in Section 8.2.1 of [53]. Depending on the number of layers (i.e., number of data streams) supported in a PSSCH, the multiplexed 2nd-stage SCI and TB are mapped to one or two

⁸This can be used for soft combining in HARQ with incremental redundancy [58], where each retransmission provides a different RV of the data.

⁹The sides of the zones are configurable per required communication range and resource pool and can be equal to 5, 10, 20, 30, 40 or 50 m [49].

¹⁰The defined set of possible values include: 20, 50, 80,100, 120, 150, 180, 200, 220, 250, 270, 300, 320, 350, 370, 400, 420, 450, 480, 500, 550, 600, 700, 1000 m [49].

layers and precoded before being mapped to L_{PSSCH} sub-channels of the PSSCH [42]. With M_{sub} PRBs per sub-channel, the number of PRBs spanned by the PSSCH is $M_{\text{PSSCH}} = L_{\text{PSSCH}} \cdot M_{\text{sub}}$ PRBs, starting from the lowest PRB within the sub-channel carrying the corresponding PSCCH.

PSSCH can be transmitted from the second SL symbol up to the second to last SL symbol in a slot. As 7 to 14 SL symbols can be (pre-)configured in a slot, PSSCH can be sent in 5 to 12 consecutive SL symbols. The number of PSSCH symbols depends on the number of SL symbols in a slot and whether PSCCH is sent in the slot. In the 2 or 3 SL symbols which carry PSCCH, the PSSCH can be multiplexed in the frequency domain with PSCCH (if the PSCCH does not span the entire L_{PSSCH} sub-channels), resulting in 2 or 3 PSCCH/PSSCH symbols. In the SL symbols without PSCCH, the PSSCH spans all the L_{PSSCH} sub-channels as shown in Fig. 5. The second SL symbol (containing the first PSCCH or PSCCH/PSSCH symbol) is duplicated in the first SL symbol that is used for automatic gain control (AGC¹¹) purposes. In addition, the symbol after the last PSSCH symbol is used as a guard symbol. Any remaining SL symbols can be used for PSFCH or for a further guard symbol as discussed in Section V.B.4).

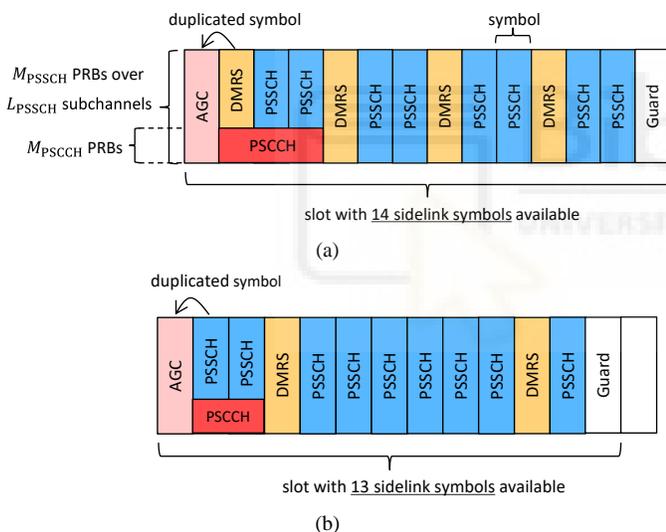


Fig. 5. Example of 2 configurations for a PSCCH/PSSCH slot (for normal CP). (a) Slot with 3 PSCCH, 12 PSSCH including 4 PSSCH DMRS symbols. (b) Slot with 2 PSCCH, 11 PSSCH including 2 PSSCH DMRS symbols.

For demodulating the PSSCH, DMRS is sent within PSSCH. The design of the PSSCH DMRS follows some aspects of the DMRS associated with the physical uplink shared channel (PUSCH) in Rel. 15 NR Uu and others from PDCCH DMRS [44]. The PSSCH DMRS reuses the pseudo-random sequence used for the Rel. 15 PUSCH DMRS [36], with the sequence initialization based on a (pre-)configured value per resource pool. In the frequency domain, the type 1 pattern configuration of PDSCH DMRS is employed for PSSCH DMRS [36].

The PSSCH DMRS can be transmitted in 2, 3, or 4 sidelink

symbols at different locations within a slot, i.e., with different time patterns. The different time patterns for the PSSCH DMRS depend on the number of symbols for PSCCH, the number of symbols with PSSCH DMRS and the number of symbols for PSSCH within a slot. The time patterns supported for PSSCH DMRS in NR V2X are listed in Table 8.4.1.1.2-1 of [36]. For a resource pool, one or more time patterns for PSSCH DMRS can be (pre-)configured. In case multiple patterns are (pre-)configured, the DMRS time pattern used in a PSSCH is indicated in the associated 1st-stage SCI.

Fig. 5 depicts two examples of PSCCH/PSSCH slots for different number of PSCCH symbols, number of DMRS and number of PSSCH symbols, which result in different DMRS time patterns¹² based on Table 8.4.1.2.2-1 of [36]. In the figure, a PSSCH symbol with DMRS is shown just as DMRS. For both examples, the first and last sidelink symbol consist of an AGC symbol and a guard symbol, respectively.

3) SL CSI-RS and SL PT-RS

In NR V2X, the transmission of SL CSI-RS is supported for unicast transmissions only. The SL CSI-RS is sent in the PSSCH region of the slot. The design of the SL CSI-RS is based on the CSI-RS design of Rel. 15 NR Uu [36]. In addition, the resource mapping of SL CSI-RS in a PRB is based on the CSI-RS resource mapping patterns in NR Uu, which support up to two antenna ports (as in NR V2X SL up to two streams can be supported in a PSSCH). Each PRB within PSSCH uses the same pattern for the SL CSI-RS. SL CSI-RS is not transmitted on symbols containing PSCCH, the 2nd-stage SCI or PSSCH DMRS [46]. The SL CSI-RS configuration includes the resource mapping pattern and the number of antenna ports for SL CSI-RS. The SL CSI-RS configuration is selected by the TX UE and provided to the RX UE via PC5-RRC configuration (see Section VII.C).

The transmission of SL CSI-RS by a TX UE along with a CSI request sent in the 1st-stage SCI triggers the RX UE of a unicast link to feed back a CSI report. The TX UE can configure aperiodic CSI reporting from the RX UE [36]. The RX UE can measure the CSI based on the SL CSI-RS sent by the TX UE as it will be discussed in Section V.C.4). The CSI is fed back over a PSSCH sent from the RX UE to the TX UE. To avoid outdated CSI, the RX UE is expected to feed back the CSI report within a maximum amount of time. This maximum amount of time is referred as the latency bound. The latency bound is determined by the TX UE and signalled to the RX UE via PC5-RRC [54].

The transmission of SL PT-RS is supported for FR2 only. SL PT-RS is sent in the PSSCH region of the slot. The reception of SL PT-RS, allows an RX UE to track phase offsets due to Doppler effects and hardware impairments [59]. The design of SL PT-RS is based on the PT-RS design of Rel. 15 NR Uu [36]. The resource mapping of SL PT-RS is based on the PT-RS resource mapping in NR Uu.

¹¹The received signal can vary over a wide dynamic range depending on the channel attenuation and interference. AGC is used to adjust the strength of the received signal in order to reduce the quantization error or the clipping of the signal at the analog to digital converter (ADC) [60].

¹²In Table 8.4.1.1.2-1 in [36], the number of PSSCH symbols listed in the first column includes the first sidelink symbol used for AGC as well as the PSSCH DMRS symbols.

4) Physical Sidelink Feedback Channel (PSFCH)

In NR V2X, the sole purpose of PSFCH is to carry the HARQ feedback from RX UE(s) to a TX UE. Within a resource pool, resources for PSFCH can be (pre-)configured periodically with a period of 1, 2 or 4 slot(s), i.e., there is a slot with PSFCH every 1, 2 or 4 slot(s) within a resource pool. PSFCH is sent in one symbol among the last SL symbols in a PSCCH/PSSCH slot as shown in the example in Fig. 6. Prior to the PSFCH symbol, one AGC symbol is used consisting of a copy of the PSFCH symbol. The symbol after the PSFCH symbol is used as a guard symbol. The three SL symbols associated with a PSFCH come after the PSSCH symbols as illustrated in Fig. 6. As a result, the number of PSSCH symbols (without the AGC and guard symbol) can be at most 9 symbols when a slot carries PSFCH. The DMRS time pattern of the slot in Fig. 6 results from Table 8.4.1.1.2-1 of [36], based on the given example with 2 PSCCH symbols, 7 PSSCH symbols and 2 PSSCH DMRS symbols.

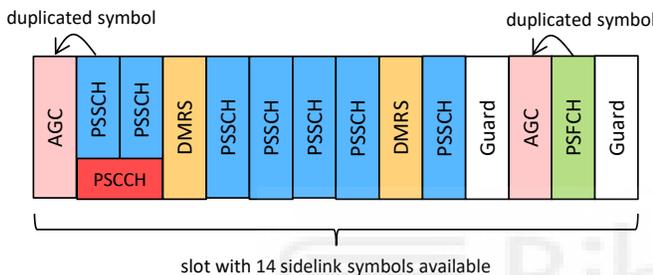


Fig. 6. Example of a PSCCH/PSSCH slot with a PSFCH, 2 PSCCH, 9 PSSCH including 2 PSSCH DMRS symbols (for normal CP).

In a resource pool, the set of PRBs in a PSFCH symbol that are available for PSFCH is indicated using a bitmap. NR V2X supports that all or only a subset¹³ of the PRBs in a PSFCH symbol can be used for PSFCH. For a transmission of PSFCH, one PRB in a PSFCH symbol is used that carries a Zadoff-Chu sequence [61] based on the sequences used for the physical uplink control channel (PUCCH) in Rel. 15 NR Uu (PUCCH format 0) [36]. The procedure for selecting a PRB and the sequence in order to send HARQ feedback in a PSFCH is discussed in Section V.C.2).

5) Sidelink Synchronization Signal Block (S-SSB): PSBCH, S-PSS and S-SSS

Synchronization information can be transmitted by a SyncRef UE in the sidelink [52], e.g., to expand the synchronization coverage of a synchronization source and to enable that nearby UEs have the same sidelink timing reference. This allows SL communication to/from the SyncRef UE as well as SL communication between nearby UEs. The sidelink synchronization information is carried in an S-SSB that consists of PSBCH, S-PSS and S-SSS. The S-SSB occupies one slot and uses the same numerology as the one (pre-)configured in the SL BWP, i.e., the same numerology as for PSCCH/PSSCH. For a normal CP or extended CP, the PSBCH, S-PSS and S-SSS are

carried in the first 13 or 11 symbols of an S-SSB slot, respectively. An S-SSB slot for a normal CP is depicted in Fig. 7. The last symbol in an S-SSB slot is used as a guard symbol. The S-SSB is not frequency multiplexed with any other sidelink physical channel within the SL BWP, i.e., S-SSBs are not transmitted in the slots of a resource pool. In the frequency domain, the S-SSB spans $M_{S-SSB} = 11$ common RBs within the SL BWP [36]. Since an RB consists of 12 subcarriers, the S-SSB bandwidth spans $11 \times 12 = 132$ subcarriers. The frequency location of an S-SSB is (pre-)configured within a SL BWP [62]. As a result, a UE does not need to perform blind detection in the frequency domain to find an S-SSB [62].



Fig. 7. S-SSB including PSBCH, S-PSS and S-SSS (for normal CP).

The S-PSS and S-SSS are jointly referred to as the sidelink synchronization signal (SLSS). The SLSS is used for time and frequency synchronization. By detecting the SLSS sent by a SyncRef UE, a UE is able to synchronize to the SyncRef UE and estimate the beginning of the frame and carrier frequency offsets. The UE can use the SL timing reference provided by the SyncRef UE for SL transmissions with nearby UEs that are using the same timing reference. Thus, a TX UE does not need to transmit also S-SSBs, i.e., not every TX UE needs to be a SyncRef UE. The S-PSS and the S-SSS consist each of sequences of 127 bits. The S-PSS is generated from maximum length sequences (m -sequences) using the same polynomial and initial values as the ones used for generating the m -sequences in the primary synchronization signal (PSS) in Rel. 15 NR Uu [36]. In NR Uu, there are three candidate sequences for PSS. However, only two candidate sequences are used for S-PSS¹⁴. The sidelink-SSS (S-SSS) is generated from Gold sequences that use the same design (i.e., generator polynomials, initial values and cyclic shifts) which is utilized for generating the Gold sequences for the secondary synchronization signal (SSS) in Rel. 15 NR Uu [36]. This results in 336 candidate sequences for S-SSS like for the SSS in NR Uu.

For the transmission of SLSS within an S-SSB, a SyncRef UE selects an S-PSS and S-SSS out of the candidate sequences based on an SLSS ID. The SLSS ID represents an identifier of the SyncRef UE and conveys a priority of the SyncRef UE as in LTE V2X. The procedure for selecting an SLSS ID is discussed in Section V.C.6). Each SLSS ID corresponds to a unique combination of an S-PSS and S-SSS out of the 2 S-PSS and 336 S-SSS candidate sequences, with the following relation $SLSS\ ID = 336 \times N_{S-PSS} + N_{S-SSS}$, where $N_{S-PSS} = \{0,1\}$ and $N_{S-SSS} = \{0,1, \dots, 335\}$ represent an identifier of the S-PSS and S-SSS among the S-PSS and S-SSS candidate sequences,

¹³As discussed in Section V.C.2), PRBs in a PSFCH symbol are distributed equally among several sub-channels for the HARQ feedback. This may result in leaving some PRBs unused, depending on the number of PRBs in a symbol.

¹⁴The two candidate sequences are generated from the same base sequence with two different cyclic shifts, namely, 25 and 65. The concept of cyclic shifts is discussed in Section V.C.2).

respectively [36]. Thus, there are $2 \times 336 = 672$ unique SLSS IDs. Once a UE has detected the SLSS and identified the S-PSS and S-SSS transmitted by a SyncRef UE, it can determine the SLSS ID of the SyncRef UE. Thus, the transmission of SLSS within an S-SSB also conveys the SLSS ID of a SyncRef UE.

The S-PSS and S-SSS are modulated with BPSK such that each sequence occupies 127 subcarriers in a symbol within the S-SSB bandwidth as shown in Fig. 7. The S-PSS and S-SSS do not occupy the entire S-SSB bandwidth of 132 subcarriers, and are allocated from the third subcarrier relative to the start of the S-SSB bandwidth up to the 129th subcarrier. The S-PSS and S-SSS are both repeated in two consecutive symbols in an S-SSB. This provides higher coverage and allows a UE receiving an S-SSB to perform phase tracking between two consecutive symbols carrying the same sequence [45]. Independently of the CP, the S-PSS is sent in the second and third symbol of the S-SSB, while the S-SSS is sent in the fourth and fifth symbol of the S-SSB as shown in Fig. 7. PBSCCH is sent in the first symbol and on the eight or six symbols after the S-SSS for a normal or extended CP, respectively. The first symbol of the S-SSB (first PSBCH symbol) is used for AGC. For a given CP, the structure of an S-SSB slot is fixed, in contrast to a PSCCH/PSSCH slot.

The main purpose of the PSBCH is to provide system-wide information and synchronization information that is required by a UE for establishing a sidelink connection. The information carried by the PSBCH includes a one bit indicator whether the SyncRef UE is in coverage of a network or of GNSS [44]. A UE is in coverage of GNSS when GNSS is reliable at the UE. GNSS is reliable if the UE can meet the accuracy requirements specified in [52]. SL transmissions are organized in frames identified by the direct frame number (DFN). Consequently, the PSBCH also indicates the DFN and the slot index as timing information [44]. The DFN enables a UE to synchronize its radio frame transmissions according to the SL timing reference. For a SyncRef UE in network coverage, the DFN can be derived based on the system frame number (SFN), where the SFN provides an indexing of the frames based on the cell timing reference. When a UE is out of network coverage, the DFN can be derived based on the coordinated universal time (UTC) provided by GNSS (i.e., GNSS timing) [63]. For a UE in network coverage but using GNSS as synchronization source for SL communication, the DFN can be derived from the GNSS timing and a timing offset with respect to the cell timing reference [63]. This timing offset enables to align the DFN and SFN timing to achieve a unified SL timing [64]. The time division duplex (TDD) configuration is also carried in the PSBCH [44]. The TDD configuration provides the TDD UL/DL configuration and/or the indication of the slots that can be contained in a resource pool. A UE that is out of coverage of a gNB can receive an S-SSB sent by a SyncRef UE that is in coverage of the gNB. From the PSBCH within the S-SSB, the out of network coverage UE also receives the TDD UL/DL configuration used by the gNB within its cell. This enables the out of network coverage UE to be aware of the slots that the gNB uses for the downlink within the cell. This UE can then avoid causing interference to UEs that are receiving a downlink transmission in the cell. To do so, it should avoid using the

downlink slots of the cell for its SL transmissions [65]. The payload size of the PSBCH consists of 56 bits, including the 1 bit in coverage indicator, 12 bits for indicating the TDD configuration, 10 bits for the DFN and 7 bits for the slot index. The PSBCH payload also includes 2 reserve bits for future purposes as well 24 bits for a CRC. The PSBCH payload is encoded using the same Polar coding employed for the physical broadcast channel (PBCH) in Rel. 15 NR Uu [53]. The PSBCH occupies 11 RBs of the S-SSB bandwidth as depicted in Fig. 7.

After successfully detecting the SLSS sent by a SyncRef UE, a UE proceeds with decoding the PSBCH. For this purpose, DMRS are transmitted in every PSBCH symbol and on every fourth subcarrier (starting from the first subcarrier), following a similar structure as the DMRS in the PBCH in Rel. 15 NR Uu [36]. The PSBCH DMRS sequence generation is also based on the design of PBCH DMRS for Rel. 15, which results in 31-length Gold sequences using QPSK modulation [36]. The SLSS ID is used for the PSBCH DMRS sequence initialization [49].

S-SSBs are sent with a fixed periodicity of 160 ms (i.e., 16 frames). This is in contrast to the SSB structure in Rel. 15 NR Uu, which supports multiple periodicities. Similar to the SSB structure in NR Uu, multiple S-SSBs can be sent within one S-SSB period. Within one S-SSB period, the number of S-SSBs is (pre-)configurable depending on the SCS and the frequency range as given in Table V [44]. For larger SCSs, a larger number of S-SSBs within a period is supported such that sufficient SL coverage can be provided also for larger SCSs. S-SSBs are distributed within an S-SSB period according to two (pre-)configured parameters, namely the offset from the start of the S-SSB period to the first S-SSB and the interval between consecutive S-SSBs (see Fig. 8) [44]. The beginning of an S-SSB period corresponds to the first slot in a frame with the SFN satisfying $\text{SFN mod } 16 = 0$, where mod is the modulo operator. Each frame within a period of 1024 frames can be identified by the SFN. The SFN allows defining transmission periods which are longer than one frame [41], e.g., as the S-SSB transmissions. When GNSS is used as the synchronization reference, the DFN is used instead of SFN [63].

TABLE V
NUMBER OF S-SSB WITHIN AN S-SSB PERIOD (160 MS)

Frequency Range	Subcarrier Spacing (SCS)	Number of S-SSBs per period
FR1	15 kHz	1
	30 kHz	1, 2
	60 kHz	1, 2, 4
FR2	60 kHz	1, 2, 4, 8, 16, 32
	120 kHz	1, 2, 4, 8, 16, 32, 64

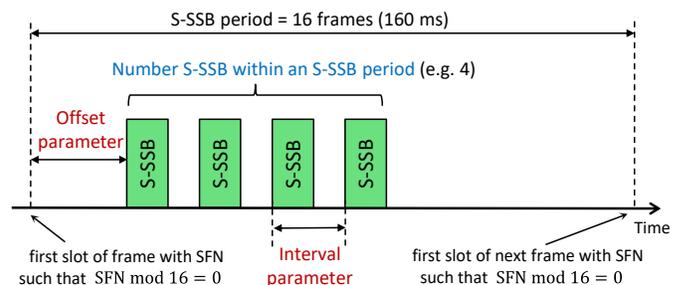


Fig. 8. Multiple S-SSB transmissions within an S-SSB period.

TABLE VI
OVERVIEW OF NR V2X PHYSICAL CHANNELS AND SIGNALS

Physical Channels	Purpose	Associated Physical Signals	Additional Aspects
PSCCH (sent within first SL symbols of first sub-channel occupied by associated PSSCH)	PSCCH carries the <u>1st-stage SCI</u> which indicates [53]: <ul style="list-style-type: none"> Resource allocation, modulation and coding scheme, and priority of associated PSSCH Resource reservation period Time pattern and number of ports for PSSCH DMRS Size and format of 2nd-stage SCI 	<ul style="list-style-type: none"> PSCCH DMRS [36]: based on Rel. 15 NR PDCCH DMRS 	<ul style="list-style-type: none"> Coding [53]: Polar coding used for Rel. 15 NR PDCCH Modulation: QPSK
PSSCH (sent in one or multiple sub-channels in a slot within a resource pool)	PSSCH carries a <u>TB</u> and the <u>2nd-stage SCI</u> which indicates [53]: <ul style="list-style-type: none"> HARQ process ID, new data indicator and redundancy version Source ID and Destination ID HARQ enabled/disabled indicator Transmission type and CSI request (SCI format 2-A) TX UE's zone ID and required communication range (SCI format 2-B) 	<ul style="list-style-type: none"> PSSCH DMRS [36]: based on Rel. 15 NR PUSCH DMRS and PDSCH DMRS SL CSI-RS [36]: based on Rel. 15 NR CSI-RS SL PT-RS [36]: based on Rel. 15 NR PT-RS 	<u>2nd-stage SCI</u> : <ul style="list-style-type: none"> Coding [53]: Polar coding used for Rel. 15 NR PDCCH Modulation: QPSK <u>Transport block</u> : <ul style="list-style-type: none"> Coding [53]: LDPC coding used for Rel. 15 NR PDSCH Supported modulation: QPSK, 16-QAM, 64-QAM, 256-QAM
PSBCH (sent periodically with the S-PSS and S-SSS in an S-SSB, not on slots of a resource pool)	PSBCH carries [44]: <ul style="list-style-type: none"> TDD configuration, Direct frame number Slot index in network/GNSS coverage indicator 	<ul style="list-style-type: none"> PSBCH DMRS [36]: based on Rel. 15 NR PBCH DMRS S-PSS [36]: <i>m</i>-sequence (based on Rel. 15 NR PSS) S-SSS [36]: Gold sequence (as Rel. 15 NR SSS) 	<ul style="list-style-type: none"> Coding for PSBCH [53]: LDPC coding used for Rel. 15 NR PBCH Modulation for PSBCH: QPSK Modulation for S-PSS and S-SSS: BPSK
PSFCH (sent periodically, in a symbol near end of a PSCCH/PSSCH slot)	As response to a PSSCH reception [66], PSFCH carries HARQ feedback: <ul style="list-style-type: none"> For unicast: ACK/NACK feedback For groupcast: NACK-only feedback (option 1) or ACK/NACK feedback (option 2) 		For a transmission of PSFCH, one PRB in a PSFCH symbol used carrying a <ul style="list-style-type: none"> Zadoff-Chu sequence [36]: based on Rel. 15 NR PUCCH format 0

As the SSB structure in NR Uu, the S-SSB structure supports the transmission of each S-SSB with a different beam, e.g. with beam sweeping [64]. Although the S-SSB structure is similar to the SSB structure, the SSBs transmission by a gNB are also part of the initial access procedure in NR Uu [66]. However, there is no such procedure in NR V2X sidelink. In NR Uu, after a UE has identified the best beamformed SSB (e.g., best beam from the gNB to reach the UE), the UE indicates this to the gNB via the physical random access channel (PRACH) in the uplink. This procedure is referred as random access and enables the gNB to be aware of the UE and of the best beam to reach the UE. As there is no similar procedure after the S-SSB transmissions in NR V2X SL, a SyncRef UE is not aware of the UEs that received its S-SSBs transmissions.

The main characteristics of the NR V2X SL physical channels and the S-SSB as well the associated signals and the two stage SCI design are summarized in Table VI.

C. Physical Layer Procedures for NR V2X sidelink

Rel. 16 NR V2X introduces unicast and groupcast communications for the sidelink. To support them, NR V2X introduces mechanisms for sidelink HARQ feedback and for determining the PSFCH to send the HARQ feedback. In addition, sidelink power control has also been introduced for NR V2X SL, as well as CSI acquisition for unicast transmissions. Similar to LTE V2X, NR V2X also supports

procedures for SL transmission of synchronization information.

1) Sidelink HARQ feedback

HARQ increases the reliability of TB transmissions by employing forward error correction (FEC) and error detection codes in combination with a retransmission strategy [69]. With HARQ, a receiver requests a HARQ retransmission if it is not able to correct all transmission errors via the FEC code and it detects errors in the transmission via the error detection code. The HARQ retransmission can include data and parity bits to aid the successful reception of a TB. After a transmission, the transmitter waits for HARQ feedback from the receiver. The reply can consist of: (i) an acknowledgement (ACK) if the reception was successful; (ii) a negative acknowledgement (NACK) if the reception was unsuccessful; or (iii) no response if the control information associated with the transmission was not received successfully within a predefined time. In case of NACK or no response, the transmitter can perform a HARQ retransmission, with multiple retransmissions being possible.

Retransmissions in the SL can aid in meeting the reliability requirements for NR V2X use cases. For this purpose, blind retransmissions can also be considered [44], i.e., a certain number of retransmissions can be performed without any HARQ feedback. However, sidelink HARQ feedback can prevent unnecessary blind retransmissions that would otherwise waste SL resources and increase the channel load.

NR V2X supports ACK/NACK feedback for the SL HARQ feedback in unicast: an RX UE sends ACK if it has successfully decoded the TB carried in a PSSCH or it sends NACK if it has not decoded the TB after decoding the 1st-stage SCI. For groupcast, two options (option 1 and option 2) are supported for the SL HARQ feedback in NR V2X. For option 1, an RX UE transmits NACK if it has not successfully decoded the TB (after decoding the 1st-stage SCI) and if its relative distance to the TX UE (referred as Tx-Rx distance) is less than or equal to the required communication range (indicated in the 2nd-stage SCI). Otherwise, the RX UE does not transmit any HARQ feedback. As the HARQ feedback for this option would only consist of NACK, option 1 is referred to as NACK-only feedback. Option 1 is illustrated in Fig. 9, which depicts scenarios where RX UEs do not transmit HARQ feedback because they have successfully decoded a TB or because they are outside the minimum required communication range (even if they did not successfully decode the TB). Option 2 for groupcast supports the ACK/NACK feedback from all RX UEs: an RX UE sends ACK if it has successfully decoded the TB or it sends NACK if it has not decoded the TB, after decoding the 1st-stage SCI. For unicast or any of the groupcast options, an RX UE does not send a reply if the RX UE does not decode the 1st-stage SCI.

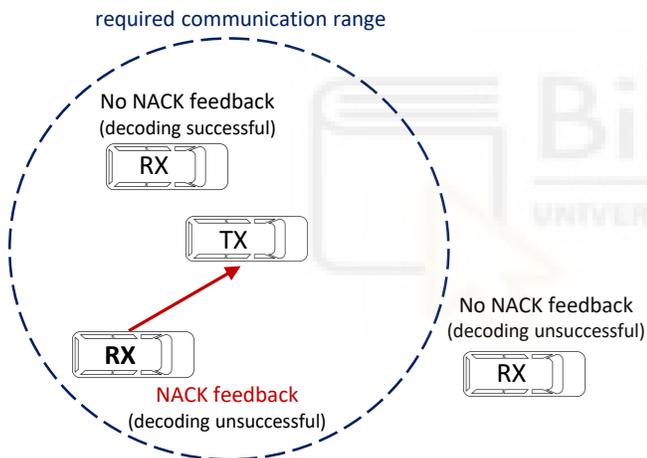


Fig. 9. NACK-only feedback for groupcast NR V2X sidelink (option 1).

With option 1, the RX UEs of a transmission share a resource for sending their NACK-only feedback, while with option 2 each RX UE sends its ACK/NACK feedback on a separate resource. Thus, with option 1 a TX UE cannot identify which RX UEs sent NACK. If a TX UE receives at least one NACK with option 1, it is aware that at least one RX UE within the required communication range did not correctly decode the transmission. In addition, if a TX UE receives no reply with option 1, it cannot distinguish whether the RX UEs within the required communication range have successfully received the transmission or if some have not successfully decoded the corresponding 1st-stage SCI. On the other hand, a TX UE can distinguish the HARQ feedback of the RX UEs with option 2. This enables the TX UE to perform a retransmission tailored to

specific RX UE(s). With option 2, if a TX UE receives no reply on the feedback resource corresponding to a given RX UE, it is aware that the RX UE has not successfully decoded the corresponding 1st-stage SCI. Consequently, option 2 enables higher reliability for the transmissions. The above advantages of option 2 over option 1, however, come at the expense of more resources required for the groupcast HARQ feedback. In NR V2X, the use of HARQ feedback and the choice whether to use option 1 or option 2 for groupcast HARQ feedback is up to UE implementation.

For groupcast communications, the TX UE indicates in the 2nd-stage SCI whether NACK-only feedback (option 1) or ACK/NACK feedback (option 2) should be used. NACK-only feedback may be considered for groupcast services where the provided information may not be so relevant for RX UEs outside the communication range, e.g., for the extended sensors use case (see Section III.A) [45]. For groupcast option 1, the Tx-Rx distance is obtained at the RX UE based on the location of the TX UE¹⁵. This location is indicated via the TX UE's zone ID carried in the 2nd-stage SCI. The Tx-Rx distance is derived at the RX UE based on its own location and the center of the nearest zone¹⁶ with the indicated zone ID.

HARQ feedback for a TB sent on a PSSCH in a resource pool is carried on a PSFCH within the same resource pool. To disable the HARQ feedback for all SL transmissions in the resource pool, no resources for PSFCH are configured within the resource pool. Resources for PSFCH can be (pre-)configured periodically with a period of $N = 1, 2$ or 4 slot(s), i.e., there is a PSCCH/PSSCH slot with a PSFCH symbol every N slots within the resource pool. Even if resources for PSFCH are (pre-)configured, the 2nd-stage SCI indicates whether HARQ feedback is enabled or not for a TB sent in a given PSSCH [56].

As discussed in Section V.B.4), one PRB is used for a PSFCH transmission. In one PRB, NR V2X supports code division multiplexing (CDM) between PSFCH transmissions of multiple RX UEs. With CDM, multiple transmissions share the same frequency resource simultaneously. To this aim, the transmissions should employ signals with low cross correlation (ideally orthogonal) [61]. This property allows a receiver to distinguish uncorrelated transmission signals that have been sent on the same time and frequency resources. Signals with low cross correlation can be generated from a periodic root or base sequence that has very good autocorrelation properties [70]. A set of uncorrelated signals for CDM can be generated from cyclically shifted versions of the base sequence [71]. The base sequence used for PSFCH is (pre-)configured per resource pool in NR V2X and it corresponds to a Zadoff-Chu sequence based on a format used for PUCCH in Rel. 15 NR Uu [36].

In a PRB used for PSFCH, the transmissions that can be multiplexed with CDM correspond to the HARQ feedback from several RX UEs. CDM is also used to distinguish the ACK or NACK feedback from an RX UE in the same PRB. To this aim, a pair of cyclic shifts (within a PRB) are used to distinguish the ACK or NACK from an RX UE [66]. The cyclic shift

¹⁵ NACK-only feedback is used only if the TX UE knows its position.

¹⁶As there is a limited number of zone IDs, the set of available zones repeat over a given area and hence there can be an ambiguity when indicating the TX

UE's position with a zone ID. Out of all possible zones with the indicated zone ID, an RX UE assumes that the TX UE is in the zone closest to the RX UE.

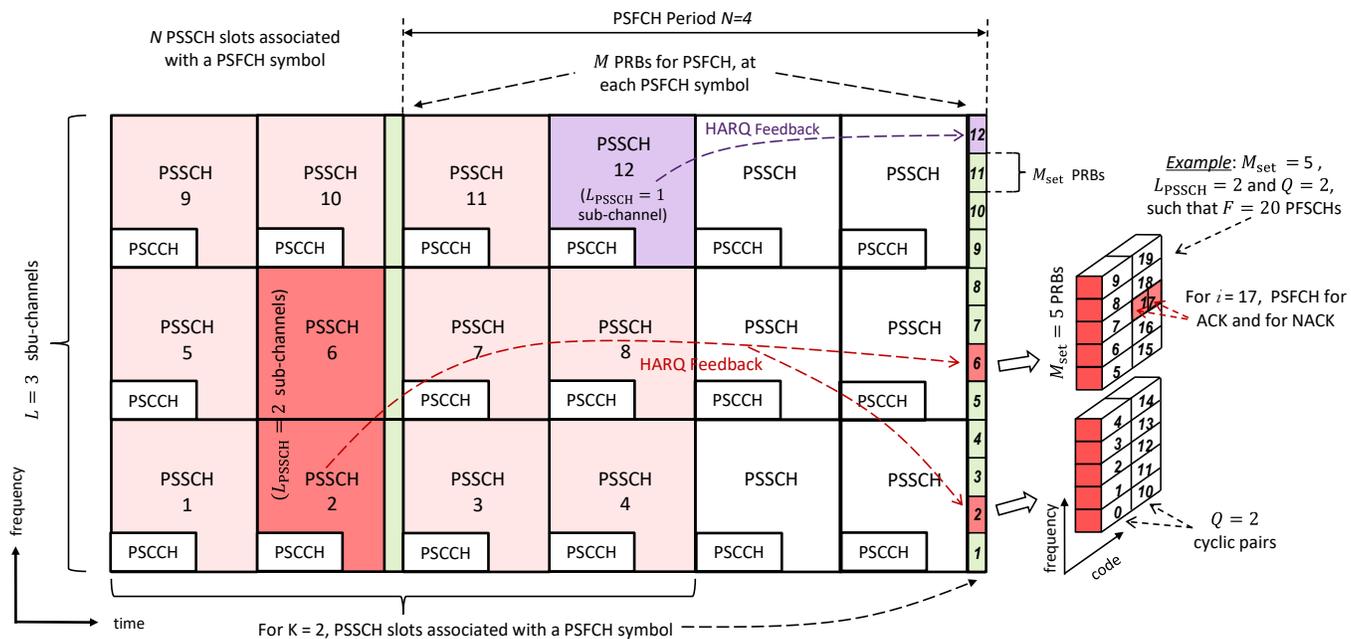


Fig. 10. PSFCHs for HARQ feedback associated with different transmissions.

corresponding to ACK is not defined in the case of NACK-only feedback for groupcast (option 1). Thus, each PSFCH in NR V2X is mapped to a time resource (PSFCH symbol), a frequency resource (one PRB), and a code resource (one cyclic shift among a cyclic shift pair). Next, we discuss how the PSFCH symbol, the PRB and cyclic shift are determined for the HARQ feedback for a given PSSCH transmission.

2) Resource selection for PSFCH

The PSFCH symbol that can be used for the HARQ feedback for a given PSSCH transmission corresponds to the PSFCH symbol in the first slot with PSFCH after a (pre-)configured number of K slots after the PSSCH. K represents the minimum number of slots within the resource pool between a slot with a PSSCH transmission and the slot containing PSFCH for the HARQ feedback of this transmission (see Fig. 11 for $K=3$). Let us consider that the last symbol of a PSSCH transmission is on slot n . The HARQ feedback for this transmission is expected in slot $n+a$, where a is the smallest integer equal or higher than K such that slot $n+a$ contains PSFCH. In the example shown in Fig. 11, the earliest possible slot for the HARQ feedback (slot $n+3$) does not contain PSFCH, so the HARQ feedback is sent at the next slot with PSFCH (slot $n+6$). The time gap of at least K slots allows considering the RX UE's processing delay in decoding the PSCCH and generating the HARQ feedback. K can be equal to 2 or 3, and a single value of K can be (pre-)configured per resource pool¹⁷. This allows several RX UEs using the same resource pool to utilize the same mapping of PSFCH resource(s) for the HARQ feedback. With the parameter K , the N PSSCH slots associated with a slot with PSFCH can be determined. In the example with $K=3$ shown in Fig. 11, the $N=4$ PSSCH slots associated with the PSFCHs at

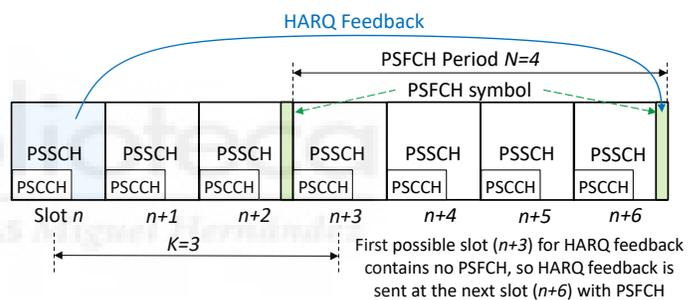


Fig. 11. PSSCH-to-HARQ feedback timing based on at least $K=3$ slots. For simplicity, the figure depicts only one sub-channel within the resource pool. We also omit the detailed structure of a PSCCH/PSSCH slot with or without PSFCH including PSSCH DMRS, AGC symbols and guard symbols.

slot $n+6$ correspond to PSSCH slots $n, n+1, n+2$, and $n+3$.

With L sub-channels in a resource pool and N PSSCH slots associated with a slot containing PSFCH, there are then $N \cdot L$ sub-channels associated with a PSFCH symbol. With M PRBs available for PSFCH in a PSFCH symbol, there are M PRBs available for the HARQ feedback of transmissions over $N \cdot L$ sub-channels. With M configured to be a multiple of $N \cdot L$ [49], then a distinct set of $M_{\text{set}} = M/(N \cdot L)$ PRBs can be associated with the HARQ feedback for each sub-channel within a PSFCH period. The first set of M_{set} PRBs among the M PRBs available for PSFCH are associated with the HARQ feedback of a transmission in the first sub-channel in the first slot. The second set of M_{set} PRBs are associated with the HARQ feedback of a transmission in the first sub-channel in the second slot and so on. This is illustrated in Fig. 10 with $N = 4, L = 3$ and with all PRBs in a PSFCH symbol available for PSFCH. In this example, the HARQ feedback for a transmission at PSSCH x is sent on the set x of M_{set} PRBs in the corresponding PSFCH

¹⁷In NR Uu, the PDSCH-to-HARQ timing (similar to the K slots for the SL HARQ feedback) is signaled in the DCI for each DL transmission [53]. In NR

V2X, the PSSCH-to-HARQ feedback timing is not indicated in the SL control information, as it is (pre-)configured per resource pool with the parameter K .

$n+2$, it would then expect to receive the PSFCH from UE D at the same slot (slot $n+6$) where UE A is expected to send PSFCH to UE B and/or UE C. However, the half-duplex constraint prevents a UE from transmitting a PSFCH and receiving a PSFCH at the same time. In case a UE is expected to transmit and receive PSFCHs in the same PSFCH symbol, a UE selects to transmit or receive PSFCH(s) also based on the priority of the associated PSSCH(s) [66].

4) Sidelink CSI Reporting

NR V2X can use link adaptation to increase the efficiency and reliability of unicast transmissions. Link adaptation adjusts the MCS for a transmission according to the quality of the radio link²⁰ [73]. In NR V2X unicast²¹, link adaptation can be based on the feedback of a channel quality indicator (CQI) from an RX UE to the TX UE. To determine the CQI, an RX UE measures the sidelink channel from the TX UE using the SL CSI-RSs transmitted by the TX UE within the PSSCH region of the slot. The CQI is selected from a CQI table that is derived based on the configured MCS table for a PSSCH [46]. The CQI provides an indication of the highest modulation and coding scheme that can be supported by the SL channel measured at the RX UE.

To improve the spectral efficiency, rank adaptation of unicast transmissions in NR V2X can also be applied by adjusting the number of streams that can be sent in a PSSCH with a multi-antenna transmission. Based on channel measurements of SL CSI-RSs sent from up to two antenna ports of the TX UE, the RX UE can determine the rank indicator (RI) that corresponds to the rank of the measured SL channel. The rank of the channel determines the number of streams that can be supported by the channel²². As up to two streams can be supported for a PSSCH transmission in NR V2X SL, the RI can be equal to 1 or 2. The RI is determined jointly with the CQI. The combination of CQI and RI represents the CSI that can be fed back from an RX UE to a TX UE for link and rank adaptation of a unicast PSSCH transmission. This is in contrast to the CSI that can be fed back by a UE to the gNB for adapting a downlink transmission in Rel. 15 NR Uu, where a UE can also feed back a precoding matrix indicator (PMI) besides the CQI and RI [41]. The PMI represents the index of the most suitable precoding matrix (out of a pre-defined codebook) that can be used to enable closed-loop multi-antenna transmission to the UE. As no PMI feedback is supported in NR V2X sidelink, open-loop multi-antenna SL transmissions can be performed by a TX UE based on the CQI and RI fed back by an RX UE.

An RX UE feeds back to the TX UE the CSI (CQI and RI) via CSI reporting. The CSI report is carried in the MAC control element (CE) [74] over a PSSCH sent from the RX UE to the TX UE. The latency bound of the CSI report is signaled by the TX UE to the RX UE via PC5-RRC [54]. It is expected that the RX UE feeds back the CSI report within the latency bound [54].

²⁰For a poor quality link, a lower modulation scheme (e.g., QPSK) and lower code rate could be used for higher robustness. Higher modulation orders and code rates can be used when the SNR is high to increase the spectral efficiency.

²¹Adapting a groupcast transmission to RX UEs with different channel qualities requires CSI feedback from multiple RX UEs and could lead to adapting the transmission for the worst CSI, which may not be efficient. The

The transmission of SL CSI-RS by a TX UE and the indication of CSI request in the 1st-stage SCI trigger the CSI reporting from an RX UE [42].

5) Selection of the Synchronization Reference

For NR V2X sidelink, there are two main sources for synchronization, namely GNSS and a gNB or eNB (referred to as gNB/eNB). In addition, a UE can use a SyncRef UE or its own internal clock as its synchronization reference [52]. Thus, there are four types of synchronization references for NR V2X sidelink (Fig. 13): gNB/eNB, GNSS, SyncRef UE, or a UE's own internal clock. In NR V2X, a UE selects its synchronization reference following the same procedure as in LTE V2X [55], i.e., based on the different priorities of gNB/eNB, GNSS and SyncRef UEs.

The priority of a SyncRef UE is given by its SLSS ID and the in coverage indicator that is carried on the PSBCH within the S-SSB transmissions of a SyncRef UE. We refer to the in coverage indicator as I_{IC} . $I_{IC} = 0$ indicates that a SyncRef UE is not in coverage of GNSS or a gNB/eNB and $I_{IC} = 1$ indicates that it is in coverage of GNSS or a gNB/eNB. The 672 SLSS IDs available in NR V2X are divided into two sets [44]:

- 1) In coverage SLSS IDs = $\{0, 1, \dots, 335\}$,
- 2) Out of coverage SLSS IDs = $\{336, 337, \dots, 671\}$.

SLSS IDs within the set of in coverage SLSS IDs are used by SyncRef UEs that are synchronized to a gNB/eNB, GNSS, or to another SyncRef UE which in turn is synchronized to a gNB/eNB or GNSS. As in LTE V2X, the SLSS ID = 0 is used in NR V2X by SyncRef UEs which are synchronized to GNSS or to another SyncRef UE that is synchronized to GNSS ([44]). SLSS IDs within the set of out of coverage SLSS IDs are used by SyncRef UEs that are synchronized to their own internal clock or to a SyncRef UE that is out of coverage of a gNB/eNB or GNSS. The SLSS IDs = 336, 337 are used by SyncRef UEs that are indirectly synchronized to GNSS ([44], [19]).

Based on the SLSS ID and I_{IC} , SyncRef UEs can be classified into five groups with different priorities:

- A first group with SyncRef UEs directly synchronized to GNSS (e.g., UE D in Fig. 13). SyncRef UEs in this group are identified by $I_{IC} = 1$ and SLSS ID = 0.
- A second group with SyncRef UEs that are out of coverage of GNSS or a gNB/eNB and are synchronized to a SyncRef UE from the first group (e.g., UE E in Fig. 13). SyncRef UEs in this group have $I_{IC} = 0$ and SLSS ID = 0.
- A third group with SyncRef UEs directly synchronized to a gNB/eNB (e.g., UE A and UE B in Fig. 13). SyncRef UEs in this group have $I_{IC} = 1$ and SLSS ID = $\{1, \dots, 335\}$.
- A fourth group with SyncRef UEs that are out of coverage of GNSS or a gNB/eNB and are synchronized to a SyncRef UE from the third group (e.g., UE C in Fig. 13). SyncRef UEs in this group have $I_{IC} = 0$ and SLSS ID = $\{1, \dots, 335\}$.

reliability of a groupcast transmission in NR V2X can be improved by considering sidelink HARQ feedback and retransmissions.

²²A channel with a high spatial correlation among the antennas has a lower rank compared to a channel with lower spatial correlation, which results in a larger rank. The rank of a channel is at most the minimum number of antennas ports at the TX UE or RX UE.

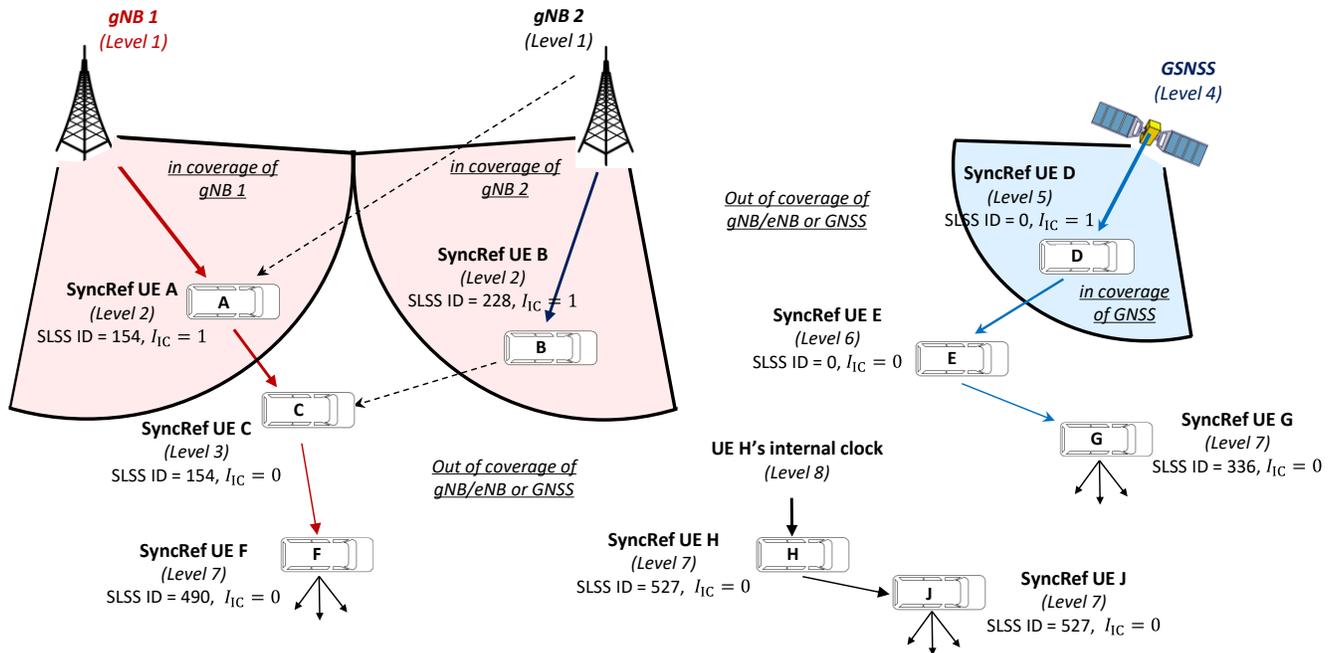


Fig. 13. Synchronization references with different priority levels.

- A fifth group with any other SyncRef UEs (e.g., UE F, UE G, UE H and UE J in Fig. 13). It includes SyncRef UEs that are out of coverage of GNSS or a gNB/eNB, and that are synchronized to their own internal clock (e.g., UE H in Fig. 13), to a SyncRef UE that uses its own internal clock or to a SyncRef UE from the second or fourth group. SyncRef UEs in this group have $I_{IC} = 0$ and $SLSS\ ID = \{336, \dots, 671\}$.

The SLSS ID and I_{IC} help distinguish whether a SyncRef UE is directly or indirectly synchronized to GNSS or a gNB/eNB. When a SyncRef UE is indirectly synchronized, the SLSS ID and I_{IC} also indicate whether the SyncRef UE is either one hop or several hops away from a gNB/eNB or GNSS, where each hop corresponds to another SyncRef UE. The larger the number of hops to GNSS or a gNB/eNB, the lower is the quality of the synchronization provided by the SyncRef UE. Thus, SyncRef UEs that are directly synchronized to GNSS (i.e., from the first group) have a higher priority as synchronization reference than those that are one or more hops away from GNSS (i.e., from the second or fifth group). In the same way when considering a gNB/eNB, SyncRef UEs in the third group have a higher priority over SyncRef UEs in the fourth group and fifth group. A priority comparison between SyncRef UEs in the first or second group with SyncRef UEs in the third or fourth group depends on whether GNSS or gNB/eNB has higher priority as synchronization source.

For the selection of a synchronization reference, NR V2X defines two sets of priority hierarchies (shown in Table VII) [45]: one with GNSS as the highest priority (GNSS-based synchronization) and another one with gNB/eNB as the highest priority (gNB/eNB-based synchronization). In both sets, SyncRef UEs are distinguished following the five groups previously described. Fig. 13 depicts examples of synchronization references according to the set of priority levels for gNB/eNB-based synchronization. Depending on

which of the two set of priorities is (pre-)configured, a UE selects as its reference the available synchronization reference with the highest priority. The selection of a gNB/eNB is based on the cell selection procedure in NR Uu [66], while the selection of GNSS is based on the reliability of GNSS [52].

TABLE VII
DIFFERENT SETS OF PRIORITIES FOR A SYNCHRONIZATION REFERENCE

Priority Level	GNSS-based synchronization	gNB/eNB-based synchronization
Level 1	GNSS	gNB/eNB
Level 2	SyncRef UE in network coverage and directly synchronized to GNSS, i.e., with $I_{IC} = 1$ and $SLSS\ ID = \{0\}$	SyncRef UE directly synchronized to gNB/eNB, i.e., with $I_{IC} = 1$ and with $SLSS\ ID = \{1, \dots, 335\}$
Level 3	SyncRef UE out of GNSS/network coverage and one hop away from GNSS, i.e., with $I_{IC} = 0$ and $SLSS\ ID = \{0\}$	SyncRef UE out of GNSS/network coverage and one hop away from a gNB/eNB, i.e., with $I_{IC} = 0$ and with $SLSS\ ID = \{1, \dots, 335\}$
Level 4	gNB/eNB	GNSS
Level 5	SyncRef UE directly synchronized to a gNB/eNB, i.e., with $I_{IC} = 1$ and with $SLSS\ ID = \{1, \dots, 335\}$	SyncRef UE directly synchronized to GNSS, i.e., with $I_{IC} = 1$ and $SLSS\ ID = \{0\}$
Level 6	SyncRef UE out of GNSS/network coverage and one hop away from a gNB/eNB, i.e., with $I_{IC} = 0$ and with $SLSS\ ID = \{1, \dots, 335\}$	SyncRef UE out of GNSS/network coverage and one hop away from GNSS, i.e., with $I_{IC} = 0$ and $SLSS\ ID = \{0\}$
Level 7	SyncRef UE out of GNSS/network coverage and two or more hops away from a gNB/eNB or GNSS, i.e., with $I_{IC} = 0$ and with $SLSS\ ID = \{336, 337, \dots, 671\}$	
Level 8	UE's own internal clock	

The selection of a SyncRef UE is based on the SLSS ID and the in coverage indicator I_{IC} . Depending on the priority level, a UE first searches for SyncRef UEs with certain SLSS IDs, i.e.,

SyncRef UEs transmitting a certain combination of S-PSS and S-SSS. For each SLSS ID, the UE measures an RSRP based on PSBCH DMRS sent by a SyncRef UE. Depending on UE implementation, the S-SSS can also be used for computing the RSRP [46]. For SLSS IDs with RSRP above a preconfigured threshold, the UE checks the value of the in coverage indicator I_{IC} carried in the PSBCH. In this way, a UE can determine the SyncRef UEs which are its candidate synchronization reference for a given priority level. For multiple candidate SyncRef UEs in the same priority level, the SyncRef UE with the highest RSRP has higher priority. Independently of the set of priorities utilized, a UE that is unable to find any other synchronization reference (i.e., GNSS, a gNB/eNB or a SyncRef UE) uses its own internal clock as synchronization reference.

Examples for the selection of synchronization reference are given next. If a UE is (pre-)configured to use gNB/eNB-based synchronization and it is out of coverage of gNB/eNB, it checks if there are SyncRef UEs with SLSS ID = {1, ..., 335}. For the detected SLSS IDs with an RSRP above a preconfigured threshold, the UE checks whether $I_{IC} = 1$. If the UE has found multiple such SyncRef UEs (i.e., with priority level 2), the UE selects the SyncRef UE with the highest RSRP as its synchronization reference. In Fig. 13, this situation can be seen for UE C, which identifies that SyncRef UE A and B have priority level 2 but different RSRP at UE C. As SyncRef UE A is closer to UE C, the RSRP of SyncRef UE A is higher than that of SyncRef UE B. Thus, UE C selects SyncRef UE A as its synchronization reference. In another example, if a UE is out of coverage of GNSS or gNB/eNB, it selects as synchronization reference the SyncRef UE with the highest RSRP within the highest priority level (among priority levels 2, 3, 5, 6 and 7).

A gNB/eNB may not be synchronized to GNSS. The use of priority levels 4 to 6 can then be disabled for GNSS-based synchronization [45]. It can also happen that the situation of a UE changes regarding available synchronization references, e.g., a higher priority synchronization reference becomes available. Thus, a UE frequently searches for synchronization references with the objective to select always the highest priority synchronization reference that is available.

6) Triggering of S-SSB Transmissions

In NR V2X, there are two general procedures for triggering S-SSB transmissions at a UE (i.e., for the UE becoming a SyncRef UE) [63]: (i) a UE is configured by the network to become a SyncRef UE; or (ii) a UE decides on its own whether to become a SyncRef UE when in or out of network coverage. A SyncRef UE sends S-SSBs based on the timing reference provided by its synchronization reference.

A UE that is in coverage of a gNB/eNB can be configured by the network to transmit or not to transmit S-SSBs. In the former case, the UE is network configured to become a SyncRef UE, while in the latter case the UE is network configured not to act as a SyncRef UE. A network configured SyncRef UE sends S-SSBs irrespective of whether it has any data to transmit in the sidelink. The SyncRef UE knows which SLSS ID and resources to use for the S-SSB transmissions based on the sidelink system information provided by the network to UEs in the cell [63].

The network configuration of a UE to send or not to send S-SSBs is optional. Consequently, a UE in network coverage may not be configured to transmit or not to transmit S-SSBs. If a UE in network coverage has not received such configuration and the UE has data to transmit in the sidelink, the UE decides on its own whether to transmit S-SSBs or not [63]. This is in contrast to the network configured scenario where the decision regarding S-SSB transmissions is taken by the network, e.g., regardless whether the UE has data to transmit in the sidelink. A UE decides on its own to become a SyncRef UE, by comparing its RSRP of the serving gNB/eNB with a threshold provided by the network. The UE measures the RSRP based on reference signals (e.g., PBCH DMRS) associated with the synchronization signal sent by the serving gNB/eNB. If the RSRP is below the provided threshold, the UE can become a SyncRef UE and transmit S-SSBs. The SyncRef UE knows which SLSS ID and resources to use for the S-SSB transmissions based on the sidelink system information provided by the network in the cell [63]. If the RSRP of the serving gNB/eNB is above or equal to the threshold, the UE does not send S-SSBs. This RSRP-based triggering of S-SSB transmissions results in having UEs close to the cell edge becoming SyncRef UEs, if they have data to transmit in the sidelink; this is illustrated in Fig. 14 for UE A. Triggering S-SSB transmissions by a UE close to cell edge allows expanding the synchronization coverage of the serving gNB/eNB as shown in Fig. 14. This allows UEs with very poor network coverage or out of network coverage (e.g., UE C in Fig. 14) to use the timing reference of the serving gNB/eNB, i.e., for SL communication with UEs inside the cell. On the other hand, there may be little or no benefit in allowing UEs which are in good network coverage (i.e., close to the gNB/eNB and with the RSRP \geq threshold) to transmit S-SSBs. The area that these UEs (e.g., UE B in Fig. 14) could cover with their S-SSB transmissions is already within the coverage of the serving gNB/eNB.

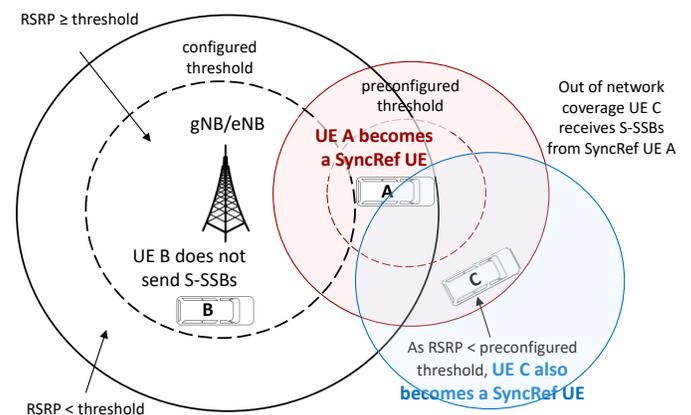


Fig. 14. Triggering of S-SSB transmissions when RSRP of synchronization reference is lower than a threshold, where the threshold is (pre-)configured.

If a UE is synchronized to a SyncRef UE and has data to transmit in the sidelink, it decides on its own whether to transmit S-SSBs or not [63]. A UE synchronized to a selected SyncRef UE can itself become a new SyncRef UE if the RSRP of its selected SyncRef UE is below a preconfigured threshold.

The RSRP is measured based on PSBCH DMRS sent in S-SSB transmissions of the selected SyncRef UE. The resources on which the new SyncRef UE sends S-SSBs are preconfigured [63], with the new SyncRef UE transmitting S-SSBs on different resources than the ones used by the selected SyncRef UE. If the selected SyncRef UE has an SLSS ID from the set of out of coverage SLSS IDs, the new SyncRef UE transmits S-SSBs with this same SLSS ID [63] (as shown for UE J in Fig. 13). If the selected SyncRef UE has an SLSS ID from the set of in coverage SLSS IDs, the new SyncRef UE transmits the S-SSBs with an SLSS ID equal to the SLSS ID of the selected SyncRef UE plus 336 [63] (as shown for UE F and UE G in Fig. 13). If the UE measures an RSRP of its selected SyncRef UE which is above or equal to the (pre-)configured threshold, the UE does not send S-SSBs. This RSRP-based triggering of S-SSB transmissions results in having UEs near the edge of the synchronization coverage becoming SyncRef UEs, if they have data to transmit in the sidelink. This enables expanding the synchronization coverage (e.g., as shown with UE C in Fig. 14) which aids nearby UEs in out of network coverage to share the same timing reference for SL communication.

If a UE has data to transmit in the sidelink and uses its internal clock as synchronization reference, the UE becomes a SyncRef UE. The SyncRef UE transmits S-SSBs with an SLSS ID randomly chosen from the set of out of coverage SLSS IDs excluding 336 and 337 [63]; this case is depicted in Fig. 13 with UE H. The resources used to send S-SSBs are preconfigured.

7) Sidelink Power Control

In NR V2X, SL power control is supported for PSSCH, PSSCH, PSFCH and S-SSB transmissions. As transmission power control (TPC) commands are not supported for NR V2X SL, the SL power control scheme is open-loop. For the SL power control, a maximum transmit power P_{MAX} is (pre-)configured at the TX UE [44]. SL power control is supported for unicast and groupcast transmissions in NR V2X.

For a unicast transmission, the PSSCH power control can be configured to use the DL pathloss PL_{DL} (between the gNB and TX UE) only, the SL pathloss PL_{SL} (between TX UE and RX UE) only, or both DL pathloss PL_{DL} and SL pathloss PL_{SL} [66]. The PSSCH power control can be based on the DL pathloss PL_{DL} when the TX UE is in network coverage. This allows mitigating the interference at the gNB (for uplink reception), similar to the SL power control in LTE V2X [75]. If the PSSCH power control is based on the DL pathloss only, TX UEs near the gNB transmit PSSCH at a lower power than TX UEs farther away from the gNB. The DL pathloss-based PSSCH power control can be enabled or disabled by the gNB [66]. The DL pathloss can be derived at the TX UE based on measurements of reference signals (e.g., CSI-RS or SSB) sent by the gNB [41].

For unicast, the PSSCH power control can also be based on the SL pathloss PL_{SL} between the TX UE and the RX UE. This allows compensating the attenuation in the SL channel. For instance, a TX UE that is far away from the gNB may transmit

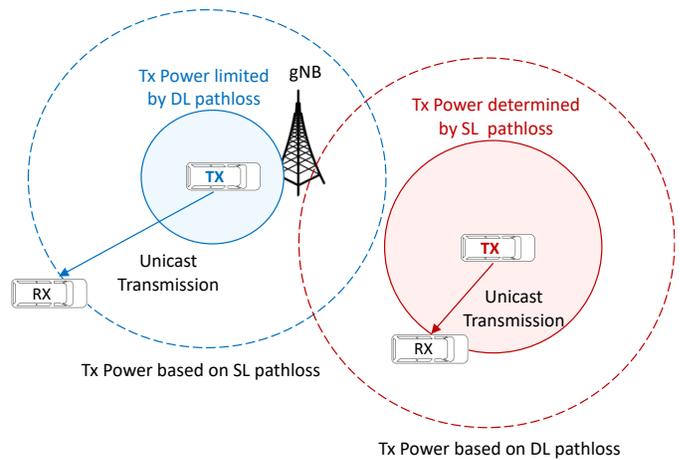


Fig. 15. Sidelink power control based on DL pathloss and SL pathloss.

PSSCH at a larger power than necessary when the PSSCH power control is configured to use DL pathloss only. However, if the PSSCH power control takes the SL pathloss also into account, this may avoid that a TX UE transmits at a large power as shown in Fig. 15. The SL pathloss-based PSSCH power control can be used when the TX UE is in or out of network coverage. The SL pathloss-based PSSCH power control can be enabled or disabled via (pre-)configuration. For this power control scheme, the TX UE requires an estimate of the SL pathloss that can be obtained from feedback of the RX UE [44]. Based on PSSCH DMRS transmitted by the TX UE, the RX UE can obtain an average RSRP over several RSRP measurements [76] in order to mitigate fluctuations on the received power. The RX UE cannot derive the SL pathloss based on the RSRP measurements since the transmit power of the PSSCH DMRS is not indicated to the RX UE [44]. Thus, the RX UE feeds back the average RSRP to the TX UE using higher layer signaling. The TX UE uses the fed back average RSRP along with the average transmit power of the PSSCH DMRS [49] to derive the sidelink pathloss PL_{SL} (in dB) as follows:

$$PL_{SL} = \text{Average Tx power PSSCH DMRS (in dBm)} - \text{Average RSRP (in dBm)}. \quad (2)$$

For a groupcast transmission, the PSSCH power control can be configured to use the DL pathloss. Although the use of SL pathloss for groupcast PSSCH power control was discussed during the development of Rel. 16 NR V2X [77], it was finally agreed that SL pathloss-based PSSCH power control is not supported for groupcast in NR V2X. For groupcast PSSCH power control based on SL pathloss, RX UEs need to feed back their RSRP to the TX UE, which can lead to a large overhead²³.

When the PSSCH power control is configured to use both the DL pathloss and the SL pathloss, the transmit power for PSSCH is determined at the TX UE as follows (in dBm):

$$P_{PSSCH,1} = \min(P_{MAX}, P_{0,DL} + 10 \log_{10}(2^{\mu} M_{PSSCH}) + \alpha_{DL}, PL_{DL}, P_{0,SL} + 10 \log_{10}(2^{\mu} M_{PSSCH}) + \alpha_{SL} PL_{SL}) \quad (3)$$

²³To address this issue, some schemes were proposed to group RX UEs for their feedback. This enables several RX UEs to share a feedback resource [77]. The grouping of the RX UEs could be done based on the distance and/or RSRP.

TABLE VIII
POWER CONTROL FOR THE DIFFERENT CHANNELS AND TRANSMISSIONS IN NR V2X

Physical Channel or Transmission		Power control expression
PSSCH	PSSCH symbols without PSCCH	When in network coverage and configured to use DL pathloss (for unicast and groupcast): $P_{\text{PSSCH},1} = \min(P_{\text{MAX}}, P_{0,\text{DL}} + 10 \log_{10}(2^{\mu} M_{\text{PSSCH}}) + \alpha_{\text{DL}} PL_{\text{DL}})$ [dBm]
		When in or out of network coverage and configured to use SL pathloss (for unicast): $P_{\text{PSSCH},1} = \min(P_{\text{MAX}}, P_{0,\text{SL}} + 10 \log_{10}(2^{\mu} M_{\text{PSSCH}}) + \alpha_{\text{SL}} PL_{\text{SL}})$ [dBm]
	When in network coverage and configured to use DL pathloss and SL pathloss (for unicast): $P_{\text{PSSCH},1} = \min(P_{\text{MAX}}, P_{0,\text{DL}} + 10 \log_{10}(2^{\mu} M_{\text{PSSCH}}) + \alpha_{\text{DL}} PL_{\text{DL}}, P_{0,\text{SL}} + 10 \log_{10}(2^{\mu} M_{\text{PSSCH}}) + \alpha_{\text{SL}} PL_{\text{SL}})$ [dBm]	
	PSSCH symbols with PSCCH	$P_{\text{PSSCH},2} = P_{\text{PSSCH},1} + 10 \log_{10}\left(\frac{M_{\text{PSSCH}} - M_{\text{PSCCH}}}{M_{\text{PSSCH}}}\right)$ [dBm]
PSCCH		$P_{\text{PSCCH}} = P_{\text{PSSCH},1} + 10 \log_{10}\left(\frac{M_{\text{PSCCH}}}{M_{\text{PSSCH}}}\right)$ [dBm]
PSFCH	When in network coverage:	$P_{\text{PSFCH}} = \min(P_{\text{MAX}}, P_{0,\text{DL}}^{\text{PSFCH}} + 10 \log_{10}(2^{\mu}) + \alpha_{\text{DL}}^{\text{PSFCH}} PL_{\text{DL}})$ [dBm]
	When out of network coverage:	$P_{\text{PSFCH}} = P_{\text{MAX}}$ [dBm]
S-SSB	When in network coverage:	$P_{\text{S-SSB}} = \min(P_{\text{MAX}}, P_{0,\text{DL}}^{\text{S-SSB}} + 10 \log_{10}(2^{\mu} \cdot 11) + \alpha_{\text{DL}}^{\text{S-SSB}} PL_{\text{DL}})$ [dBm]
	When out of network coverage:	$P_{\text{S-SSB}} = P_{\text{MAX}}$ [dBm]

where the second and third term in the argument of the min expression are associated with the DL pathloss and SL pathloss, respectively. The power control parameters (i.e., the nominal power P_0 and the parameter α used for fractional power control [78]), are (pre-)configured separately when considering the DL pathloss and SL pathloss: $P_{0,\text{DL}}$ and α_{DL} are associated with the DL pathloss and $P_{0,\text{SL}}$ and α_{SL} are associated with the SL pathloss. μ is the SCS configuration factor of the configured numerology. M_{PSSCH} is the number of PRBs for the PSSCH (in symbols without PSCCH), where $M_{\text{PSSCH}} = L_{\text{subCH}} \cdot M_{\text{sub PRBs}}$ with L_{subCH} sub-channels for PSSCH and $M_{\text{sub PRBs}}$ PRBs per sub-channel. If the sidelink power control is not based on the DL pathloss or the SL pathloss, then the second or the third term, respectively, is not included in the argument of the min expression. As mentioned before, SL pathloss-based PSSCH power control is only supported for unicast transmissions. In case both the DL pathloss and SL pathloss are disabled for the PSCCH power control, the transmit power for PSSCH is equal to the (pre-)configured maximum transmit power P_{MAX} .

For transmissions with two streams in PSSCH, the PSSCH transmit power is equally shared between the two streams [49]. The expression in (3) represents the PSSCH transmit power for PSSCH symbols without PSCCH. In NR V2X, a TX UE transmits PSSCH and PSCCH with the same power spectral density (i.e., with the same power over a PRB) in all symbols with PSCCH, PSSCH or PSCCH/PSSCH [49]. As a result, the transmit power $P_{\text{PSSCH},1}$ given by (3) is shared between PSSCH and PSCCH in the PSCCH/PSSCH symbols. With M_{PSCCH} PRBs for PSCCH, there are $M_{\text{PSSCH}} - M_{\text{PSCCH}}$ PRBs available for PSSCH in a PSCCH/PSSCH symbol. Based on this, the PSSCH transmit power $P_{\text{PSSCH},2}$ and PSCCH transmit power P_{PSCCH} in a PSCCH/PSSCH symbol is given by (in dBm) [66]:

$$P_{\text{PSSCH},2} = P_{\text{PSSCH},1} + 10 \log_{10}\left(\frac{M_{\text{PSSCH}} - M_{\text{PSCCH}}}{M_{\text{PSSCH}}}\right) \quad (4)$$

$$P_{\text{PSCCH}} = P_{\text{PSSCH},1} + 10 \log_{10}\left(\frac{M_{\text{PSCCH}}}{M_{\text{PSSCH}}}\right). \quad (5)$$

Reference signals sent in the PSSCH region and PSCCH region of the slot (i.e., SL CSI-RS, SL PT-RS, PSCCH DMRS, PSCCH DMRS) are also sent with the same power spectral density based on $P_{\text{PSSCH},1}$ [49]. The different variations for the PSSCH power control and the resulting PSCCH power control are summarized in Table VIII.

For a UE transmitting PSFCH, the PSFCH power control can be based on the pathloss between the gNB and the UE (i.e., the DL pathloss) if the UE is in network coverage [66]. While using the SL pathloss may avoid transmitting at a power larger than necessary, SL pathloss-based PSFCH power control is not supported in NR V2X [44]. For a SL pathloss-based power control for PSFCH, the RX UE(s) of a transmission need to know the SL pathloss to the TX UE. However, as a TX UE does not indicate its SL transmit power, an RX UE cannot derive the SL pathloss and cannot perform SL pathloss-based power control for PSFCH in NR V2X. The PSFCH power control parameters associated with the DL pathloss (i.e., $P_{0,\text{DL}}^{\text{PSFCH}}$ and $\alpha_{\text{DL}}^{\text{PSFCH}}$) are configured separately from the parameters used for the PSCCH/PSSCH power control [44]. The transmit power for PSFCH is determined at the UE as follows²⁴ (in dBm) [66]:

$$P_{\text{PSFCH}} = \min(P_{\text{MAX}}, P_{0,\text{DL}}^{\text{PSFCH}} + 10 \log_{10}(2^{\mu}) + \alpha_{\text{DL}}^{\text{PSFCH}} PL_{\text{DL}}). \quad (6)$$

If a UE sending a PSFCH is out of network coverage, it sends PSFCH with the (pre-)configured maximum transmit power P_{MAX} as shown in Table VIII. If a UE is to send multiple PSFCH

²⁴A PSFCH consists of one PRB. The bandwidth factor is then $10 \log_{10}(2^{\mu})$.

simultaneously, the power is equally shared among them in order of priority of the received PSSCH (see Section V.B.2)), up to the maximum number of PSFCHs the UE can transmit.

For a SyncRef UE, the S-SSB power control can be based on the pathloss between the gNB and the SyncRef UE (i.e. the DL pathloss) if the SyncRef UE is in coverage. The S-SSB power control parameters associated with the DL pathloss (i.e., $P_{0,DL}^{S-SSB}$ and α_{DL}^{S-SSB}) are configured separately from the PSSCH or PSFCH power control parameters. The SyncRef UE sends S-SSBs with transmit power²⁵ (in dBm) [66]:

$$P_{S-SSB} = \min(P_{MAX}, P_{0,DL}^{S-SSB} + 10 \log_{10}(2^{\mu} \cdot 11) + \alpha_{DL}^{S-SSB} PL_{DL}). \quad (7)$$

In case the SyncRef UE is out of network coverage, the SyncRef UE sends S-SSBs with the (pre-)configured maximum transmit power P_{MAX} as shown in Table VIII.

In NR V2X sidelink, a UE transmits S-SSBs, PSSCH, PSCCH, or PSFCH with the same transmit power in all the corresponding symbols [66]. This is also the case when the UE transmits simultaneously on the uplink and sidelink. When a TX UE transmits simultaneously on the sidelink and uplink, it adjusts the transmit power based on the transmission that is prioritized [66]. If the sidelink is prioritized, the TX UE reduces the uplink transmit power before the start of the transmission if the UE's total transmit power would be larger than P_{MAX} when the sidelink and uplink transmissions overlap. The same process is inversely applied if the uplink is prioritized. In either case, it is up to UE implementation how to reduce the uplink or sidelink transmit power. Sidelink transmissions can also be dropped in some overlapping symbols when the uplink transmission is prioritized and the UE cannot achieve the same sidelink transmit power in the symbols. It is up to UE implementation the symbols that are dropped within the overlapping symbols. If a UE is not capable to simultaneously transmit on the sidelink and uplink, it transmits only on the link with highest priority.

VI. RESOURCE ALLOCATION FOR 5G NR V2X SIDELINK

Rel. 16 defines two new modes (modes 1 and 2) for the selection of sub-channels in NR V2X SL communications using the NR V2X PC5 interface. These two modes are the counterparts to modes 3 and 4 in LTE V2X (Section II). However, LTE V2X only supports broadcast SL communications while NR V2X supports broadcast, groupcast, and unicast SL communications.

A. Mode 1

Similar to mode 3 in LTE V2X, the gNB or eNB assigns and manages the SL radio resources for V2V communications under mode 1 using the NR (or LTE) Uu interface. UEs must therefore be in network coverage to operate using mode 1. SL radio resources can be allocated from licensed carriers dedicated to SL communications or from licensed carriers that share resources between SL and UL communications. The SL

radio resources can be configured so that mode 1 and mode 2 use separate resource pools. The alternative is that mode 1 and mode 2 share the resource pool. Pool sharing can result in a more efficient use of the resources, but it is prone to potential collisions between mode 1 and mode 2 transmissions. To solve this, mode 1 UEs notify mode 2 UEs of the resources allocated for their future transmissions as it is described below.

Mode 1 uses dynamic grant (DG) scheduling like LTE V2X mode 3, but replaces the semi-persistent scheduling in LTE V2X mode 3 with a configured grant scheduling [45]. With DG, mode 1 UEs must request resources to the base station for the transmission of every single TB (and possible blind or HARQ retransmissions²⁶) [44]. To this aim, the UEs must send a Scheduling Request (SR) to the gNB using the PUCCH, and the gNB responds with the DCI over the PDCCH. The DCI indicates the SL resources (i.e., the slot(s) and sub-channel(s)) allocated for the transmission of a TB and up to 2 possible retransmissions of this TB. The UE informs other UEs about the resources it will use to transmit a TB and up to 2 possible retransmissions using the 1st-stage SCI. Nearby UEs operating under mode 2 can then know which resources UEs in mode 1 will utilize.

Requesting resources for each TB increases the delay. Mode 1 includes the configured grant scheduling option to reduce the delay by pre-allocating SL radio resources. With this scheme, the gNB can assign a set of SL resources to a UE for transmitting several TBs. This set is referred to as a configured grant (CG). The UE sends first a message with UE assistance information to the gNB indicating information about the expected SL traffic including: periodicity of TBs, TB maximum size and QoS information. The QoS information includes KPIs such as the latency and reliability required by the TBs and their priority. This information is used by the gNB to create, configure and allocate a CG to the UE that satisfies the requirements of the SL traffic. The CG is configured using a set of parameters that includes the CG index, the time-frequency allocation and the periodicity of the allocated SL resources. Fig. 16 shows two examples of configured CGs. The two CGs have different periodicities and different time-frequency allocations. The time-frequency allocation indicates the slot(s) and sub-channel(s) that are assigned periodically to the UE in a CG. A UE can be assigned a maximum of 3 SL resources during each period of the CG. The UE informs other vehicles of the resources allocated by the gNB for a particular CG period using again the 1st-stage SCI²⁷. The UE can decide how to use the SL resources of an assigned CG. However, it can only transmit one new TB in each CG period. The SL resources of a CG period can also be used to retransmit the new TB transmitted in this CG, or to retransmit other TBs initially transmitted in previous CG periods. The maximum number of retransmissions per TB in a CG is associated with the priority of the TB. HARQ retransmissions in a CG are only possible if the resource pool of the CG has a PSFCH configured by the gNB (see Section

²⁵An S-SSB consists of 11 PRBs. The bandwidth factor is then $10 \log_{10}(2^{\mu} \cdot 11)$.

²⁶This is in contrast to mode 3 that only allows one blind retransmission.

²⁷This SCI only informs about the resources allocated in the current CG period. Extensions to inform about the resources allocated in the next CG period are being studied in [119].

V.B.4). It should be noted that a gNB can assign multiple SL CGs to a UE. The configuration of each CG can be adapted to the characteristics or demands of different V2X applications. However, the transmission and retransmissions of a TB must always utilize resources of a single CG [44].

Mode 1 defines two types of CG schemes for SL: CG type 1 and CG type 2. Both are configured using Radio Resource Control (RRC) signaling (Section IX.B). CG type 1 can be utilized by the UE immediately until it is released by the base station (also using RRC signalling). SL CG type 2 can be used only after it is activated by the gNB and until it is deactivated. To this aim, the gNB notifies of the activation and deactivation using DCI signalling. The DCI also includes the CG index and time-frequency allocation of CG type 2. CG type 2 can configure multiple CGs for a UE and only activate a subset of the CGs based on the UE needs. Resources in non-active CGs can be allocated to other UEs. CG type 1 can also configure multiples CGs. However, it forces UEs to activate CGs at the time of their configuration. CG type 1 reduces the signalling and the time needed to initiate a transmission compared to CG type 2. However, if any of the CGs type 1 are not used by the UE, the resources cannot be allocated to other UEs.

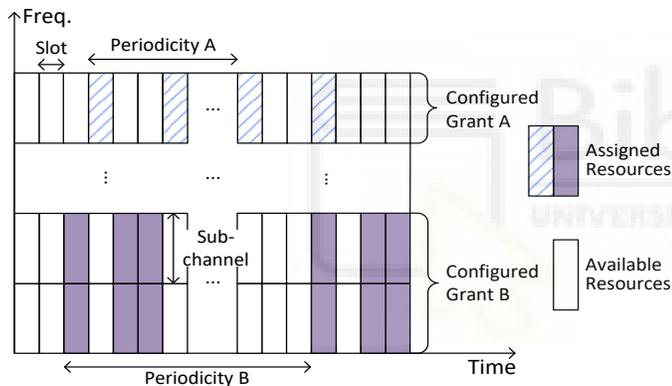


Fig. 16. Examples of CGs under NR V2X SL mode 1.

Fig. 17 shows an example that compares how transmissions are organized and scheduled when utilizing the DG and CG schemes to transmit two TBs (TB₁ generated at t_0 and TB₂ generated at t_5). We discuss first the case of the DG scheme. The UE sends an SR at t_1 to request resources for transmitting TB₁. The gNB responds with a DCI at t_3 that indicates the resources the user can utilize at t_4 . The same process takes place starting at t_5 when TB₂ is generated. TB₂ can be transmitted using the next set of resources allocated at t_{10} . For CG, the UE does not request resources but waits until the gNB grants it a CG. We suppose the UE receives the granted CG at t_4 . This CG includes a set of resources assigned periodically to the UE at t_4 , t_8 and t_{11} . The time period is adjusted to the time between TBs indicated by the UE in the provided UE assistance information. The UE uses the resources at t_4 to transmit TB₁ and the resources at t_8 to transmit TB₂. The CG scheme reduces the time needed to transmit the two TBs compared to DG. However, the DG scheme can utilize resources more efficiently when handling non periodic traffic since resources are only allocated when needed to transmit TBs.

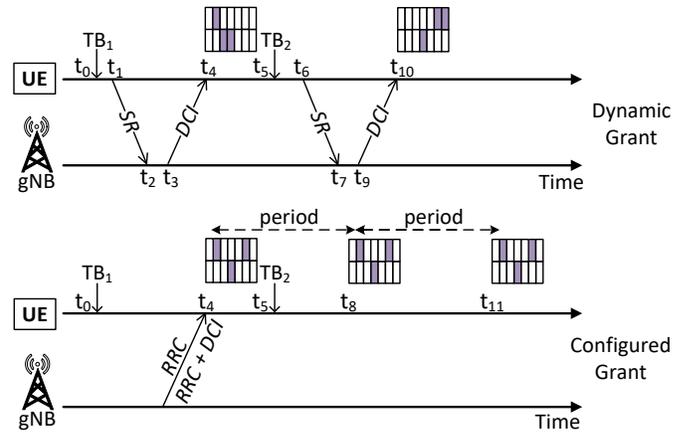


Fig. 17. Mode 1 dynamic and configured grant schemes.

Mode 1 defines three different MCS tables that are used to determine the MCS of a TB (Table 5.1.3.1-1, Table 5.1.3.1-2 and Table 5.1.3.1-3) in Section 5.1.3.1 of [42]. There are different ways for determining the MCS of a TB. The gNB can configure a UE to use one or more MCSs. In the latter case, the UE autonomously selects the MCS. The gNB can also configure a UE to use one, two or three of the MCS tables. In this case, the UE again autonomously selects the MCS from the configured tables. The UE indicates the selected MCS and MCS table (when applicable) for a TB in the associated 1st-stage SCI.

In mode 1, the gNB can enable HARQ retransmissions for groupcast and unicast communications. In this case, a RX UE can use a PSFCH associated with a received PSSCH to send HARQ feedback as described in Section V.C.2). The TX UE can inform the gNB about the feedback received from RX UEs with a feedback report. The gNB configures a PUCCH for the TX UE to send the feedback report (Section 16.5 in [66]). This feedback report in the PUCCH assists the gNB in the allocation of SL resources. The feedback report informs whether the SL transmissions of the last allocated SL resources were successful or not. Therefore, it is transmitted in the PUCCH configured after the PSFCH associated to the last allocated resource in the case of DG. For the case of the CG, the feedback report is transmitted in the PUCCH configured after the PSFCH associated to the last resource allocated for each CG period. It should be noted that the feedback report uses a single bit of the PUCCH to inform the gNB about the success of the potentially last three SL transmissions (both for DG, or a CG period). When the feedback report includes information of the transmission and retransmissions of the same TB, the TX UE reports an ACK in the PUCCH only if all RX UEs receive correctly at least one of the (re-)transmissions and they positively acknowledge it to the TX UE. Otherwise the TX UE reports a NACK in the PUCCH. When the feedback report includes information of the transmission and retransmissions of different TBs, the TX UE only reports an ACK if all RX UEs receive correctly the different TBs and they positively acknowledge them to the TX UE. After receiving the feedback report, the gNB evaluates if it has to allocate new SL resources to the TX UE for additional HARQ retransmissions.

B. Mode 2

Like with mode 4 in LTE V2X, UEs can autonomously select their SL resources (one or several sub-channels) from a resource pool when using mode 2 in NR V2X. In this case, UEs can operate without network coverage. The resource pool can be (pre-)configured by the gNB or eNB when the UE is in network coverage. Mode 2 and mode 4 differ on the scheduling scheme. Mode 4 operates following the sensing-based SPS scheme explained in Section II. On the other hand, mode 2 can operate using a dynamic or a semi-persistent scheduling scheme that differs from the one designed for mode 4. The dynamic scheme selects new resources for each TB and can only reserve resources for the retransmissions of that TB. We should note that we distinguish in this section between a selected resource (using dynamic or semi-persistent schemes) and a reserved resource. A reserved resource is a selected resource that a UE reserves for a future transmission by notifying neighboring UEs using the 1st-stage SCI. A UE can select and reserve resources for the transmission of several TBs (and their retransmissions) when utilizing the semi-persistent scheme [45]. It is important to note that the semi-persistent scheme can be enabled or disabled in a resource pool by (pre-)configuration.

Mode 2 uses almost the same procedure to select resources for the dynamic and semi-persistent schemes [45]. The differences result from the fact that the dynamic scheme only selects resources for a TB while the semi-persistent scheme selects resources for a number of consecutive *Reselection Counter* TBs. The time period between the resources selected for the transmission of consecutive TBs in the semi-persistent scheme is defined by the Resource Reservation Interval (*RRI*). The possible values of the *RRI* are $\{0, [1:99], 100, 200, 300, 400, 500, 600, 700, 800, 900, 1000\}$ ms. It should be emphasized that NR V2X mode 2 provides higher flexibility to fit the requirements and characteristics of different eV2X services by allowing any integer *RRI* between 1 and 99 ms. This is in contrast to LTE V2X mode 4 that only considered the following *RRI* values: $\{0, 20, 50, 100, 200, 300, 400, 500, 600, 700, 800, 900, 1000\}$ ms. Like in mode 4, a list of permitted *RRI*s is (pre-)configured in the resource pool for mode 2 [63]. The list can include a maximum of 16 different *RRI*s. A UE selects an *RRI* from the list when it selects new SL resources²⁸. The selection of the *RRI* is left to UE implementation. The UE can select the *RRI* that best suits the characteristics of the traffic to be transmitted. The selected *RRI* also determines the *Reselection Counter* that is randomly set within an interval that depends on the selected *RRI*. If $RRI \geq 100$ ms, this counter is randomly set within the interval $[5, 15]$. If $RRI < 100$ ms, the counter is randomly set within the interval $[5 * C, 15 * C]$ where $C = 100 / \max(20, RRI)$ (Section 5.22.1 in [74]).

A UE can select new SL resources when it generates a new TB²⁹. A new selection can also be triggered (only for the semi-persistent scheme) because a new TB does not fit in the previously reserved resources. To select new SL resources (for

both dynamic and semi-persistent schemes), a UE first defines the selection window where it looks for candidate resources to transmit a TB. The selection window includes all resources within the range of slots $[n+T_1, n+T_2]$ (Fig. 18) (Section 8.1.4 in [42]), where n is the resource (re-)selection trigger or slot at which new resources must be selected. T_1 is the processing time (in slots) required by a UE to identify candidate resources and select new SL resources for transmission. T_1 is equal to or smaller than $T_{proc,1}$. $T_{proc,1}$ is equal to 3, 5, 9 or 17 slots for a SCS of 15, 30, 60 or 120 kHz, respectively³⁰. The value of T_2 is left to UE implementation but must be included within the range $T_{2min} \leq T_2 \leq PDB$, where PDB is the Packet Delay Budget (in slots). PDB is the latency deadline by which the TB must be transmitted. The deadline is established by the V2X application generating the packet to be transmitted in the TB. The value of T_{2min} depends on the priority of the TB and the SCS. Possible values of T_{2min} include $\{1, 5, 10, 20\} * 2^\mu$ slots, where μ is the SCS configuration factor as defined in Section V.A.1). This results in values of T_{2min} equivalent to $\{1, 5, 10, 20\}$ ms. NR V2X mode 2 can then guarantee minimum latency levels of 1 ms compared to 10 ms in LTE V2X mode 4.

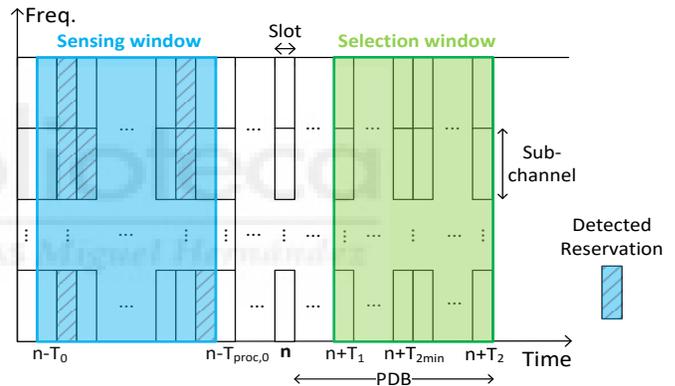


Fig. 18. Sensing and selection windows of NR V2X SL mode 2 when $T_2 = PDB$.

Once the selection window is defined, the UE must identify the candidate resources within the selection window. A candidate resource is defined by a slot in time and L_{PSSCH} contiguous sub-channels in frequency. L_{PSSCH} is an integer in the range $1 \leq L_{PSSCH} \leq \max(L_{PSSCH})$, where $\max(L_{PSSCH})$ is the total number of sub-channels per slot in the selection window. However, the value of $\max(L_{PSSCH})$ can be modified by the congestion control process explained in Section VI.C. L_{PSSCH} must be selected such that the TB and its associated SCI (including the 1st-stage and 2nd-stage SCI)³¹ fit in the candidate resource. L_{PSSCH} therefore depends on the size of the TB and the SCI as well as on the utilized MCS. For mode 2, a UE is (pre-)configured to use one, two or three of the MCS tables that are supported for TB and mode 1 (see Section VI.A). The UE selects the MCS from the (pre-)configured MCS tables. The UE indicates the MCS used for the TB in the associated 1st-stage SCI. The 1st-stage SCI also includes the selected MCS table if

²⁸ This does not apply for the selection of new SL resources with the re-evaluation or pre-emption mechanisms that are explained below.

²⁹ It may also select new SL resources in mode 2 with the re-evaluation or pre-emption mechanisms that are explained below.

³⁰ $T_{proc,1}$ is equivalent to 3, 2.5, 2.25 or 2.125 ms for a SCS of 15, 30, 60 or 120 kHz, respectively.

³¹ The TB and SCI are transmitted in the same slot as discussed in Section V.

the UE is (pre-)configured with more than one table.

When a UE is not transmitting, it senses the SL resources during the sensing window (Fig. 18) in order to identify available candidate resources. The sensing window is the time interval defined by the range of slots $[n-T_0, n-T_{proc,0})$, where n is the resource (re-)selection trigger or slot at which new resources must be selected. T_0 is an integer defined in number of slots that depends on the SCS configuration but that must be set to a value (in number of slots) equivalent to 1100 ms or 100 ms. The selected value is determined by the (pre-)configuration of the resource pool. $T_{proc,0}$ is the time required to complete the sensing procedure and is equal to one slot for a SCS of 15 or 30 kHz, and equal to 2 or 4 slots for a SCS of 60 or 120 kHz, respectively³². During the sensing process, the UE decodes the 1st-stage SCI received from other UEs in the sensed SL resources. The 1st-stage SCIs indicate the SL resources that other UEs have reserved for their TB and SCI transmissions in the PSSCH and PSCCH. In particular, the 1st-stage SCIs can indicate the SL resources reserved for retransmissions of the TB associated to the 1st-stage SCI, and resources reserved for the initial transmission and retransmissions of the next TB. The UE also measures the RSRP [45] of the transmissions associated to the 1st-stage SCIs received from other UEs. The UE stores the sensed information (the decoded 1st-stage SCI and the RSRP measurements) and uses it to determine which candidate resources from the selection window should be excluded when a new selection is triggered.

To select new resources, mode 2 defines a 2-step algorithm that is used by both the dynamic and semi-persistent scheduling schemes [44] with small differences that are highlighted when applicable. Step 1 excludes candidate resources in the selection window. The first exclusions relate to the half-duplex operation since a UE cannot sense the reservations from other UEs announced in the slots of the sensing window where the UE was transmitting. Let us suppose the UE was transmitting in a slot s_i in the sensing window. The UE could not sense the reservations of other UEs at slot s_i , and it excludes all candidate resources at any slot $s_i+q*RRI_i$ in the selection window in case other UEs have reserved them. RRI_i represents all the possible values of the RRI (represented here in slots) based on the list of permitted $RRIs$ in the resource pool. The parameter q is an integer that takes values in the range $1 \leq q \leq Q$. Q is equal to $\left\lceil \frac{T_2}{RRI_i} \right\rceil$ when the following two conditions are fulfilled: 1) $RRI_i < T_2$ (RRI_i and T_2 in ms) and 2) $n-s_i \leq RRI_i$ (RRI_i in slots), where n is the resource (re-)selection trigger or slot at which new resources must be selected and the value of T_2 (in slots) is left to UE implementation but must be included within the range $T_{2min} \leq T_2 \leq PDB$. Otherwise Q is set equal to 1. Q indicates the estimated number of periodic transmissions from other UEs for each potential RRI . If the UE operates with the semi-persistent scheme, it should exclude additionally all candidate resources

in any slot s_j in the selection window such that any slot $s_j+j*RRI_{TX}$ overlaps with any slot $s_i+q*RRI_i$ where: RRI_{TX} is the RRI (in slots) selected by the UE that is selecting new resources, s_i is the slot during which the UE was transmitting in the sensing window, j is an integer that can take values in the range $1 \leq j \leq 10*ReselectionCounter-1$ ³³, and RRI_i represents all the possible values of the RRI (in slots) based on the list of permitted $RRIs$ in the resource pool (Section 8.1.4 in [42]).

The UE also excludes the candidate resources based on the reservations received from other UEs in the 1st-stage SCIs detected during the sensing window. In this case, the candidate resources are only excluded if the UE has measured an RSRP associated with the reservation that is higher than an RSRP threshold. Note that a list of RSRP thresholds is (pre-)configured in the resource pool. The utilized RSRP threshold depends on the priority of the TB for which the UE is selecting new resources and on the priority (included in the 1st-stage SCI) of the other UE that reserved the resource [44]. The procedure to exclude resources based on the resources reserved by other UEs for transmitting their next TB can be formally expressed as follows. Let us consider that a 1st-stage SCI was received (with a measured RSRP higher than the corresponding RSRP threshold) by the UE in a slot s_k in the sensing window, and that the 1st-stage SCI is reserving a resource R_x for the slot s_k+RRI_{RX} . RRI_{RX} is the RRI included in the 1st-stage SCI and it is represented in number of slots³⁴. Then, the UE excludes every candidate resource that overlaps in frequency with R_x in any slot $s_k+q*RRI_{RX}$, where q is the variable previously defined but substituting s_j by s_k and RRI_i by RRI_{RX} in the related expressions. If the UE that is selecting the SL resources is operating with the semi-persistent scheme, it excludes additionally every candidate resource that overlaps in frequency with R_x and is located in any slot s_p in the selection window such that any slot $s_p+j*RRI_{TX}$ overlaps with any slot $s_k+q*RRI_{RX}$ (Section 8.1.4 in [42]). Remember that RRI_{TX} is the RRI (in slots) selected by the UE that is selecting new resources, and j is an integer that can take values in the range $1 \leq j \leq 10*ReselectionCounter-1$. This exclusion is applied so that the UE avoids that any of its following $10*ReselectionCounter-1$ transmissions collide with a potential transmission from the UE that sent the 1st-stage SCI.

After executing all exclusions in step 1, the UE (with dynamic or semi-persistent scheme) checks whether the percentage of remaining available candidate resources (i.e., those that have not been excluded in step 1) in the selection window is equal or higher than $X\%$. If not, the RSRP thresholds are increased by 3 dB, and the process is repeated iteratively until the percentage of available candidate resources in the selection window is at least equal to $X\%$. Possible values of X are 20, 35 or 50. The selected value depends on the priority of the TB for which the UE is selecting new SL resources.

In Step 2, the UE randomly selects the SL resource from the list of available candidate resources. Mode 2 eliminates the

³² $T_{proc,0}$ is equivalent to 1 ms for a SCS of 15 kHz and 0.50 ms for the rest of SCS configurations.

³³ Note that even though the UE would perform a number of consecutive *Reselection Counter* transmissions, j can take up to 10 times this value. This is the case to avoid potential collisions with transmissions from other UEs in case

the UE transmits *Reselection Counter* TBs and then maintains the same selected resources based on the probability parameter P that is defined later.

³⁴ Note that in contrast to the previous procedure that considered the list of $RRIs$ (i.e. RRI_i) to exclude candidate resources, in this case only the RRI received in the 1st-stage SCI is taken into account.

third step included in mode 4 in LTE V2X for selecting resources. This third step estimated the average RSSI of the candidate resources to select one from a list of candidate resources with the lowest RSSI.

A UE selects N candidate resources ($N \leq N_{MAX}$ with $1 \leq N_{MAX} \leq 32$) within the same selection window for the initial transmission of a TB and its $N-1$ blind or potential HARQ retransmissions³⁵. The process to select N candidate resources applies to both the dynamic and semi-persistent schemes and follows the 2-step algorithm previously explained. A UE can only perform HARQ retransmissions for groupcast or unicast transmissions if a PSFCH is (pre-)configured in the resource pool (see Section V.C.1). If not, only blind retransmissions are possible. We should note that if a TB is positively acknowledged when using HARQ retransmissions, the TX UE drops all next scheduled retransmissions of the TB in the selected resources (including the reserved ones) (Section 5.22.1 in [74]). The value of N is left to UE implementation. The value of N_{MAX} is (pre-)configured and can vary depending on the channel utilization or load following the congestion control mechanism (see Section VI.C). The UE cannot select for N a value higher than the number of available candidate resources after step 1. Once the UE has selected the value of N , it 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 resource reservations located within a window W of 32 slots, which conditions the candidate resources that can be selected. The first slot of the window W is where the 1st-stage SCI is transmitted. In addition, a 1st-stage SCI can only notify about a maximum number of N_{SCI} resources. The maximum value of N_{SCI} (i.e. $\max(N_{SCI})$) is (pre-)configured in the resource pool and can be equal to 2 or 3 resources. A 1st-stage SCI therefore notifies about the transmission (or retransmission) of the TB happening in the same slot as the transmission of this 1st-stage SCI, and $N_{SCI}-1$ reserved resources for the following $N_{SCI}-1$ retransmissions of the TB (Fig. 19) (Section 8.1.5 in [42]).

Once a UE has identified the list of available candidate resources following the step 1 previously described, it executes step 2 to select the N candidate resources. To do so, the UE first selects randomly one of the N candidate resources³⁶. Let us consider that the first candidate resource is selected at the slot m_1 (see Fig. 19). Then, the UE selects also randomly the second candidate resource, but with the restriction that the gap between this candidate resource and the first selected candidate resource must be smaller than a window W of 32 slots. This means that the second candidate resource will be located (at the slot m_2) within the range of slots $[m_1-31, m_1+31]$. This guarantees that the 1st-stage SCI of one of the two selected candidate resources is able to reserve the other candidate resource. If $N > 2$, the UE selects also randomly the third candidate resource but with the restriction that it is located (at the slot m_3) within the range $[m_1-$

$31, m_1+31]$ or the range $[m_2-31, m_2+31]$. This chain procedure is repeated with the aim that all candidate resources for the TB retransmissions can be reserved by a previous SCI. However, if the UE can only select a subset of the N candidate resources following this procedure, the remaining ones are selected randomly within the selection window even if they do not meet the previous 1st-stage SCI limitations. It is also important to note that 3GPP standards specify that there should be a minimum time gap t_{GAP} (in number of slots) between any pair of consecutive selected resources for HARQ retransmissions so that the TX UE can receive and process the acknowledgements from the RX UEs, and prepare the next HARQ retransmission (Section 5.22.1 in [74]). This minimum gap does not apply to blind retransmissions if a PSFCH is not (pre-)configured in the resource pool. If it is (pre-)configured, the gap applies although blind retransmissions do not use HARQ feedback.

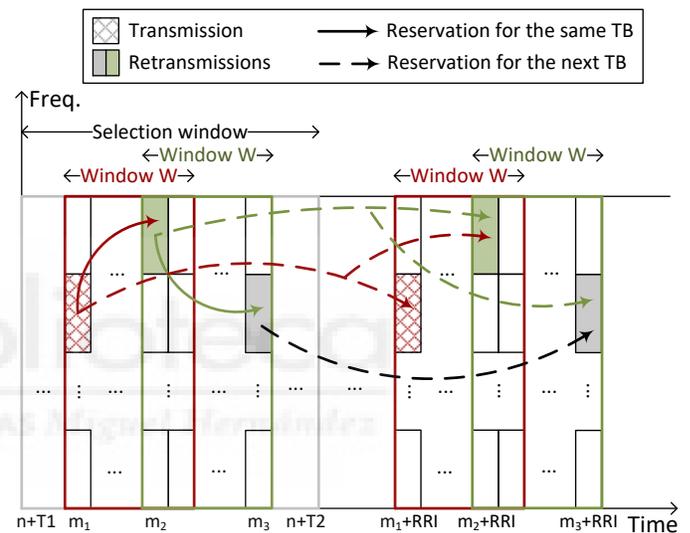


Fig. 19. Resource reservation for the retransmissions of the same TB and for the transmission and retransmissions of the next TB in NR V2X mode 2 (illustrative example when $N=3$, $\max(N_{SCI})=3$ and RRI (in number of slots)). The transmission of the TB at m_1 does not announce the reservation of the retransmission at m_3 because it is outside of the window W . This reservation for the second retransmission is announced by the first retransmission at m_2 .

When a UE operates with the dynamic scheme, it needs to select new SL resources for the transmission of each TB and its retransmissions. On the other hand, a UE selects resources for a number of consecutive *Reselection Counter* TBs in case of the semi-persistent scheme. Remember that the time period between the selected resources for consecutive TBs is determined by the *RRI*. The *Reselection Counter* is randomly chosen every time new resources must be selected³⁷. The *Reselection Counter* is decremented by one after finishing the transmission of a TB, i.e. after transmitting the TB and all its possible retransmissions. When *Reselection Counter* is equal to zero, the UE must select new resources with probability $(1-P)$. Otherwise, the UE continues using the same SL resources with

³⁵ This is in contrast to LTE V2X mode 4 that only allows one blind retransmission (Section II).

³⁶ The first selected candidate resource is not necessarily the candidate resource for the initial transmission of the TB. It could be used for a retransmission of the TB.

³⁷ This does not apply for the selection of new SL resources with the re-evaluation or the pre-emption mechanism that are explained below.

the same *RRI* for a number of consecutive *Reselection Counter* TBs³⁸. Each UE can set up P between 0 and 0.8. Before transmitting the last TB that sets the *Reselection Counter* to zero, the TX UE evaluates (i.e. with probability $1-P$) whether it must select new resources for the next TB. If this is the case, the UE sets the *RRI* equal to 0 ms in the 1st-stage SCI to indicate other UEs that it is not reserving the same resources for the next TB. If not, the UE keeps the same resources for the next TB and the same *RRI* included in the 1st-stage SCI.

Even if the *Reselection Counter* is higher than zero, new resources must be selected if the new TB to be transmitted does not fit in the resources previously selected, or if those resources cannot fulfil the latency requirement for transmitting the new TB. A TX UE implementing the semi-persistent scheme utilizes the *RRI* included in the “resource reservation period” field of each transmitted 1st-stage SCI (Section 8.1.4 in [42]) to inform nearby UEs about the resources it is reserving for the initial transmission or a retransmission of the next TB. The 1st-stage SCI can indicate up to N_{SCI} resources, so the TX UE can indicate with each transmitted 1st-stage SCI the resources reserved for the following $N_{SCR}-1$ retransmissions of the next TB (Fig. 19).

1) Re-evaluation mechanism

A UE that has selected one new SL resource (e.g. at the slot m) will continue sensing the transmissions from other UEs during the selection window. A UE can decide whether to execute again step 1 to check if the selected resource is still available. The decision to execute again step 1 (and how often it is executed) is left to UE implementation. The UE can execute again step 1 at slot $m-T_3$ or before. T_3 is the maximum time allowed for a UE (in slots) to complete the resource selection process and is equal to $T_{proc,1}$. A UE can execute again step 1 after slot $m-T_3$ only if its processing capabilities are sufficient to complete step 1 and step 2 before m . The process to execute again step 1 works as follows. We denote n' as the slot at which a UE initiates a new execution of step 1. The UE defines a new selection window SW' that starts at slot $n'+T_1$ and ends at slot $n'+T_2'$. T_2' must be within the range $T_{2min} \leq T_2' \leq PDB-(n'-n)$ ³⁹. The UE executes then step 1 over the candidate resources in SW' in order to evaluate the currently available and excluded resources. If the selected resource at slot m is now excluded, then the UE has detected what is called in 3GPP standards a re-evaluation (Section 8.1.4 in [42]). Re-evaluation applies to both dynamic and semi-persistent schemes. This re-evaluation triggers the execution of step 2 to select a new SL resource among the currently available resources in SW' (Section 5.22.1 in [74]) (Fig. 20). The UE does not execute step 2 if the initially selected resource remains available. We should note that the UE could have selected N candidate resources in the initial selection window SW (for the transmission of the TB and its retransmissions as described previously). If this is the case, the UE can decide again whether to execute again step 1 to check if the selected resources are still available. If the UE detects that a subset of M ($M \leq N$) of the selected resources are not available anymore (i.e. they have been excluded after executing again

step 1), then the UE executes step 2 to select M new candidate resources within the selection window SW' . It is important to highlight that the re-evaluation is an important novelty introduced in NR V2X SL mode 2 compared to LTE V2X mode 4. It provides higher flexibility in the management of resources and better capacity to cope with the variability in transmitted messages and previously undetected interference problems between vehicles.

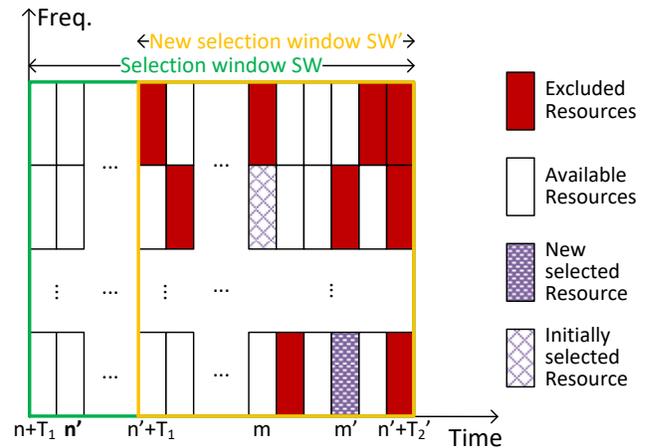


Fig. 20. New step 2 execution after a new step 1 execution and re-evaluation detection in NR V2X SL mode 2.

2) Pre-emption mechanism

Mode 2 introduces a pre-emption mechanism to prioritize traffic⁴⁰ [45] that is an important novelty compared to mode 4. It offers a more flexible management of resources under the presence of traffic with different priorities. With pre-emption, a UE with low priority traffic must free its reserved resource if it estimates that another UE with higher priority will use the reserved resource⁴¹. When a priority threshold is (pre-)configured in the resource pool, then the UE only frees its reserved resource if the priority of the other UE is higher than the priority threshold. The UE frees its reserved resource (e.g. at the slot r) to avoid a collision with the estimated transmission of the other UE that has higher priority. Note that pre-emption applies to both the dynamic and semi-persistent schemes. Note also that the UE can only identify that other UEs will use its reserved resources if it executes again the step 1. The decision to execute the step 1 (and how often it is executed) is left to UE implementation. With pre-emption, the step 1 is only executed after a TB is generated and it is waiting for the reserved resource to be transmitted. This is the case because the UE can only know the priority of its TB once this is generated. The UE can execute the step 1 at slot $r-T_3$ or before. It could execute the step 1 after slot $r-T_3$ only if its processing capabilities are sufficient to complete step 1 and step 2 before r . Let us consider that the UE initiates the step 1 at the slot n'' . The UE defines a new selection window SW'' that starts at slot $n''+T_1$ and ends at slot $n''+T_2''$. T_2'' must be within the range $T_{2min} \leq T_2'' \leq PDB-(n''-n_G)$. n_G is the slot at which the current TB has been generated and PDB is the deadline for the transmission of the TB. If the reserved resource at slot r is now excluded with step 1, then the

³⁸ The *Reselection Counter* is also randomly chosen in this case.

³⁹ $(n'-n)$ represents the time elapsed (in number of slots) from the initial resource (re-)selection trigger at slot n .

⁴⁰ The priority of a transmission is included in the 1st-stage SCI.

⁴¹ Note that pre-emption can be enabled or disabled in the resource pool.

UE checks the pre-emption condition (Section 8.1.4 in [42]), i.e. if the other UE that wants to utilize the reserved resource has higher priority or not (and higher than the priority threshold if it is (pre-)configured). If it has higher priority, the UE must free the reserved resource at slot r and execute again step 2 to select a new SL resource among the currently available resources in SW'' (Section 5.22.1 in [74]). We should note that the UE could have initially reserved N_R resources in SW''. If this is the case, the UE checks whether the N_R resources are available or not after executing step 1 again. If the UE detects that a subset of M_R ($M_R \leq N_R$) resources have been excluded after step 1 and they fulfil the pre-emption condition, then the UE frees the M_R resources and selects new M_R resources within the selection window SW'' following step 2.

C. Congestion Control

Rel. 16 supports congestion control for NR V2X SL communications in mode 2⁴². Like in Rel. 14 [23], Rel. 16 does not specify a particular congestion control algorithm but defines related metrics and possible countermeasures to reduce the channel congestion. These metrics are the Channel Busy Ratio (CBR) and the Channel occupancy Ratio (CR) that are referred to in Rel. 16 as SL CBR and SL CR.

In LTE V2X, the CBR is defined as the ratio of occupied sub-channels within the previous 100 subframes (Section 5.1.30 in [24]). A sub-channel is considered occupied if the measured RSSI exceeds a (pre-)configured threshold⁴³. In NR V2X, the size of the time window for estimating the SL CBR is equal to 100 slots (for any μ) or $100 \cdot 2^\mu$ slots by (pre-)configuration per resource pool [76], where μ is the SCS configuration factor. Rel. 16 also adapts the calculation of the RSSI to the granularity of NR V2X in the time and frequency domains (Section 5.1.25 in [76]).

The CR estimates the channel occupancy generated by a TX UE. In Rel. 14 LTE V2X, the CR is computed as the ratio of sub-channels that the TX UE utilizes or selects within a period or window of 1000 subframes or 1 second (Section 5.1.31 in [24]). The CR at subframe n is estimated over subframes $[n-a, n+b]$. It is up to the UE to decide the values of a and b subject to $b \geq 0$, $a+b+1=1000$, and $a \geq 500$. Rel. 16 adapts the computation of SL CR so that the $[n-a, n+b]$ range is defined in slots rather than subframes. Rel. 16 also adapts the size of the window to compute the SL CR (i.e., $a+b+1$) that is set equal to 1000 slots (for any μ) or $1000 \cdot 2^\mu$ slots by (pre-)configuration per resource pool [76]. a and b are determined by UE implementation considering that a must be positive, $b < (a+b+1)/2$ and $n+b$ shall not exceed the last selected resource (Section 5.1.26 in [76]). When the UE evaluates the SL CR in slots $[n+1, n+b]$, it shall assume that the transmission parameters used at slot n are also utilized in the scheduled (re-)transmissions within this range. The UE shall also assume that packet dropping will not occur in the range $[n+1, n+b]$.

A TX UE uses the measured SL CBR and SL CR to identify whether it has to modify its transmission parameters to reduce the channel load. This is done using a (pre-)configured lookup table that includes up to 16 CBR ranges. Each range is linked with a maximum SL CR (CR_{limit}) that the TX UE cannot surpass⁴⁴. CR_{limit} increases as the CBR range decreases. 3GPP establishes that the value of the CR_{limit} for each CBR range should be a function of the priority of the TB and the absolute speed of the TX UE [63]. The standard specifies that the TX UE evaluates whether it is exceeding the CR_{limit} and has to modify its transmission parameters at each (re-)transmission. To this aim, the TX UE computes the SL CR and SL CBR at slot $n-N_{proc}$ for a scheduled (re-)transmission at slot n ⁴⁵. N_{proc} depends on the UE processing capability. The standard differentiates UEs with fast processing capability or processing capability 1, and UEs with slow processing capability or processing capability 2 (Section 8.1.6 in [42]). When $\mu=0$, N_{proc} is equal to 2 slots for any UE processing capability. For any other μ , N_{proc} is equal to 2^μ slots for a UE with processing capability 1 and equal to $2 \cdot 2^\mu$ slots for a UE with processing capability 2. A TX UE has to modify its transmission parameters if the SL CR exceeds the CR_{limit} . The objective is to reduce the SL CR and control the channel load generated by a TX UE. A TX UE can modify the following transmission parameters per resource pool:

- 1) MCS: the TX UE can reduce the channel load using a higher order MCS that reduces the number of sub-channels necessary to transmit a TB.
- 2) Number of sub-channels: the UE can reduce its CR by limiting $\max(L_{PSSCH})$ or the number of sub-channels it can utilize. If a UE needs to fit a TB to a reduced number of sub-channels, it can utilize, e.g., a higher order MCS.
- 3) Number of (re-)transmissions: the UE can reduce its CR by limiting N_{MAX} or the number of (re-)transmissions.
- 4) Transmission power: the UE can decrease the CBR by reducing its transmission power. If the CBR decreases to values within lower CBR ranges, the UE can utilize a higher CR_{limit} .

VII. QoS FRAMEWORK FOR 5G NR V2X

V2X communications over the LTE-Uu and LTE-PC5 interfaces have QoS support. For LTE-Uu based V2X communications, the LTE QoS model that is based on QoS class identifiers (QCIs) is utilized. The QCI is a scalar that references specific QoS characteristics (e.g., packet loss rate) to be provided to a service data flow. This LTE QoS model was extended to introduce standardized QCI values specific for unicast and multicast/broadcast V2X messages (Section 4.4.5.2 in [22]). V2X SL communications over the LTE-PC5 interface leverage the QoS handling specified by ProSe. In particular, LTE V2X manages the QoS of SL communications on a per-

⁴²The gNB controls the channel load under mode 1. To this aim, it can request each UE to report (periodically or on demand) its measured CBR [62].

⁴³A range of values can be found in Rel. 14 [79]. Rel. 16 also specifies a range of values for this threshold that are defined as $(-112 + n \cdot 2)$ dBm, where n is an integer in the range $0 \leq n \leq 45$.

⁴⁴For LTE V2X, this table is defined in Europe by ETSI in Table 1 of Section 4.4 in [80]. This table has not yet been defined for NR V2X. However, a new Work Item (NR-V2X access layer - REN/ITS-00446) is defined in ETSI to extend the LTE V2X access layer specifications to include NR V2X.

⁴⁵This is in contrast to LTE V2X that uses the CBR and CR values computed at subframe $n-4$ for each scheduled (re-)transmission at subframe n .

packet basis. V2X packets generated at the Application Layer are associated to priority (PPPP, Proximity Service Per-Packet Priority) and, optionally, reliability (PPPR, Proximity Service Per-Packet Reliability) values [81], [34]. It is up to UE-implementation how to map at the Access Stratum (AS) Layer (i.e., PHY and Layer 2) the PPPP and PPPR values to the Sidelink Radio Bearer (SLRB) configuration as well as how to configure the corresponding SL logical channels. This means there are no unified rules for managing QoS among UEs in LTE V2X SL, and different UEs may handle the prioritization of logical channels differently.

NR V2X is designed to support diverse and stringent QoS requirements of eV2X services. These requirements are defined in terms of priority, transmission rate, latency, reliability, data rate, and communication range [51]. Rel. 16 identified that using per-packet QoS management for NR V2X SL based on PPPP and PPPR only might not be sufficient to account for all QoS requirements of advanced eV2X services. Consequently, 3GPP has adopted for NR V2X SL communications the more advanced 5G QoS model that is defined for NR Uu (Section 5.7 in [35], Section 5.4 in [51], [82]). The 5G QoS model is described in Section VII.A. QoS management in NR V2X SL is configured by the network and is not dependent on UE implementation. It is based on *QoS Flows*. QoS Flows are associated with the QoS requirements of the eV2X applications following QoS Profiles (Section VII.B). The QoS Profiles are specified using QoS parameters and QoS characteristics [35]. Section VII.C describes the PC5-RRC sublayer that is introduced in NR V2X for QoS profiling and management of unicast NR V2X SL communications. For NR V2X communications over the Uu interface, Rel. 16 introduces advanced mechanisms to support service continuity in case of QoS changes (Section VII.D) and report analytics to monitor expected changes of the application QoS (Section VII.E).

A. QoS management for NR V2X sidelink

Rel. 16 NR V2X QoS model is based on the 5G QoS model. The model is based on 5G QoS flows that are defined in [35] as “*the finest granularity for QoS forwarding treatment in the 5G System. All traffic mapped to the same 5G QoS Flow receive the same forwarding treatment (e.g., scheduling policy, queue management policy, rate shaping policy, RLC configuration, etc.). Providing different QoS forwarding treatment requires separate 5G QoS Flow.*”. The 5G QoS model for NR V2X SL is illustrated in Fig. 21. First, the V2X application packets are mapped to QoS flows at the V2X layer taking into account the applications’ QoS requirements (i.e., QoS parameters and QoS characteristics that are defined in Section VII.B). The mapping is done using the PC5 QoS rules following the mapping configuration provided by the 5G Core Network (Section VII.B). Each QoS flow is identified with a PC5 QoS Flow ID (PFI). The QoS flows are then mapped to sidelink radio bearers⁴⁶ using rules provided to the UE by the network via NR Uu RRC signaling. The mapping from QoS flows to radio bearers is performed at the SDAP (Service Data Adaptation

Protocol) layer. SDAP was first introduced in Rel. 15 and it is used in Rel.16 to support the QoS management of NR V2X SL communications. SDAP is also responsible for marking the transmitted packets with the corresponding PFI. Finally, the SLRB is established with the peer node (e.g., another UE for SL unicast communications) and configured. Configuring an SLRB implies configuring all parameters of the AS Layer (i.e., Packet Data Convergence Protocol (PDCP), Radio Link Control (RLC), MAC, and PHY layer). It should be noted that the 5G QoS model applies when UEs are in network coverage and in RRC connected state. Rel. 16 NR V2X extends the 5G QoS model to also support out-of-coverage operation and QoS provisioning when the UE is in RRC the idle state ([51], Section 7 in [62]). In these cases, the UE could use (pre-)configured mapping rules or SLRB configurations provided by the network. Rel. 16 adapts the application of the 5G QoS model to NR V2X SL based on the cast type (unicast, broadcast and groupcast). Information on the specific SLRB configuration process (based on the 5G QoS model) for NR V2X SL unicast, broadcast, and groupcast communications is provided below.

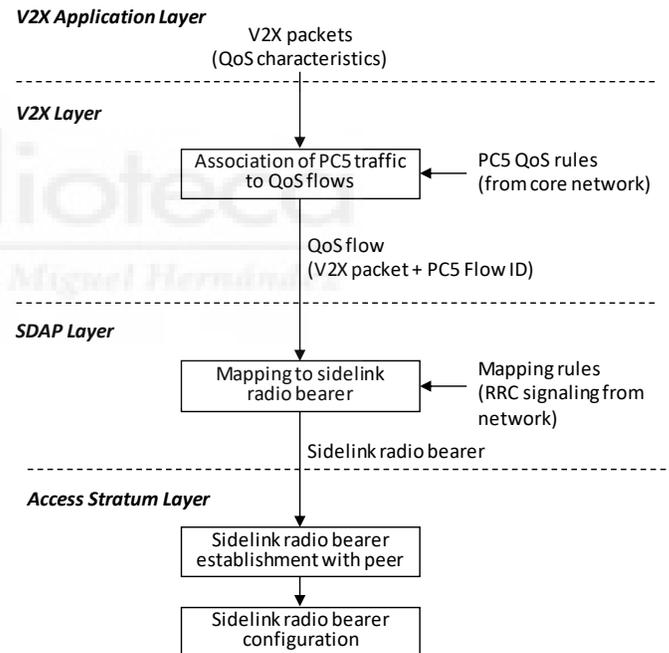


Fig. 21. Per-flow 5G QoS Model.

1) SL unicast

For SL unicast, two procedures are defined for configuring the SLRB (Section 6.3 in [51]). The first one is based on the 5G QoS model described in Fig. 21, and is referred to as a UE-specific SLRB configuration. The process starts with the provisioning to the UE of the PC5 QoS rules as part of the service authorization process. The PC5 QoS rules are used at the UE to map the application layer eV2X packets to PFI(s). The UE uses the derived PFI(s) to request from the network a configuration for the SLRB(s). The network uses the provided PFI(s) to derive the PC5 QoS profile(s) (i.e., a set of specific

⁴⁶A radio bearer is defined in [83] as “*the service provided by the Layer 2 for transfer of user data between UE and radio access network.*” The radio

bearer specifies the configuration of Layer 2 and PHY layer to meet the QoS requirements.

PC5 QoS parameters) and notifies the configuration of the SLRB to the UE. This UE-to-network request/response process is executed via NR Uu RRC signaling. With the configuration received from the network, the UE establishes and configures the SLRB with the peer UE via specific PC5-RRC signaling (see Section VII.C), and maps the PC5 QoS flow(s) to the configured SLRB. The mapping from PC5 QoS flow to SLRB occurs in the SDAP layer of the UE. The SL unicast communication between the two UEs is then performed as per the SLRB configuration. In the second procedure, the UE leverages NR SL (pre-)configured mapping rules (either signaled by the network or pre-configured in the UE) to configure the SLRB. This second procedure is used when the UE is out-of-coverage. The second procedure is referred to as pre-configured based SLRB configuration. This procedure maps (at the V2X layer) the eV2X application packets to PFI(s) based on (pre-)configured PC5 QoS rules. The NR SL (pre-)configuration is also used at the UE's AS layer to configure the SLRB based on the derived PFI(s).

2) SL groupcast and broadcast

Rel. 16 has defined three different procedures for the configuration of the SLRB for groupcast and broadcast communications (Section 6.3 in [51]). SL groupcast and broadcast transmissions are managed on a per-packet basis but with the advanced QoS profiles presented in Section VII.B. The V2X layer identifies the adequate PC5 QoS profile for each V2X packet to be transmitted. In the first procedure (referred to as UE-specific), the UE requests the network via RRC signaling for a specific SLRB configuration for the V2X packets, which are characterized by a PC5 QoS profile (e.g., range). The received SLRB configuration is used to establish the SLRB(s) and map the packets to the established SLRB(s). The second method (referred to as cell-specific) is based on a cell-specific configuration of the SLRB. In this case, the network uses V2X specific System Information Block (SIB) messages to broadcast the SLRB configuration associated with each possible PC5 QoS profile. The packets to be transmitted are then mapped to the SLRB configuration included in the SIB. Finally, NR SL can also use (pre-)configured based SLRB for SL groupcast and broadcast communications. In this procedure, the SLRB configuration associated with all PC5 QoS profiles is (pre-)configured either via signaling messages from the network or pre-configured in the UE. The UE uses this (pre-)configuration to map the V2X packets to be transmitted to the most adequate SLRB.

B. PC5 QoS Profiles for V2X sidelink

As described in Section VII.A, PC5 QoS Rules are used to classify and mark the eV2X application packets of SL user plane traffic, i.e., associate them to PC5 QoS Flows. A PC5 QoS Rule consists of the PFI of the associated PC5 QoS Flow, a precedence value, and a PC5 Packet Filter Set. The precedence

⁴⁷Intelligent Transportation Systems Application Identifier (ITS-AID) is a globally unique number used to identify an ITS application (e.g., 36 corresponds to CAM), as specified by ISO [84]. ITS-AID is named as PSID (Provider Service Identifier) in the IEEE WAVE specifications.

TABLE IX
STANDARDIZED PQI TO QoS CHARACTERISTICS MAPPING (BASED ON TABLE 5.4.4-1 IN [51])

PQI Value	Resource Type	Default priority Level	Packet Delay Budget	Packet Error Rate	Example V2X Services
21	GBR	3	20 ms	10 ⁻⁴	Platooning between UEs – Higher level of automation (LoA); Platooning between UE and RSU – Higher LoA
22		4	50 ms	10 ⁻²	Sensor sharing – Higher LoA
23		3	100 ms	10 ⁻⁴	Information sharing for automated driving – between UEs or UE and RSU - Higher LoA
55	Non-GBR	3	10 ms	10 ⁻⁴	Cooperative lane change – Higher LoA
56		6	20 ms	10 ⁻¹	Platooning informative exchange – low LoA; Platooning – information sharing with RSU
57		5	25 ms	10 ⁻¹	Cooperative lane change – lower LoA
58		4	100 ms	10 ⁻²	Sensor information sharing – lower LoA
59		6	500 ms	10 ⁻¹	Platooning – reporting to an RSU
90	Delay Critical GBR	3	10 ms	10 ⁻⁴	Cooperative collision avoidance; Sensor sharing – Higher LoA; Video sharing – Higher LoA
91		2	3 ms	10 ⁻⁵	Emergency trajectory alignment; Sensor sharing – Higher LoA

value determines the order in which the PC5 QoS Rules are evaluated when classifying the traffic. The PC5 Packet Filter Set contains the information on the V2X service type (e.g., the ITS-AID⁴⁷ or PSID), Source/Destination Layer-2 (L2) ID⁴⁸, and Application Layer ID (e.g., Station ID).

A UE checks for any service request or packet coming from the V2X application layer if there is any existing PC5 QoS Flow matching the request, based on the PC5 QoS Rules for the existing PC5 QoS Flows. If not, the UE derives the PC5 QoS parameters taking into account the V2X application requirements and service type according to the mapping configuration. The mapping configuration could be pre-configured at the UE or provided/updated by the V2X application server or the core network. The UE then creates a new PC5 QoS Flow and assigns it a PFI [51].

QoS Flows are characterized by QoS Parameters and QoS Characteristics that, overall, are referred to as QoS Profiles. For each cast type (unicast, broadcast, and groupcast), the UE maintains the mappings of PFIs to the PC5 QoS parameters and PC5 QoS Rules per destination (identified by Destination L2 ID). The assigned PFI and the associated PC5 QoS parameters are provided from the V2X layer to the AS layer. For unicast communications, the peer UEs negotiate the PFI and PC5 QoS parameters as described in Section VII.C.

Rel. 16 defines the following PC5 QoS Parameters associated to the PC5 QoS Flows (Section 5.4.2 in [51]):

- 1) *PC5 5G NR Standardized QoS Identifier (PQI)*. PQI is an identifier (a scalar value) that refers to specific PC5 QoS characteristics associated with V2X services. It is

⁴⁸Source and Destination Layer-2 IDs identify the sender and the target of the NR SL communication, respectively. One of three types of transmission modes (unicast, groupcast, and broadcast) is supported for a pair of a Source L2 ID and a Destination L2 ID.

introduced in Rel. 16 as a special 5G NR Standardized QoS Identifier (5QI).

- 2) *PC5 Flow Bit Rates*. This parameter is specified for QoS Flows with guaranteed bit rate (GBR) only. It consists of Guaranteed Flow Bit Rate (GFBR) and Maximum Flow Bit Rate (MFBR) parameters that are used to control the bitrate on the PC5 link. GFBR denotes the lowest bitrate that the V2X service can tolerate. MFBR sets the maximum bitrate expected by the QoS Flow. The excess traffic may get discarded or delayed by a rate shaping or policing function at the UE, RAN or CN. The measurements of GFBR and MFBR are done over an Averaging Time Window that is specified as a part of the associated QoS characteristics.
- 3) *PC5 Link Aggregated Bit Rates*. This parameter is defined for non-GBR links, and includes per link Aggregate Maximum Bit Rate (PC5 LINK-AMBR). PC5 LINK-AMBR limits the aggregate bit rate (measured over an AMBR averaging window) for all non-GBR QoS flows over a single PC5 unicast link with a peer UE.
- 4) *Range*. It is defined for groupcast communications over PC5. It indicates a minimum range (in meters) between the TX UE and the RX UEs for which the rest of PC5 QoS parameters must be guaranteed. For RX UEs beyond the indicated range, the communication is best effort⁴⁹.

The PQI identifies the following PC5 QoS characteristics that define the QoS profile and control how packets are managed from the QoS perspective:

- 1) *Resource Type* GBR, Delay-critical GBR or Non-GBR GBR flows require the dedicated allocation of network resources. Delay-critical GBR is introduced in 5G, and has specific definitions of Packet Delay Budget and Packet Error Rate that differ from the ones defined for GBR flows. Delay-critical GBR is also characterized by Maximum Data Burst Volume (MDBV).
- 2) *Priority Level*. It is similar to the priority value of LTE PC5 (PPPP, [34]), but improves backward compatibility for NR PC5. It is utilized to prioritize PC5 services (with the lowest value corresponding to the highest priority) if QoS requirements cannot be fulfilled for all of them.
- 3) *Packet Delay Budget (PDB)*. It indicates the upper bound delay. For Delay-Critical GBR flows, a packet delayed more than PDB is counted as lost (if the data burst does not exceed the MDBV within PDB, and the flow is not exceeding the GFBR). For GBR QoS flows, 98% of the packets shall not experience a delay over the PDB (if the flow is not exceeding GFBR). Services using Non-GBR QoS flows can experience packet drops and delays during congestions. Yet, in uncongested scenarios, 98% of the packets should not experience a delay over PDB.
- 4) *Packet Error Rate (PER)*. It is the upper bound for the rate of packets processed by the RLC layer of the TX UE but not successfully delivered by the RX UE to its PDCP layer. PER is not related to congestion losses but to radio errors. For Delay-critical GBR QoS Flows, packets delayed more

than PDB are counted as lost and included in the PER estimation (unless the GFBR or MDBV is exceeded).

- 5) *Averaging window* (for GBR and Delay-critical GBR resource type only). It specifies the duration over which the GFBR and MFBR are calculated.
- 6) *Maximum Data Burst Volume* (for Delay-critical GBR resource type only). MDBV denotes the largest amount of data that the PC5 link is required to serve for QoS Flows with Delay-critical GBR within a period of PDB.

Table IX provides the mapping between standardized PQI values and PC5 QoS characteristics following Table 5.4.4-1 in [51]. The mapping is done for PC5 QoS characteristics that correspond to QoS requirements of V2X services provided as examples. PQI values of 21-23 refer to GBR flows mapped to services with higher degree of vehicular automation. Non-GBR flows are represented by PQI values of 55-59, and provide a wider range of options for the PDB and PER constraints. Delay Critical GBR flows are indicated by PQIs 90 and 91, and correspond to the most stringent PDB and PER requirements. Note that GBR and Delay Critical GBR PQIs can only be used for unicast PC5 communications. MDBV is specified for all Delay Critical GBR flows by default as 2000 bytes. The averaging window for GBR and Delay Critical GBR is set to 2000 ms. The default values of MDBV and averaging window could be overridden if indicated by the application. Similarly, a PQI could be used with a priority level (indicated by an application) that is different than its default value.

C. PC5-RRC

PC5-RRC is introduced in NR V2X to provide functionalities to support SL unicast communications [62]. Note that in LTE V2X, PC5-RRC has a single functionality, which is to support exchanging synchronization-related information between the UEs on the Sidelink Broadcast Control Channel (SBCCH) [19]. PC5-RRC comprises RRC protocol and signaling that runs over the AS PC5 control plane protocol stack composed of the RRC, PDCP, RLC, MAC and PHY sublayers [45]. The PC5-RRC functionalities in NR V2X mainly consist of exchanging AS-level information that needs to be aligned between TX UE and RX UE to support SL unicast communications. These include:

- SL UE capability information to indicate additional, i.e., non-mandatory, features and parameters that the UE supports for each AS radio protocol (only utilized for the unicast NR SL communication), such as the frequency band, SCS, and MCS.
- SLRB configuration information to establish, modify, and release SLRBs (cf. Section VII.B).
- SL measurement configuration provided by the TX UE to the RX UE. The RX UE should be able to report RSRP measurements to the TX UE (without necessarily notifying the gNB of the result). Based on the RSRP measurements, the TX UE could then adjust its SL transmission power. The reporting could be event-triggered or periodical (for more information, see

⁴⁹ For example, transmission reliability is improved in HARQ groupcast feedback option 1 with HARQ retransmissions based on NACK sent only by RX UEs within the required communication range.

Section V.C.7).

For the signaling of the above information between the peer UEs, different PC5-RRC procedures and messages are specified [87]. UEs can request each other's capability information via Capability Enquiry message along with sending their own, whereas, SL RRC Reconfiguration procedure is used for the exchange of SLRB and SL measurement configurations. Signaling is done over the Sidelink Signaling Radio Bearer (SRB) on the logical channel SCCH. In case a configuration fails, UEs can utilize explicit failure message and timer based indications to notify the failure to the peer UE.

A PC5-RRC connection is defined as a logical connection between a pair of Source and Destination L2 IDs for unicast NR SL communication [87]. A UE may have multiple PC5-RRC connections, i.e., unicast connections, with one or more UEs for different pairs of Source and Destination L2 IDs. Signaling for establishing a PC5-RRC connection is initiated after a corresponding unicast link is established. The PC5-RRC connection and the corresponding SLRBs are released when the unicast link is released. Establishment and release of the unicast link is specified in [51].

PC5-RRC also supports the detection of SL Radio Link Failures (RLF) over the unicast NR SL communications. This is important to determine whether or when to release an SL unicast connection, for example, due to the degradation of the link as UEs move away from each other. The RLC triggers the SL RLF declaration when the maximum number of retransmissions to a specific destination has been reached. Upon the declaration, the UE releases the PC5-RRC connection immediately and discards any associated SL UE context.

D. Alternative QoS Profiles and Service Requirements for V2X communication over Uu

Rel. 16 also introduces advanced mechanisms to support service continuity in NR V2X communications over Uu. Service continuity is an important requirement of critical V2X services. Rel. 16 facilitates service continuity by supporting V2X applications with a range of different configurations and QoS characteristics (e.g., different bitrates or delay requirements). This is useful since the applications can continue to be operational even if the initial QoS profile is not available and only an alternative QoS profile (i.e., with a lower QoS) can be used. An alternative QoS profile is a combination of QoS parameters and QoS characteristics to which the application traffic is able to adapt and that has the same format as the QoS profile for that QoS Flow (as described in Section VII.B). Note that Section VII.B presents PC5 QoS Profiles, i.e., PQI. Equivalent profiles for 5G communications over the Uu are reported in Section 5.7.4 of [35] and are referred to as 5QI.

To support V2X applications that can operate with different configurations, the V2X AS provides to the 5GS the requested level of service requirements and the alternative service requirements. The alternative QoS profile contains one or more QoS reference parameters in a prioritized order for service operation. This enables the 5GS to use the alternative service profiles in case of the QoS changes. The NG-RAN notifies the 5GC that an alternative QoS profile can be supported, and then

the 5GC can provide this notification to the V2X AS.

The V2X AS includes the alternative service requirements as specified in Section 6.1.3.22 of [88] when providing service information to the PCF (i.e., through the NEF in case the V2X AS is outside of the 3GPP network domain, see Section IV). The V2X AS subscribes to receive notifications from the PCF when the QoS targets (the initial or the alternative ones) can no longer/again be fulfilled. The PCF enables QoS notification control and includes the derived alternative QoS parameter sets (in the same prioritized order indicated by the V2X AS) in the information sent to the SMF. If the NG-RAN receives from 5GC (through the SMF) a list of alternative QoS profile(s) for a QoS flow, it checks if a QoS profile of the list can be supported and sends a notification to the PCF. When the PCF notifies the V2X AS that the QoS targets can no longer/again be fulfilled, it includes the alternative QoS parameter set.

E. QoS Sustainability Analytics for V2X communication over Uu

The experienced QoS over the Uu interface may be affected by various factors (e.g., UE density, interference, mobility, handover, and roaming transitions). Rel. 16 introduces in the 5GS the mechanisms to monitor, collect and report information of the experienced QoS. In particular, the 5GS can notify the V2X application (upon request and through a V2X AS) of an expected or estimated change of QoS before it actually occurs [51]. This procedure is referred to as QoS Sustainability Analytics in 3GPP standards and helps the V2X application to decide in a proactive and safe manner if there is need for an application change (e.g., safely stop a service, adapt an application, etc.) when the QoS degrades. A V2X AS may request notifications on QoS sustainability analytics for an indicated geographic area and time period. The V2X AS then provides the network with location information in the form of a path of interest or a geographical area where to receive notifications of potential changes in QoS. The V2X AS also sends to the 5GS the QoS parameters that should be monitored as well as the corresponding QoS thresholds (e.g., minimum acceptable data rate, maximum acceptable packet error rate) for efficient and safe operation of an application. The 5GS compares the QoS thresholds with predicted values of QoS parameters to decide if it should notify the V2X AS of an expected change of QoS.

The V2X AS can either subscribe to notifications (i.e., a Subscribe-Notify model) or request a single notification (i.e., a Request-Response model) by the 5GS. Fig. 22 depicts the procedure to provide notifications on QoS Sustainability Analytics. This procedure is described in detail in Section 6.4.1 of [51] and Section 6.9 of [89] and consists of the next steps:

- 1) The V2X AS collects application layer information (e.g., V2X service, path, path start time and QoS requirements and thresholds).
- 2) The V2X AS requests or subscribes to analytics information on QoS Sustainability provided by the NWDAF through the NEF (see Section IV). NWDAF is responsible for on demand provision of analytics.
- 3) The NWDAF collects statistics provided by the

Operations, Administration and Maintenance (OAM) entity that is responsible for management plane functions including network performance monitoring.

- 4) The NWDAF computes the requested analytics or prediction about expected change of QoS for the requested area and time period. The NWDAF can detect the need to notify a potential QoS change by comparing the requested analytics of the target QoS profile against the threshold(s) provided by the V2X AS.
- 5) The NWDAF replies to the V2X AS (via the NEF) with the following information:
 - Applicable Area. A list of Tracking Area Identifiers (TAIs) or Cell IDs where the provided analytics applies.
 - Applicable Time Period. The time period during which the analytics applies.
 - Crossed Threshold(s). The QoS parameters with thresholds that are met or exceeded compared with the analytics or prediction computed by the NWDAF.
 - Confidence. Confidence of computed analytics or prediction.
- 6) The V2X application can take the necessary decisions based on the notification received from the network (these decisions are outside of 3GPP scope).

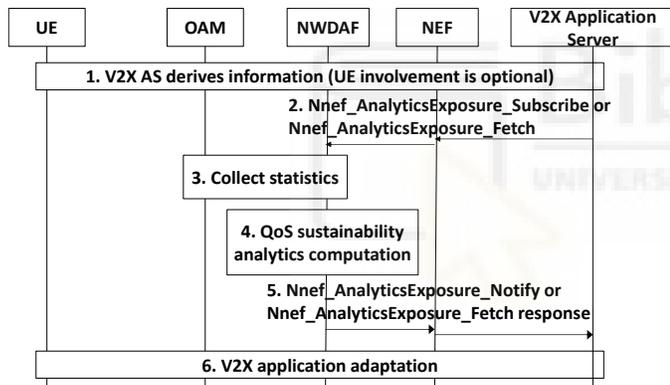


Fig. 22. Notification on QoS Sustainability Analytics to the V2X Application Server (based on 3GPP TS 23.287 and 3GPP TS 23.288).

VIII. 5G NR V2N ENHANCEMENTS

A. Uu enhancements

Rel. 16 NR V2X supports V2N communication over the Uu interface. Rel. 16 introduces enhancements to the Rel. 15 NR Uu and LTE Uu interfaces to meet the requirements of advanced eV2X services. It should be noted though, that neither Rel. 15 nor Rel. 16 NR Uu support broadcast and multicast V2N communication. If any eV2X service requires broadcast or multicast transmissions over the Uu interface, this can be supported using the MBSFN (Multimedia Broadcast Single Frequency Network) and SC-PTM (Single-Cell Point-To-Multipoint) technologies introduced in Release 13 (Rel. 13) [90]. This section presents the main enhancements to the NR Uu interface introduced for NR V2X in Rel. 16.

1) Multiple active UL configured grants

Uplink configured grants (UL CGs) are a set of periodic

resources utilized to schedule V2N transmissions semi-persistently. Similarly to CG in sidelink (Section VI.A), UL CGs avoid the signaling necessary to dynamically schedule each transmission. Uplink semi-persistent scheduling (UL SPS) was first introduced in Release 8 (Rel. 8) [91]. Rel. 15 LTE V2X supports an enhanced version of Rel. 8 UL SPS with up to eight UL SPS configurations. These configurations are established by the eNB using the UEs' reports on UL traffic characteristics (Section 5.2.2 in [15]). A Rel. 15 UL SPS configuration is defined, among others, by the semi-persistent scheduling periodicity (indicated in number of subframes or short transmission time intervals (TTIs)) and the number of permitted blind retransmissions (see [19] for a complete list of configuration parameters). Out of the eight UL SPS configurations, a UE can only have one configuration active at a time (Section 6.1.2 in [62]). The eNB dynamically selects the most suitable configuration based on the information sent on the UE's report (including the periodicity of TBs and their maximum size) [19]. The eNB activates and releases an UL SPS configuration using the DCI (Section 5.3.3.1.1 in [93]). The DCI also indicates the resources assigned to a UE and the MCS used for all its UL transmissions.

Rel. 15 LTE V2X UL SPS may not efficiently support eV2X services with varying traffic patterns and strict QoS requirements. Supporting eV2X services with UL SPS from Rel. 15 LTE V2X may result in an overprovisioning of resources to satisfy the most demanding QoS requirements. It can also result in frequent re-configurations of the UL CGs to adjust the resource allocation and MCS to the packets to be transmitted. In this context, Rel. 16 includes the support for multiple active UL CGs in order to efficiently and simultaneously support eV2X services with distinct requirements (e.g., latency, packet size, and reliability). These grants are also activated, configured, and released using the DCI. Fig. 23 shows an example in which a UE has two active UL CGs. Each grant is characterized by a time-frequency domain allocation, periodicity, and number of blind retransmissions of a TB (the number of blind retransmissions can be 2, 4, or 8). A complete characterization of each grant for Rel. 16 UL CG is available in [63], and it includes (among others) the MCS, TB size, DMRS configuration, and an indication of whether power control should be utilized or not (Section 7.1.1 in [66]). A UE with multiple active CGs can support more efficiently periodic and aperiodic transmissions with different requirements. For every UL TB transmission, the UE would select the active UL CG that most closely satisfies the service requirements. It can occur that if a UE has multiple active UL CGs, there are resources from different grants in the same slot (i.e., they coincide in time). In this case, the UE selects for this TB transmission the UL SPS configuration (from the grants that overlap on the same slot) that best fits the service requirements (Section 6.1.2 in [62]). It should be noted that a UE can only perform one transmission using any of the UL CGs at a time (i.e., at the same slot).

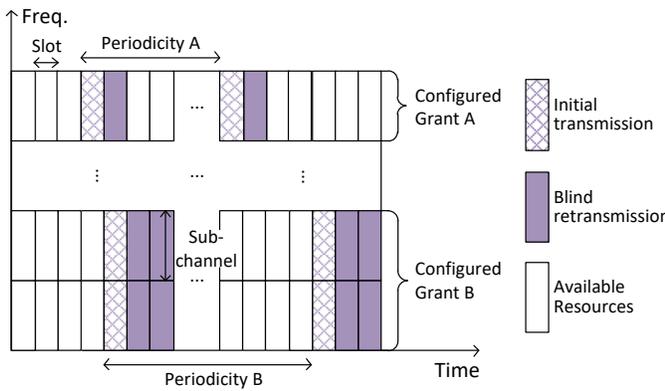


Fig. 23. Multiple active UL configured grants for NR V2N communications.

2) UE assistance information to the gNB

Rel. 16 allows UEs to report relevant assistance information to the gNB. This includes information about the Uu and SL traffic characteristics such as periodicity, packet latency requirements, and maximum TB size among others (see Section 5.7.4 in [63] for a complete list). The gNB can exploit this information to identify the UL CGs that best match the characteristics and requirements of the Uu traffic. The assistance information can also be utilized to improve V2N communications in scenarios where SL and Uu transmissions share the same radio resources. In this case, the gNB can exploit the information about the SL traffic characteristics for scheduling UL transmissions and to identify the adequate UL CGs that can also minimize the interference to the SL transmissions. The UE can also report (over the Uu interface) assistance information to the gNB that can be utilized by mode 1 and mode 2 SL scheduling [94]. These reports include information about the SL traffic (e.g., SL CBR of the SL resource pool) and UE-related geographic information (e.g., position, speed). The gNB could, for example, use UE's location to avoid assigning under mode 1 the same resources to UEs that are close to each other. On the other hand, the gNB could assign the same resources to UEs that are far enough apart in order to improve the efficiency of resource use. For mode 2 SL scheduling, [94] shows that geo-location information could be used, for example, to allocate the same resource pool to vehicles driving in the same direction on a highway.

B. Mobility Enhancements

Rel. 16 introduces mobility enhancements to manage handovers (HOs) of advanced eV2X applications supported over the Uu interface (e.g., remote driving). The HO mechanism in NR Rel. 15 is based on LTE (legacy HO). In a legacy HO, the source cell configures the moment at which the UE must execute a HO by sending a HO command. These commands are generally sent when the source cell detects that the UE receives a signal level from a potential target cell higher than from the source cell. The transmission of the HO command may fail (e.g., due to low signal level). If this happens, the UE will not perform the HO to the target cell and may suffer a RLF with the source cell and consequently a service interruption. Even if the HO command is correctly received, the legacy HO includes a mobility interruption time during the time the UE

disconnects from the source cell and connects to the target cell [16]. During this time, a UE cannot exchange user plane packets with any base station. LTE includes solutions such as Make Before Break and RACH-less HO that reduce the mobility interruption time to ~ 5 ms [19]. However, this is still not sufficient for certain eV2X applications with strict QoS requirements such as remote driving. To overcome the limitations of legacy HOs, Rel. 16 introduces for NR the Dual Active Protocol Stack (DAPS) and conditional HOs (CHO). Rel. 16 supports intra-frequency and inter-frequency HOs from NR {FR1, FR2} to NR {FR1, FR2} [52]. For all possible options, [52] defines the handover delay (or time it takes to the UE to respond to the HO command) and maximum interruption time as a function of the utilized SCS.

With DAPS, UEs can have two active links during a HO, one with the source cell and the other one with the target cell [95]. The UE can simultaneously receive data from (or transmit data to) the two cells as depicted in Fig. 24. When the UE simultaneously transmits to both cells, it must share the available transmission power between both connections. In this case, the UE receives power control commands from the source and target cells [95]. The total transmission power configured via the power control commands by the two cells cannot be higher than the UE maximum transmission power. If it is, then the UE prioritizes the power control command from the target cell and may reduce the transmission power to the source cell. The capacity to transmit (and/or receive) simultaneously to two cells is UE implementation dependent. The UE must inform the network whether it can execute DAPS for its DL and/or UL transmissions. In case a UE supports DAPS HO for DL communication only, the UE must switch the UL connection from the source to the target cell after the Random Access (RA) procedure is completed with the target cell (Section 8 in [66]).

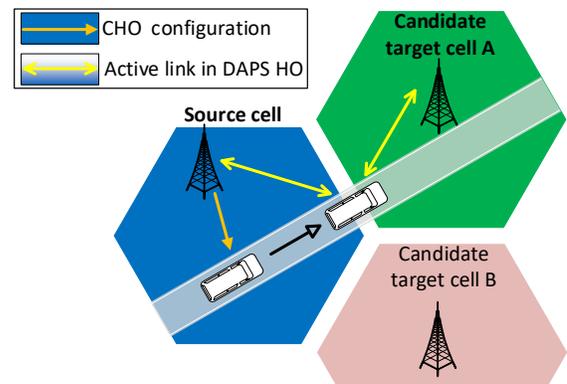


Fig. 24. Rel. 16 mobility enhancements: DAPS and CHO mechanisms.

A UE can receive ACKs or NACKs from the source and target cells during the DAPS HO. These acknowledgments may result in HARQ UL retransmissions by the UE. With DAPS, a UE can continue sending HARQ UL retransmissions to the source cell after starting the UL connection with the target cell [95]. Managing these retransmissions is not a challenge if the UE can maintain active links with both cells. However, conflicts can appear if the UE cannot simultaneously transmit to both cells. For example, a conflict occurs if both cells

schedule a UL transmission from the UE in resources that overlap in time. If such conflicts arise, the UE prioritizes the transmission to the target cell and discards the scheduled UL transmission to the source cell. The UE can perform UL HARQ retransmissions to the source cell until it receives an explicit indication from the target cell to stop the exchange of information with the source cell.

Rel. 16 also introduces the concept of CHO to improve the reliability of the handover process [95]. CHO avoids the risk of not receiving correctly the HO command. With CHO, the UE is notified of the conditions for executing a CHO before the link quality. This process is referred to as CHO configuration (Fig. 24) and is implemented through the exchange of RRC messages between the source cell and the UE. The UE constantly compares the HO measurements (i.e. RSRP, Reference Signal Received Quality (RSRQ) and Reference Signal-Signal to Noise and Interference Ratio (RS-SINR)) with the conditions on the CHO configuration to decide when to execute the HO [95]. A condition included in the CHO configuration is associated to a reference signal type (SSB or CSI-RS sent by a gNB), a trigger quantity⁵⁰ (RSRP, RSRQ or RS-SINR), and an event⁵¹. An example of an event (referred to as A3 event in the standard, Section 5.5.4 in [63]) is when the UE detects that the signal level of a target cell is higher than the source cell by a given offset. Each condition included in the CHO configuration is associated with a candidate target cell for the execution of the CHO. Multiple conditions can be associated with a candidate target cell as long as: 1) all the conditions associated with a candidate target cell must refer to a single reference signal type; and 2) all conditions use up to two triggering quantities (out of the three available ones). The UE continuously compares the conditions on the CHO with its HO measurements to decide when to execute the HO. When several candidate target cells simultaneously fulfill the conditions to execute a CHO, the selection of the candidate target cell with which to execute the CHO is left to UE implementation. Once a UE starts executing a CHO with a candidate target cell, it stops checking the conditions of other candidate target cells and ignores any new CHO configuration from the source cell [96]. It should be noted though that CHO is optional for both UEs and cells [96]. Rel. 16 still considers the use of legacy HO that is compatible with CHO. It is then possible that a source cell configures and sends a CHO configuration but then sends a HO command (i.e., legacy HO) to the UE. In this case, the legacy HO is prioritized

The execution of a CHO or a legacy HO may fail. If this happens, the UE has to select a cell following the process described within Section 5.2.3 in [97]. This process includes two procedures depending on the information (e.g., frequencies and cell parameters) stored in the UE about the neighboring cells. This information could be obtained, for example, as part of the CHO configuration for the candidate target cells. If the UE can leverage stored information, it selects a cell that fulfills

the received signal level and quality selection criterion defined in Section 5.2.3.2 of [97]. If the selected cell belongs to the set of CHO candidate target cells, the UE tries (again) to execute the CHO. If the CHO fails again, the UE does not try anymore to execute the CHO. Instead, it performs an RRC connection re-establishment procedure with the selected cell (Section 5.3.7 in [63]). The RRC connection re-establishment is also performed if the selected cell is not a CHO candidate target cell. This is also the case when the UE selects the cell without leveraging stored information about the surrounding cells. However, when the selected cell is not a CHO candidate target cell, the selected cell might not be aware of the current UE context (i.e., information needed to maintain the service). Retrieving the UE context increases the delay to resume the Uu data transmission (Section 5.3.7 in [63]).

In addition to DAPS and CHO, Rel. 16 also introduces a mechanism to improve the reliability of legacy HOs in NR V2X [95]. The mechanism is referred to as fast recovery handover failure. It can be applied when the UE is moving away from the source cell and does not receive correctly the HO command. The fast recovery handover failure includes three different timers: TTT, T310 and T312. First, the TTT timer is activated when the UE is approaching a target cell and the conditions specified in the A3 event are fulfilled. The T310 timer activates (after the activation of the TTT timer) when the UE detects a number of N310 out-of-sync signals⁵². The T312 timer activates when the TTT timer expires. The T312 timer is introduced in Rel. 16 for fast recovery handover failure, following the fast RLF recovery mechanism of LTE described in [98]. The introduction of this timer forces the execution of the RRC connection re-establishment procedure when the T312 expires. Otherwise, as the T310 timer is usually longer than the T312 timer, the RRC connection re-establishment procedure would be executed when the T310 timer expires. In case the T310 timer expires earlier than the T312 timer, the T312 timer is cancelled. It should be also noted that when the TTT timer expires, the UE sends to the source cell a measurement report, which may trigger a HO command from the source cell. If this HO command is received while the T312 timer is active, the UE performs a legacy HO to the target cell. In this case, the T310 and T312 timers are cancelled, since there is no need to execute the RRC connection re-establishment procedure.

IX. COEXISTENCE BETWEEN LTE V2X AND NR V2X SIDELINK

From the start of the work on NR V2X [9], 3GPP decided that NR V2X will complement LTE V2X and not replace it. LTE V2X is envisioned to support basic active safety applications, whereas NR V2X will support more advanced applications including connected and automated driving (Section III). Vehicles will select the adequate Radio Access Technology (RAT) based on the active V2X application. Fig.

⁵⁰Trigger quantity is the term used in the CHO mechanism to refer to the metric used to trigger the HO execution.

⁵¹An event defines the condition that should be fulfilled between the HO measurements obtained from the source and the target cell to trigger the HO execution in the CHO mechanism.

⁵²The UE detects an out-of-sync signal when the signal level of the source cell's PDCCH channel is below the threshold Q_{out} [99]. The N310 variable can be equal to {1, 2, 3, 4, 6, 8, 10, 20} and it is configured by the source cell (Section 6.3.2 in [63]).

25 illustrates the coexistence of NR V2X and LTE V2X. Vehicles exchange data and/or control information with other vehicles through the LTE or NR PC5 interfaces and with the infrastructure using the LTE or NR Uu interface. Referring to Fig. 25, vehicles A, B, and C can establish multi-RAT and multi-link V2V communications and still be able to communicate with vehicle D. In this context, Rel. 16 has defined mechanisms to facilitate the coexistence of NR V2X and LTE V2X at the vehicle/device level (referred to as in-device coexistence) in order to manage SL resources [45]. Rel. 16 also considers the coexistence and cooperation between the two RATs at the network level [45]. This coexistence is referred to as cross-RAT control.

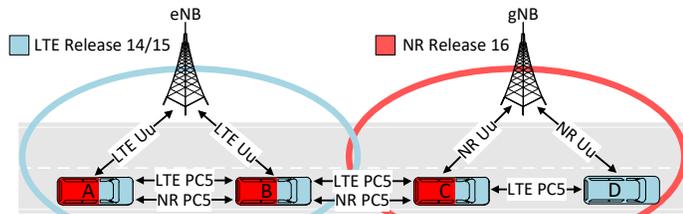


Fig. 25. Coexistence between NR V2X and LTE V2X.

A. In-Device Coexistence

The integration of two RATs within a vehicle creates a number of in-device coexistence challenges, especially if the two RATs share a part of the radio chain needed to perform transmission/reception or if the two RATs need to share the same channel/carrier/resources. Co-channel coexistence is defined in 3GPP as the concurrent usage of the same time/frequency resources by the two RATs (NR V2X and LTE V2X). Co-channel coexistence is not supported by Rel. 16 V2X since each RAT will utilize a different resource pool [100]. This decision was made at the start of the work item [10] due to the complexity that co-channel coexistence entails. However, the configuration and selection of resources in one RAT affects the other RAT for two main reasons [44]. The first reason relates to the transmission power. A vehicle has a maximum transmission power and the available power can be utilized by one RAT only during a certain amount of time (coordination is necessary between the two RATs to alternate) or it can be shared if both RATs transmit simultaneously. The first challenge is then related to the simultaneous transmission (Tx/Tx) over both RATs. The second reason relates to the potential interference at a given vehicle caused by the two RATs. This interference can appear if the LTE V2X and NR V2X resource pools are not sufficiently separated in frequency and the two RATs are simultaneously utilized. If such separation cannot be guaranteed, a vehicle will not be able to correctly receive data through RAT₁ if it is transmitting at the same time through RAT₂. The second challenge therefore occurs when the transmission of one RAT temporarily overlaps with the reception of the other RAT (Tx/Rx). Simultaneous reception (Rx/Rx) might be possible depending on the implemented solution to address the in-device coexistence challenges [68].

⁵³The co-existence with LTE V2X requires the time granularity to be at the subframe level and not the slot level.

To address these challenges, Rel. 16 proposes TDM (Time Division Multiplexing) and FDM (Frequency Division Multiplexing) solutions [62]. The solutions are applicable independently of whether NR V2X and LTE V2X operate using mode 1 and mode 3, or mode 2 and mode 4, respectively [44].

1) TDM Solutions

TDM solutions prevent vehicles from simultaneously transmitting over both RATs, thus requiring the synchronization of the two RATs [62]. Rel. 16 defines long-term and short-term TDM solutions for the coordination of the two RATs [62]. Both options are illustrated in Fig. 26.

The long-term TDM solution assigns statically the time during which each RAT can utilize its own resources (Fig. 26(a)). This allocation is pre-configured or determined by the gNB or eNB. Fig. 26(a) shows an example where LTE V2X is allowed to operate (Tx or Rx) during subframes SF_1 and SF_3 while NR V2X is allowed to operate (Tx or Rx) during subframes SF_2 and SF_4 . A vehicle can only transmit or receive in the same SF since the radios operate in half-duplex. The other RAT cannot transmit or receive during this time. This solution is simple to design, implement, and operate, but can negatively impact the QoS experienced over each RAT (e.g., the latency), since a RAT can only utilize the resources within its pool during the assigned subframes.

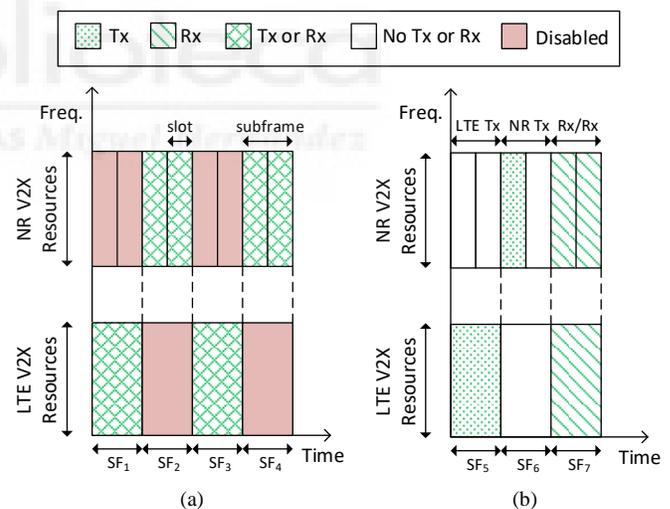


Fig. 26. TDM solutions for the coordination between RATs. (a) Long-term. (b) Short-term.

With the short-term TDM solution, each RAT can transmit or receive data in any slot or subframe of the assigned resource pool. LTE V2X and NR V2X need to dynamically coordinate the usage of their radio resources to avoid inter-RAT interference. This dynamic coordination enables a more flexible and granular usage of the resources that can better suit the requirements of applications served by each RAT. For the coordination, RATs exchange information about the resources they intend to use and the priority associated to this usage. In particular, RATs notify each other about [62]: 1) all subframes (SFs)⁵³ required for their planned transmissions (i.e., their

reserved SFs); and 2) all SFs in which they expect to receive a transmission (i.e., all SFs with detected reservations by other vehicles); and 3) the priority of these transmissions and expected receptions (if available). With this information, the RATs can detect if they were both planning to be active (Tx or Rx) at the same time. If the RAT with higher priority is planning to use its resources for a transmission, then the other RAT cannot transmit or receive data at the same time. If the RAT with higher priority is planning to use its resources for receiving data, then the other RAT can also receive data at the same time but cannot transmit [45]. The possible scenarios for the short-term TDM solution are illustrated in Fig. 26(b). In subframe SF_5 , a transmission of LTE V2X is prioritized (LTE Tx) and NR V2X cannot transmit or receive in any of the slots within SF_5 . In subframe SF_6 , a transmission of NR V2X is prioritized during the first slot of the subframe (NR Tx), but LTE V2X cannot transmit or receive during the complete SF_6 . Subframe SF_7 represents the case in which a reception (for any of the two RATs) is prioritized. In this case, the other RAT can also receive (Rx/Rx) but not transmit.

The short-term TDM solution leaves certain decisions for UE implementation. For example, let's consider the scenario where a UE has a potential conflict between the two RATs because they both want to transmit at time t_n . The conflict can be avoided using the short-term TDM coordination mechanism if the two RATs exchange their coordination information before $t_n - T$. The value of T is up to UE implementation subject to an upper bound of 4 ms⁵⁴. If the coordination information is not exchanged before the deadline, the approach to resolve the conflict between the two RATs is left to UE implementation [56]. For example, the vehicle could always prioritize LTE V2X transmissions as LTE V2X supports basic safety applications [44]. It is also left up to UE implementation the resolution of the conflict when both RATs want to transmit at the same time and have the same priority [44]. Conflicts between RATs can be frequent if the network load is high. These conflicts can be particularly delicate if a RAT is generally prioritized (e.g., due to the criticality of the data it transmits) and prevents the other RAT from being able to transmit, thus implying a QoS degradation. To prevent this from happening, 3GPP considers that the short-term TDM solution should only be utilized when the load of both RATs is below an acceptable level to avoid high performance degradation to any of the two RATs [62].

A vehicle is not required to support the short-term TDM solution. In this case, the vehicle will implement the long-term solution if it operates under TDM. The gNB or eNB configures the time during which each RAT is active (and can utilize its resource pool) if the vehicle is under cellular coverage. To this aim, the vehicle must report to the gNB or eNB whether it can execute or not the short-term TDM solution [44].

It should be highlighted that the TDM solutions help combat the power and interference challenges of in-device coexistence. Indeed, the TDM solutions (long-term and short-term) allow only one RAT to transmit at a given time. Consequently, the

transmitting RAT can operate at maximum transmission power. For the same reason, the TDM solutions prevent the interference between RATs. The two RATs can be active simultaneously only when they are in reception mode.

2) FDM Solutions

A vehicle can simultaneously transmit over NR V2X and LTE V2X if it implements FDM solutions for in-device coexistence [62]. In this case, the maximum transmission power of a vehicle or UE must be shared between the RATs. Similar to TDM, the sharing can be static or dynamic. In the latter case, the power is dynamically shared between RATs based on the priority of the data to be transmitted by each RAT [62]. An advantage of static power sharing is that it does not require any synchronization between RATs. Synchronization and coordination are necessary with dynamic power sharing to ensure each RAT adequately utilizes the corresponding transmission power at each point in time [62]. FDM solutions can also be inter-band or intra-band, depending on whether NR V2X and LTE V2X operate on the same or different frequency bands [62]. The frequency separation between the resource pools of LTE V2X and NR V2X (whether inter- or intra-band) can have a significant impact on the interference between RATs and hence on the operation and performance of FDM solutions. Sufficient spectral separation is needed not only to be able to transmit using both RATs, but also to be able to simultaneously transmit over one RAT and receive over the other RAT. Fig. 27 illustrates inter-band and intra-band FDM solutions with power sharing and inter/intra-band deployments. Static and dynamic power sharing as well as intra-band and inter-band operations were discussed during the standardization process. However, Rel. 16 specifications only include static power allocation in inter-band deployments [100]. This is the case because the inter-band FDM solution with sufficient spectral separation between the RATs is the only viable solution for avoiding interference and allowing simultaneously transmitting over one RAT and receiving over the other RAT (Tx/Rx) ([101]).

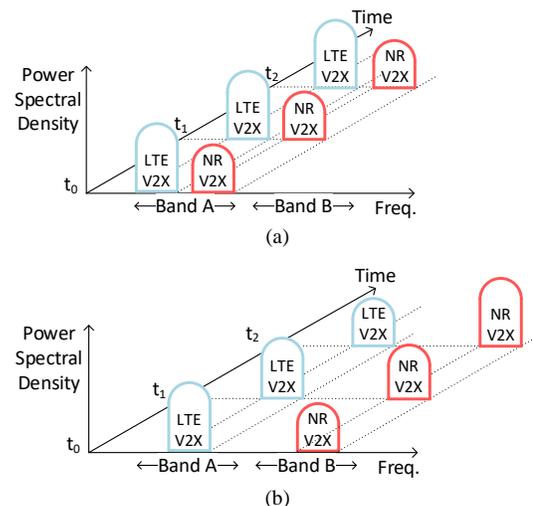


Fig. 27. FDM solutions. (a) Intra-band operation and static power sharing. (b) Inter-band operation and dynamic power sharing.

⁵⁴ The 4 ms upper bound is selected based on the maximum T_1 value required to prepare an LTE SL transmission (see Section II).

B. Support for Cross-RAT control

Rel. 16 has standardized cross-RAT control so that eNBs and gNBs can cooperate to manage the radio resources of NR V2X and LTE V2X SL communications when vehicles are under cellular coverage. In particular, the eNB can manage the resources associated with the LTE and NR PC5 interface through the LTE Uu interface. Furthermore, the gNB can manage the resources associated with both the LTE and NR PC5 interfaces through the NR Uu interface [45]. This can reduce the need to deploy gNBs to support NR V2X SL communications. With cross-RAT control, the gNB or eNB can also configure the CBR measurement and UE assistance information reporting on LTE V2X or NR V2X sidelink, respectively (see Section VIII) [63]. However, the cross-RAT control has some restrictions that depend on the type of signaling utilized and the resource allocation modes [102].

Base stations (eNBs and gNBs) send RRC signaling messages to the vehicles for managing the cross-RAT resources. Base stations can use two types of RRC signaling for cross-RAT control: dedicated RRC signaling and V2X-specific SIB [102]. With dedicated RRC signaling, a base station of RAT₁ can configure the resource pool of RAT₂ (cross-resource pool configuration) and also allocate its radio resources (cross-resource allocation). Dedicated RRC signaling can only be utilized if the vehicle has established an RRC connection with the serving base station, and it is in RRC connected state. Establishing a RRC connection is not necessary when utilizing V2X-specific SIB signalling. However, the base station can only execute the cross resource pool configuration (and not cross-resource allocation) when operating with V2X-specific SIB signalling.

Rel. 16 introduces changes to the RRC signalling so that vehicles can correctly decode the information transmitted by eNBs and gNBs [102]. In particular, Rel. 16 defines new RRC containers within the RRC messages of dedicated RRC signalling to be able to execute cross resource allocation and cross resource pool configuration. The standard introduces a container in NR RRC [63] that is compatible with LTE RRC [19] so that a gNB can control and manage the LTE V2X resources. The standard also updates the LTE RRC and includes a new container compatible with NR RRC so that eNBs can control and manage the NR V2X resources [102]. Similarly, Rel. 16 introduces two new V2X-specific SIBs for cross-RAT control: a V2X-Specific SIB compatible with LTE RRC is defined for gNBs and a V2X-Specific SIB compatible with NR RRC is defined for eNBs [102].

Cross-RAT control can be used to manage the radio resources under NR V2X mode 1 and LTE V2X mode 3, and also under NR V2X mode 2 and LTE V2X mode 4. For mode 2 and mode 4, a base station can (but does not need to) configure the resource pools of both RATs (cross-resource pool configuration). This is the case because vehicles autonomously select their radio resources and hence do not require cross-resource allocation from the base stations. For mode 1 and mode 3, a base station of RAT₁ can configure the resource pool of RAT₂ and also allocate its radio resources. However, it should be noted that an eNB can only allocate resources of NR

V2X mode 1 when operating with CG type 1 scheduling (and not with dynamic scheduling or CG type 2) [102].

X. EVALUATION METHODOLOGY

Rel. 16 NR V2X introduces new communication modes (unicast, groupcast) compared to Rel. 14/15 LTE V2X, as well as more stringent use case requirements (increased reliability and reduced latency) and new data traffic patterns requiring higher message and data rates (Section III). These differences required 3GPP to define an evaluation methodology for LTE V2X and NR V2X [103] that has been utilized during the standardization work and that should guide the community when evaluating the performance of NR V2X. The evaluation methodology includes: i) new channel models, in particular for V2V SL; and ii) assumptions for system and link level simulations.

A. Channel models

The SI that defined the evaluation methodology placed particular attention to the V2V SL channel models. For scenarios of interest (urban and highway), the goal of the study was to define propagation conditions, path loss and shadow fading, large scale parameters, and small scale parameters for V2V SL channels. Specifically, the approach to generate the V2V SL channels in Rel. 16 is based on the 12-step procedure defined in the Rel. 15 NR channel modeling study [104] and depicted in Fig. 28. However, Rel. 16 adapts several important components to the peculiarities of V2V channels. For example, the Rel. 16 V2X SL channel models introduce the NLOS_v state in addition to the LOS and NLOS states used for SL and UL/DL. The NLOS_v state describes V2V SL channels where the LOS path is blocked by another vehicle. There were several reasons leading to the introduction of NLOS_v: i) measurements reported in [105] and [106] showed that blockage of SL LOS by vehicles exhibits considerably different properties compared to NLOS channels since NLOS assumes blockage by objects considerably larger than vehicles (e.g., buildings); ii) in highways, there are no other objects that block the LOS other than vehicles; and iii) dynamics of NLOS_v blockage (e.g., in terms of temporal and spatial variation of blockage) is different to NLOS blockage.

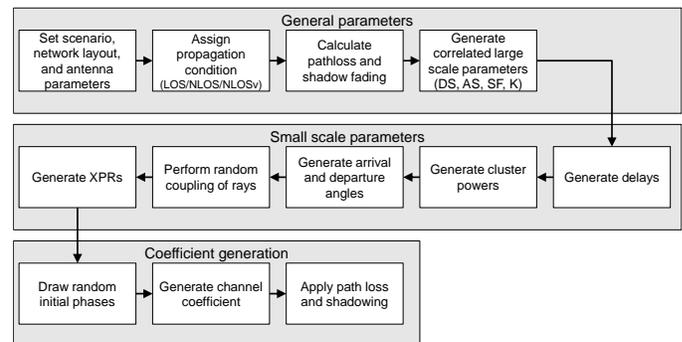


Fig. 28. Channel generation procedure for V2V SL (adapted from [104]).

Other important novelties introduced by the NR V2V SL channel models defined in Rel. 16 include:

- 1) New equations defining probability of LOS, NLOS, and

TABLE XI
PATHLOSS FOR V2V SL LINKS [103]

LOS/NLOS/NLOS _v	Pathloss [dB]		Shadow fading std. deviation σ_{SF} [dB]
LOS, NLOS _v	Highway $PL = 32.4 + 20 \log_{10}(d) + 20 \log_{10}(f_c)$	Urban grid $PL = 38.77 + 16.7 \log_{10}(d) + 18.2 \log_{10}(f_c)$	$\sigma_{SF} = 3$
NLOS	$PL = 36.85 + 30 \log_{10}(d) + 18.9 \log_{10}(f_c)$		$\sigma_{SF} = 4$

NLOS_v state for urban grid and highway.

- 2) New shadow fading and fast fading parameters for NLOS_v.
- 3) New dual mobility (i.e., both transmitter and receiver are mobile) Doppler calculations that also include variable scatterer velocities.

1) LOS probability, pathloss model, and shadow fading

Pathloss and LOS probability equations have been defined separately for all three channel states (LOS, NLOS, NLOS_v). Formally, the three channel states are defined as follows [103]:

- 1) LOS: V2V link is in LOS state if the two vehicles are in the same street and the LOS path is not blocked by vehicles.
- 2) NLOS (LOS path blocked by buildings): V2V link is in NLOS state if the two vehicles are in different streets.
- 3) NLOS_v (LOS path blocked by vehicles): V2V link is in NLOS_v state if the two vehicles are in the same street and the LOS path is blocked by vehicles.

LOS probabilities were not previously defined for SL communications. In fact, the LOS probabilities included in [104] were only defined for UL/DL. This, along with introduction of a new SL channel state (NLOS_v), required new LOS probability equations that have been defined for the evaluation methodology of NR V2X. These equations have been defined based on an extensive study reported in [108] that analyzed LOS probabilities and state transitions using maps of real cities and highways. 3GPP adopted a probabilistic approach for calculation of LOS probabilities for V2V SL communications when vehicles are in the same street. For urban grid scenarios, 3GPP assumes that the LOS path is blocked by buildings when vehicles are not in the same street and want to establish a V2V SL link (i.e., they operate under NLOS state). Consequently, LOS probability equations are only required for LOS and NLOS_v states and are defined following Table X for V2V SL communications. In the table, d denotes the relative distance between TX UE and RX UE in meters.

Furthermore, a new set of pathloss models for V2V SL

TABLE X
PROBABILITY OF LOS AND NLOS_v STATES.

Highway	
LOS	If $d \leq 475$ m, $P(\text{LOS}) = \min\{1, ad^2 + bd + c\}$ where $a=2.1013 \cdot 10^{-6}$, $b=-0.002$ and $c=1.01093$ If $d > 475$ m, $P(\text{LOS}) = \max\{0, 0.54 - 0.001(d-475)\}$
NLOS _v	$P(\text{NLOS}_v) = 1 - P(\text{LOS})$
Urban	
LOS	$P(\text{LOS}) = \min\{1, 1.05 \cdot \exp(-0.0114d)\}$
NLOS _v	$P(\text{NLOS}_v) = 1 - P(\text{LOS})$

⁵⁵ P2P stands for Pedestrian-to-Pedestrian communications.

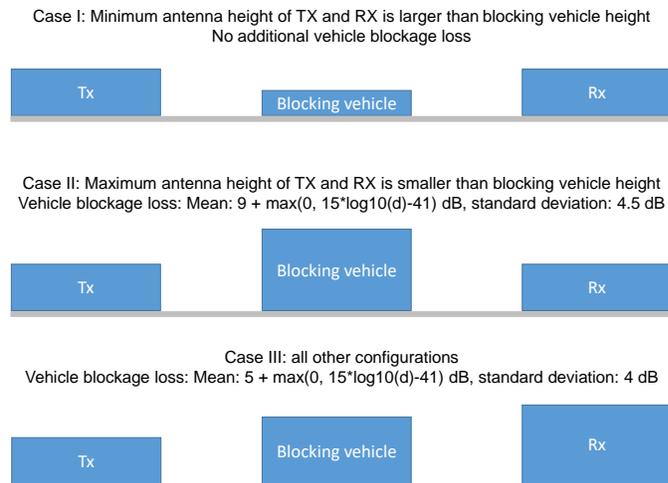


Fig. 29. Additional vehicle blockage loss under NLOS_v: scenarios and loss calculation. Vehicle blockage loss is modeled by the normal distribution for all cases.

communications in highway and urban grid have been introduced in Rel. 16, along with new values for the shadow fading standard deviation. The shadow fading is modeled with a random variable according to a lognormal distribution with zero mean. The shadow fading models for all types of SL (V2V, V2P, P2P⁵⁵, V2R, and R2R⁵⁶ links) are taken over from LTE V2X, as described in [15]. Effectively, for each V2V link, shadow fading is an independent and identically distributed lognormal random variable. The LOS shadowing model from [15] applies to NLOS_v as well. The model for spatial correlation of shadow fading defined for LTE V2X in [15] also applies to NR V2X.

Table XI contains the pathloss expressions for highway and urban grid scenarios where f_c denotes the carrier frequency in GHz and d denotes the Euclidean distance between a TX UE and a RX UE in 3D space in meters, i.e., considering also the heights of the Tx antenna and Rx antenna. Note that LOS and NLOS_v states use the same pathloss equation. However, additional *vehicle blockage loss* is introduced for NLOS_v following Fig. 29. The figure describes how much additional loss (if any) is added in case of NLOS_v for different relationships of the height of Tx antenna, Rx antenna and the blocking vehicle. The resulting pathloss plots are shown in Fig. 30 for all link types and in urban grid and highway scenarios.

2) Fast fading model

A completely new set of fast fading parameters were introduced for SL under Rel. 16. Each of the states (LOS, NLOS, NLOS_v) have a specific set of parameters associated with them, most of which are dependent on the assumed center

⁵⁶R2R stands for RSU-to-RSU communications.

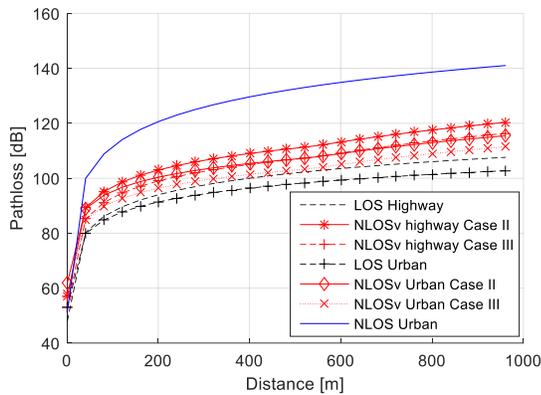


Fig. 30. V2V SL pathloss as a function of the distance between transmitter and receiver. The pathloss for NLOSv is plotted as an average value for each case.

frequency f_c . Table 6.2.3-1 in [103] lists all of the fast fading parameters.

3) Modeling dual mobility for V2V SL

In V2V SL, transmitter and receiver are highly mobile. To account for this condition, Rel. 16 NR V2X introduces for the first time a dual mobility model addressing different Doppler components. The model is an extension of the single mobility model from [104]. It takes into account the relative speed difference between a TX UE and RX UEs as well as the relative speed of delayed paths coming from scatterers. Specifically, the Doppler for the LOS path is expressed as follows [103]:

$$v_{n,m} = \frac{\hat{r}_{rx,n,m}^T \cdot \bar{v}_{rx} + \hat{r}_{tx,n,m}^T \cdot \bar{v}_{tx}}{\lambda_0} \quad (8)$$

$$\bar{v}_{rx} = v_{rx} [\sin \theta_{v,rx} \cos \phi_{v,rx}, \sin \theta_{v,rx} \sin \phi_{v,rx}, \cos \theta_{v,rx}]^T \quad (9)$$

$$\bar{v}_{tx} = v_{tx} [\sin \theta_{v,tx} \cos \phi_{v,tx}, \sin \theta_{v,tx} \sin \phi_{v,tx}, \cos \theta_{v,tx}]^T \quad (10)$$

whereas the Doppler for the delayed (scattered) paths is:

$$v_{n,m} = \frac{\hat{r}_{rx,n,m}^T \cdot \bar{v}_{rx} + \hat{r}_{tx,n,m}^T \cdot \bar{v}_{tx} + 2\alpha_{n,m} D_{n,m}}{\lambda_0} \quad (11)$$

where λ_0 denotes the wavelength of the carrier, θ is the elevation angle, Φ is the azimuth angle, $\hat{r}_{rx,n,m}$ is the spherical unit vector with azimuth arrival angle $\phi_{n,m,AOA}$ and elevation arrival angle $\theta_{n,m,ZOA}$, given by:

$$\hat{r}_{rx,n,m} = \begin{bmatrix} \sin \theta_{n,m,ZOA} \cos \phi_{n,m,AOA} \\ \sin \theta_{n,m,ZOA} \sin \phi_{n,m,AOA} \\ \cos \theta_{n,m,ZOA} \end{bmatrix}. \quad (12)$$

n denotes a cluster and m denotes a ray within cluster n . $\hat{r}_{tx,n,m}$ is the spherical unit vector with azimuth departure angle $\phi_{n,m,AOD}$ and elevation departure angle $\theta_{n,m,ZOD}$, given by:

$$\hat{r}_{tx,n,m} = \begin{bmatrix} \sin \theta_{n,m,ZOD} \cos \phi_{n,m,AOD} \\ \sin \theta_{n,m,ZOD} \sin \phi_{n,m,AOD} \\ \cos \theta_{n,m,ZOD} \end{bmatrix}. \quad (13)$$

Furthermore, $D_{n,m}$ is a random variable with uniform distribution from $-v_{scatt}$ to v_{scatt} m/s. v_{scatt} is the maximum speed (in m/s) of the vehicle in the simulation, and $\alpha_{n,m}$ ($0 \leq \alpha_{n,m} \leq 1$) is a random variable with uniform distribution. $D_{n,m}$ ensures that scatterers can have a range of speeds varying from minimum to maximum speed values, whereas $\alpha_{n,m}$ determines the proportion of scatterers that are mobile relative to the TX UE and RX UEs.

B. System level simulations

1) Evaluation scenarios, vehicle types, and dropping options

Similar to evaluation settings in Rel. 14 and Rel. 15 LTE V2X, Rel. 16 considers two evaluation environments for NR V2X: urban grid and highway, which are shown in Fig. 31 and Fig. 32, respectively. The requirement for urban grid is that the simulation is performed over at least 3x3 road grids in order to reduce the lower interference values at the edge where there are fewer communicating vehicles. The same problem in the highway scenario is tackled with the so-called wrap-around that reintroduces the simulated vehicles at one end of the highway at the other end of the simulated area. For each environment, the evaluation scenarios are divided into below and above 6

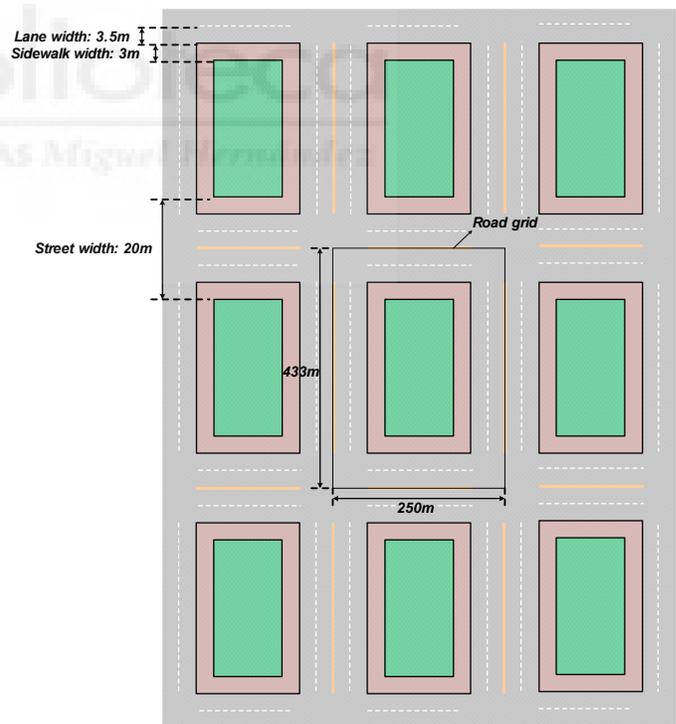


Fig. 31. Road configuration for urban grid (adapted from [103]).

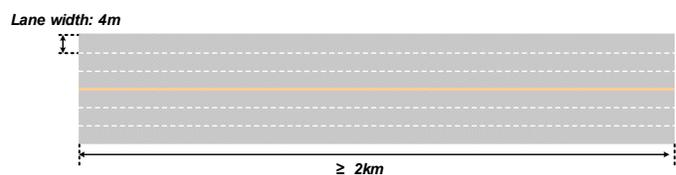


Fig. 32. Road configuration for highway scenario (adapted from [103]). Three lanes per direction. Wrap-around applied to simulation area.

TABLE XII
EVALUATION SCENARIOS AND PARAMETERS FOR NR V2X (ADAPTED FROM [103])

Parameters	Below 6 GHz		Above 6 GHz	
	Urban grid for eV2X	Highway for eV2X	Urban grid for eV2X	Urban grid for eV2X
Carrier frequency	Macro BS to/from vehicle/pedestrian UE: 4 GHz Between vehicle/pedestrian UE: 6 GHz Micro BS to/from vehicle/pedestrian UE: 4 GHz UE-type-RSU to/from vehicle/pedestrian UE: 6 GHz Note: Agreed value does not mean non-ITS band is precluded for real deployment for sidelink	Macro BS to/from vehicle/pedestrian UE: 2 GHz or 4GHz Between vehicle/pedestrian UE: 6 GHz Micro BS to/from vehicle/pedestrian UE: 4 GHz UE-type-RSU to/from vehicle/pedestrian UE: 6 GHz Note: Agreed value does not mean non-ITS band is precluded for real deployment for sidelink	Macro BS to/from vehicle/pedestrian UE: 30 GHz Between vehicle/pedestrian UE: 30 or 63 GHz Micro BS to/from vehicle/pedestrian UE: 30 GHz UE-type-RSU to/from vehicle/pedestrian UE: 30 or 63 GHz Note: Agreed value does not mean non-ITS band is precluded for real deployment for sidelink	Macro BS to/from vehicle/pedestrian UE: 30 GHz Between vehicle/pedestrian UE: 30 or 63 GHz Micro BS to/from vehicle/pedestrian UE: 30 GHz UE-type-RSU to/from vehicle/pedestrian UE: 30 or 63 GHz Note: Agreed value does not mean non-ITS band is precluded for real deployment for sidelink
Aggregated system bandwidth	Up to 200 MHz (DL+UL) Up to 100 MHz (SL)	Up to 200 MHz (DL+UL) Up to 100 MHz (SL)	Up to 1 GHz (DL+UL) Up to 1 GHz (SL)	Up to 1 GHz (DL+UL) Up to 1 GHz (SL)
Simulation bandwidth	20 or 40 MHz (DL+UL) 10 and 20 MHz (baseline for SL) 100 MHz (optional for SL)	20 or 40 MHz (DL+UL) 10 and 20 MHz (baseline for SL) 100 MHz (optional for SL)	200 MHz (DL+UL) 200 MHz (SL)	200 MHz (DL+UL) 200 MHz (SL)
BS Tx power	Macro BS: 49 dBm PA scaled down proportionally with simulation BW when system BW is higher than simulation BW. Otherwise, 49 dBm Micro BS: 24 dBm PA scaled down with simulation BW when system BW is higher than simulation BW. Otherwise, 24 dBm Note: 33 dBm for RSU is not precluded	Macro BS: 49 dBm PA scaled down proportionally with simulation BW when system BW is higher than simulation BW. Otherwise, 49 dBm Micro BS: 24 dBm PA scaled down with simulation BW when system BW is higher than simulation BW. Otherwise, 24 dBm Note: 33 dBm for RSU is not precluded	Macro BS: 43 dBm PA scaled down proportionally with simulation BW when system BW is higher than simulation BW. Otherwise, 43 dBm. EIRP should not exceed 78 dBm and is also subject to appropriate scaling	Macro BS: 43 dBm PA scaled down proportionally with simulation BW when system BW is higher than simulation BW. Otherwise, 43 dBm. EIRP should not exceed 78 dBm and is also subject to appropriate scaling
UE Tx power	Vehicle/pedestrian UE or UE-type-RSU: 23 dBm Note: 33 dBm is not precluded	Vehicle/pedestrian UE or UE type RSU: 23 dBm Note: 33 dBm is not precluded	Vehicle/pedestrian UE or UE type RSU: 23 dBm for 30 GHz, 21 dB baseline for 63 GHz, 27 dBm optional for 63 GHz. For both 30 and 63 GHz, EIRP should not exceed 43 dBm.	Vehicle/pedestrian UE or UE type RSU: 23 dBm for 30 GHz, 21 dB baseline for 63 GHz, 27 dBm optional for 63 GHz. For both 30 and 63 GHz, EIRP should not exceed 43 dBm.
BS receiver noise figure	5 dB	5 dB	7 dB	7 dB
UE receiver noise figure	9 dB		13 dB (baseline), 10 dB (optional)	

GHz with the related parameters shown in Table XII. Beyond the different carrier frequencies, a notable difference is the assumed total system bandwidth. This bandwidth is up to 200 MHz for UL+DL and 100 MHz for SL below 6 GHz. It is up to 1 GHz for UL+DL and 1 GHz for SL above 6 GHz. This indicates that the expectation of 3GPP is to have significantly larger bandwidths available in above 6 GHz (especially mmWave) frequencies.

3GPP guidelines for system level simulations distinguishes three types of vehicles:

- 1) Type 1 (passenger vehicle with lower antenna position, e.g., at the bumpers): length of 5 meters, width of 2.0 meters, height of 1.6 meters, antenna height of 0.75 meters.
- 2) Type 2 (passenger vehicle with higher antenna position, e.g., at the roof): length of 5 meters, width of 2.0 meters, height of 1.6 meters, antenna height of 1.6 meters.
- 3) Type 3 (truck/bus): length of 13 meters, width of 2.6 meters, height of 3 meters, antenna height of 3 meters (e.g., antenna at the roof).

The vehicles are dropped independently in each lane. They are dropped following an exponential distribution with mean equal to the average of the speed per lane multiplied by 2 seconds. Additionally, vehicles dropped in the same lane must always respect a minimum distance of 2 meters between the bumpers. Once dropped, vehicles maintain a fixed speed according to the speed assigned to the lane. Furthermore, to support the platooning use case, a clustered vehicle dropping is supported for Type 3 vehicles in the highway scenario. The cluster consists of vehicles in the same lane and with the same speed and a predefined distance between them. Further details on vehicle dropping options are shown in Table XIII.

2) Traffic Models

Traffic models have been defined in order to support as best as possible the diversity of the use case requirements defined in [12], while keeping the simulation complexity manageable. Three periodic and two aperiodic traffic models have been defined, as indicated in Table XIV.

TABLE XIII
VEHICLE DROPPING OPTIONS (ADAPTED FROM [103][103])

Vehicle Dropping	Urban grid	Highway
Option A	<ul style="list-style-type: none"> Vehicle type distribution: 100% vehicle type 2. Clustered dropping is not used. Vehicle speed is 60 km/h in all the lanes. In the intersection, a UE goes straight, turns left, turns right with the probability of 0.5, 0.25, 0.25, respectively. 	<ul style="list-style-type: none"> Vehicle type distribution: 100% vehicle type 2. Clustered dropping is not used. Vehicle speed is 140 km/h in all the lanes as baseline and 70 km/h in all the lanes optionally.
Option B	<ul style="list-style-type: none"> Vehicle type distribution: 20% vehicle type 1, 60% vehicle type 2, 20% vehicles type 3. Clustered dropping is not used. Vehicle speed in each lane is as follows: <ul style="list-style-type: none"> In the East-West/West-East direction: <ul style="list-style-type: none"> Speed in Lane 1: 60km/h (east-west left lane) Speed in Lane 2: 50km/h (east-west right lane) Speed in Lane 3: 25km/h (west-east left lane) Speed in Lane 4: 15km/h (west-east right lane) In the North-South/South-North direction: <ul style="list-style-type: none"> 0 km/h in all the lanes. No vehicles are dropped at the intersections in the North-South direction, so as not to block the traffic in the intersection. Vehicles do not change their direction at the intersection. 	<ul style="list-style-type: none"> Vehicle type distribution: 20% vehicle type 1, 60% vehicle type 2, 20% vehicle type 3. Clustered dropping is not used. Vehicle speed in each lane is as follows: <ul style="list-style-type: none"> Speed in Lane 1: 80km/h (east-west right lane) Speed in Lane 2: 100km/h (east-west center lane) Speed in Lane 3: 140km/h (east-west left lane) Speed in Lane 4: 40km/h (west-east left lane) Speed in Lane 5: 30km/h (west-east center lane) Speed in Lane 6: 20km/h (west-east right lane)
Option C	N/A	<ul style="list-style-type: none"> Vehicle type distribution: 0% vehicle type 1, 67% vehicle type 2, 33% vehicle type 3. Clustered dropping is used. Each cluster consists of 6 Type 3 vehicles with a gap of 2 meters. Vehicle speed is 140 km/h in all the lanes.

TABLE XIV
DATA TRAFFIC MODELS [103]

Model	Periodic			Aperiodic	
	Model 1 (low traffic intensity)	Model 2 (medium traffic intensity)	Model 3 (high traffic intensity)	Model 1 (medium traffic intensity)	Model 2 (high traffic intensity)
Inter-packet arrival time (ms)	100 ms	10 ms	30 ms	50 ms + an exponential random variable with the mean of 50 ms	10 ms + an exponential random variable with the mean of 10 ms
Packet size (bytes)	Pattern of {300 bytes, 190 bytes, 190 bytes, 190 bytes, 190 bytes} with random starting point for each UE	1200 bytes with probability of 0.2 and 800 bytes with probability of 0.8	Uniformly random in the range between 30000 bytes and 60000 bytes with the quantization step of 10000 bytes	Uniformly random in the range between 200 bytes and 2000 bytes with the quantization step of 200 bytes	Uniformly random in the range between 10000 bytes and 30000 bytes with the quantization step of 4000 bytes
Latency requirement (ms)	100 ms	10 ms	30 ms	50 ms	10 ms

3) Antenna settings

Antenna patterns for BSs have been largely carried over from [104], with heights set to 25 meters for macro BS and 5 meters for micro BS in urban grid. For highway, the height is set to 35 meters for macro BS with inter site distance (ISD) of 1732 meters, 25 meters height for ISD of 500 meters, and 5 meter height for micro BS. Antenna patterns for pedestrians are also largely carried over from [104] and the height is set to 1.5 meters. Antenna heights for vehicles are defined according to the vehicle type. Antenna patterns for vehicles are defined separately for roof and bumper locations for both rear and front, as well as for different frequencies. For full details on the patterns, we refer the reader to Section 6.1.4 of [103].

4) Performance metrics

Three performance metrics have been defined and utilized in 3GPP as shown in Table XV.

C. Link level simulations

Whenever applicable, the assumptions for system level simulation are used for link level simulation as well. The following list contains the most relevant link level simulation parameters [103]:

- Carrier frequency.
- Channel model (e.g., fast fading model).
- PHY packet size.
- Channel codes (for control and data channels).
- Modulation and code rates (for control and data channels).
- Signal waveform (for control and data channels).
- SCS.
- CP length.
- Frequency synchronization error.
- Time synchronization error.
- Channel estimation (e.g., DMRS pattern).
- Number of retransmission and combining (if applied).
- Number of antennas (at UE and BS).

TABLE XV
PERFORMANCE METRICS [103]

Packet reception ratio (PRR)	Packet Inter-Reception (PIR)	Throughput (defined in [109])
<p>PRR type 1: For one Tx packet, the PRR is calculated by X/Y, where Y is the number of UE/vehicles that are located in the range (a, b) from the TX, and X is the number of UE/vehicles with successful reception among Y. CDF of PRR and the following average PRR are used in evaluation.</p> <ul style="list-style-type: none"> CDF of PRR with $a = 0$, $b =$ baseline of 320 meters for highway and 150 meters for urban. Optionally, $b = 50$ meters for urban with 15 km/h vehicle speed. Average PRR, calculated as $(X1+X2+X3...+Xn)/(Y1+Y2+Y3...+Yn)$ where n denotes the number of generated messages in simulation. with $a = i*20$ meters, $b = (i+1)*20$ meters for $i=0, 1, \dots, 25$ 	<p>PIR type 1: For a given distance d, PIR is the time T_i elapsed between two successive successful receptions of two different packets transmitted from node A to node B for the same application, if the distance between node A and node B at the two packets' receiving time is within the range of $(0,D]$.</p> <ul style="list-style-type: none"> Average PIR within given distance d, calculated as $(T1+T2+T3+...+Tn)/n$ where n denotes the number of collected PIR in simulation. CDF of PIR with given distance D. 	<p>User throughput = amount of data (file size) / time needed to download data</p> <ul style="list-style-type: none"> Time needed to download data starts when the packet is received in the transmit buffer, and ends when the last bit of the packet is correctly delivered to the receiver
<p>PRR type 2: For one Tx packet, the PRR is calculated by S/Z, where Z is the number of UEs in the intended set of receivers, and S is the number of UE with successful reception among Z. Unicast is the special case where Z includes a single UE, where the PRR is the average of packets of the unicast link.</p>	<p>PIR type 2: PIR is the time T_i elapsed between two successive successful receptions of two different packets transmitted from node A to node B for the same application, if the node B is one of the intended set of receivers of the node A.</p> <ul style="list-style-type: none"> Average PIR with intended set of receivers, calculated as $(T1+T2+T3+...+Tn)/n$ where n denotes the number of collected PIR in simulation. CDF of PIR with intended set of receives. 	

14) Transmission diversity scheme (if applied).

15) UE receiver algorithm.

16) AGC settling time and guard period.

For all above parameters, [103] contains the suggested values or range of values. Furthermore, a set of Cluster Delay Line (CDL) models⁵⁷ has also been developed for the purpose of link level simulations. CDL models can be used for computationally efficient simulation of wireless channels, with the angular component allowing the incorporation of aspects particularly relevant for MIMO systems. The CDL models generated for Urban LOS, Urban NLOS, Urban NLOSv, Highway LOS, and Highway NLOSv channels are available in [103].

XI. OUTLOOK FOR FUTURE RELEASES

After completing Rel. 16, 3GPP has already identified new study and work items for Rel. 17. This section presents the most relevant study and work items for NR V2X SL communication as well as other possible enhancements that can benefit NR V2X SL communication, even if they are not included in the officially approved items.

A. Beamforming in Sidelink

In Rel. 16, NR V2X SL communication has been designed with a predominant focus on FR1. FR2 can be considered by using the design for FR1, but no specific optimization for FR2 (except for SL PT-RS) or beam management is supported under Rel. 16 [10]. In this context, enhancements to support high data rates (e.g., based on FR2) were initially considered in the preparation of the new work item on NR SL enhancement for Rel. 17. However, the agreed work item description for NR SL enhancement does not finally consider optimizations for FR2 [110]. Nonetheless, certain eV2X use cases such as vehicles

platooning, advanced driving and extended sensors require high data rates (50-700 Mbps) for long distances (200 meters or beyond) [12] (see Table II). These use cases could be supported with FR2, in particular with beamforming, to compensate the pathloss at higher frequencies. Beamforming is not only applicable to higher frequencies, but can also be supported for FR1. As beamforming enables directional transmission, this allows for a spatial reuse of available resources due to reduced interference. This can be leveraged for supporting a high connection density as required in certain V2X scenarios [12].

Although no specific optimization for FR2 is supported in Rel. 16 NR V2X, the S-SSB structure enables the transmission of synchronization information by a SyncRef UE in various directions using different beams (as discussed in Section V.B.5)) [111]. For a configured number of S-SSBs, the synchronization information can be sent by a SyncRef UE with the same beam multiple times or with different beams at each S-SSB, e.g., through beam sweeping. However, as the goal of transmitting S-SSBs is to expand the synchronization coverage, sending S-SSBs in all directions may not be efficient. For instance, S-SSB transmissions by a SyncRef UE in the direction of a gNB is not necessary as this area is already under the gNB's synchronization coverage as shown in Fig. 33(a). In fact, S-SSB transmissions towards the gNB may only lead to interference as a UE (which receives synchronization information from the gNB and the SyncRefUE) anyway selects the gNB as its synchronization reference, as it has higher priority as explained in Section V. To increase the synchronization coverage, a SyncRef UE should transmit S-SSBs away from the gNB as depicted in Fig. 33(b). This spatially selective transmission of synchronization information can be enabled with beamforming, i.e., by the SyncRef UE transmitting S-SSBs on certain beams.

⁵⁷Clustered Delay Line is a type of channel model where the received signal is composed of a number of separate delayed clusters. Each cluster contains a

number of multipath components with the same delay but different Angle of Departure and Angle of Arrival.

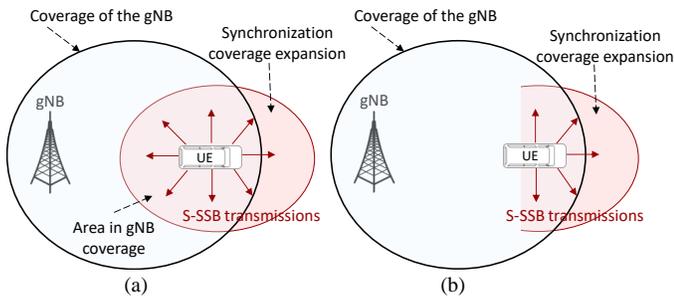


Fig. 33. Synchronization coverage expansion. (a) SyncRef UE transmits S-SSBs omnidirectionally. (b) SyncRef UE transmits S-SSBs away from the gNB.

This mechanism can also be beneficial for synchronization coverage expansion in out of coverage scenarios.

Further optimization includes extending the PSCCH/PSSCH power control to consider beam-based transmissions, similar to how the PUSCH and PUCCH power control support beamforming in NR Rel. 15 [66]. In NR Rel. 15, a UE determines the initial transmit power for sending PUSCH or PUCCH with a given beam based on the DL pathloss measured on that beam⁵⁸. With beam-based SL power control, the transmit power for a SL transmission on a given beam would depend on the SL pathloss and/or DL pathloss measured on that beam. This enables optimizing the PSCCH/PSSCH transmit power of a TX UE in the spatial domain. Consider for example a TX UE that is in network coverage with the PSCCH/PSSCH power control based on the DL pathloss (see Section V.C.7). Assuming in this case that the strongest signal (i.e., smallest DL pathloss) from the gNB is received by the TX UE on beam X (as seen in Fig. 34(b)), the transmit power of a PSCCH/PSSCH sent on another beam (i.e., beam Y) would not depend on the smallest DL pathloss which was measured on beam X . This is in contrast to having no beam-based power control as depicted in Fig. 34(a), where the transmit power of PSCCH/PSSCH (in any direction) is limited by the DL pathloss (see Fig. 15).

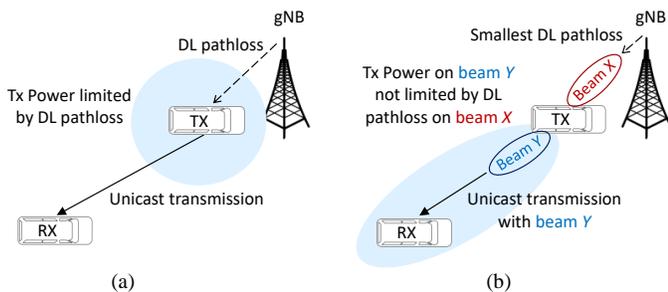


Fig. 34. Sidelink power control based on DL pathloss. (a) No beam-based power control/transmission. (b) Beam-based power control/transmission.

B. Relative positioning via sidelink measurements

A study item on scenarios and requirements of in coverage, partial coverage, and out of coverage positioning use cases has been approved for Rel. 17 [112]. The objectives of the study are to identify the positioning use cases and requirements for V2X

⁵⁸In contrast to NR V2X, the uplink power control in NR is closed loop, i.e., based on TPC commands from the gNB. However, prior to receiving any TPC commands, the uplink transmit power is based on the DL pathloss.

and public safety, as well as to identify potential deployment and operation scenarios for UEs in network coverage, partial network coverage, and out of network coverage. For instance, V2X use cases like platooning and coordinated maneuvers may impose requirements on relative positioning of nearby vehicles. In an overtaking maneuver, for example, a collision with a nearby vehicle can be avoided if the overtaking vehicle knows the relative position of the nearby vehicle as shown in Fig. 35. In such cases, relative positioning based on SL measurements (i.e., SL positioning or ranging) may be beneficial, as it can be supported for UEs independent of the network coverage.

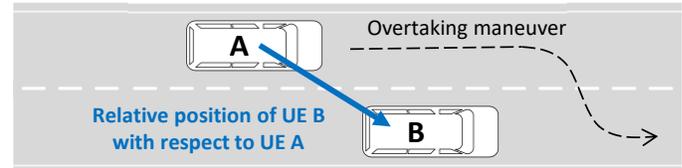


Fig. 35. Relative positioning by an overtaking vehicle (UE A) of a nearby vehicle (UE B).

Relative positioning can also be obtained with network-based positioning. However, network-based positioning requires UEs to be in network coverage and it may incur increased latency and signalling. In fact, the network first needs to obtain the absolute position of two UEs in order to derive their relative position, based on the difference between the two absolute positions. The relative position should then be shared with at least one of the UEs. In addition, deriving the relative position based on absolute positions can lead to imposing more stringent requirements for absolute positioning. This is the case as the uncertainties (i.e., positioning errors) in the absolute position of each of the two UEs add up when obtaining the relative positioning (e.g., based on the difference between the absolute positions of the two UEs shown in Fig. 36).

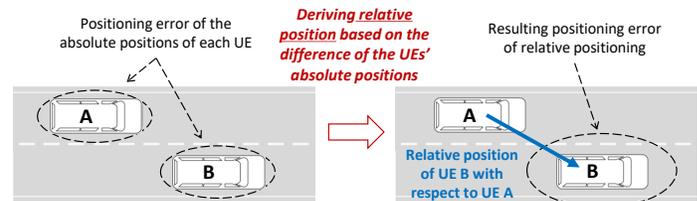


Fig. 36. Deriving relative positioning based on absolute positions of two UEs.

As some coordinated maneuvers in V2X only require avoiding a collision with a nearby vehicle, i.e., via relative positioning, the absolute position of a nearby vehicle is not really necessary. Relative positioning can be performed based on SL measurements, i.e., via SL positioning or ranging. Sidelink positioning can potentially benefit from the smaller SL pathloss between UEs (compared with the DL pathloss), the support for out of coverage, and the lower SL latency and signalling (compared with the Uu link) [113]. The SL measurements can be based on time and angular measurements

[114], similar to the measurements used for network-based positioning in Rel. 16. For V2X, the large dimensions of the vehicles can be exploited for placing distributed antennas or arrays for performing the SL measurements [114]. In addition, SL positioning can be used to complement or enhance any existing positioning schemes, e.g., based on network-based positioning, sensors, or GNSS, which may be limited due to obstructions or other factors. For this purpose, techniques for relative positioning based on SL measurements are of interest.

C. Enhancements to resource allocation

A new work item on NR Sidelink enhancement has been agreed for Rel. 17 [110]. One of the objectives of this Rel. 17 work item is to enhance Rel. 16 NR V2X resource allocation mode 2. Some planned enhancements focus on power saving and on improving KPIs such as reliability and latency. All enhancements must be able to coexist in the same resource pool (i.e., co-channel coexistence) with Rel. 16 NR V2X.

Mode 1 and mode 2 in Rel. 16 NR V2X have been designed for UEs such as vehicles or RSUs, which do not have strong power limitations. These limitations are present in other types of UEs (e.g., smartphones) that are used by VRUs such as pedestrians. With the current Rel. 16 mode 2 specifications, a UE requires long sensing intervals that severely impact the battery consumption. An enhancement adopted for study under Rel. 17 is a variant of mode 2 that would reduce power consumption by using partial sensing [110]. Partial sensing was already considered in Rel. 14 for a variant of LTE V2X mode 4, where a UE only senses a subset of subframes contained in the sensing interval [23], [115].

Another SL enhancement adopted for Rel. 17 is the support of inter-UE coordination [110]. Inter-UE coordination was analysed under Rel. 16 where it was referred to as mode 2(b) [116], [62], but it was not standardized. With inter-UE coordination, an UE can assist other UE(s) in their resource selection process. There are two types of inter-UE coordination: type A and type B [116]. Under type A, an *assisting* UE (UE₁) restricts the resources that can be used by an *assisted* UE (UE₂). Under type B, UE₁ sends recommendations on which resources should be selected by UE₂. UE₂ decides whether or not to follow these recommendations. Contents of recommendations or restrictions have not been decided by 3GPP yet, but the recommendations could contain, for example, the available resources detected by the assisting UE.

For Rel. 17, inter-UE coordination type B is considered [110]. An important benefit from inter-UE coordination is the reduction of the hidden terminal problem as illustrated in Fig. 37. Fig. 37 illustrates a scenario where UE B is under the transmission range of UE A and UE C. UE B can then receive transmissions and detect the resource reservation of UE A and UE C. However, this is not the case between UE A and UE C (or viceversa) that experience the hidden terminal problem. In this case, UE A and UE C can end up selecting the same resources for their transmissions. In Fig. 37, UE B has reserved the resource consisting of sub-channel SC_2 at slot S_2 . UE C has

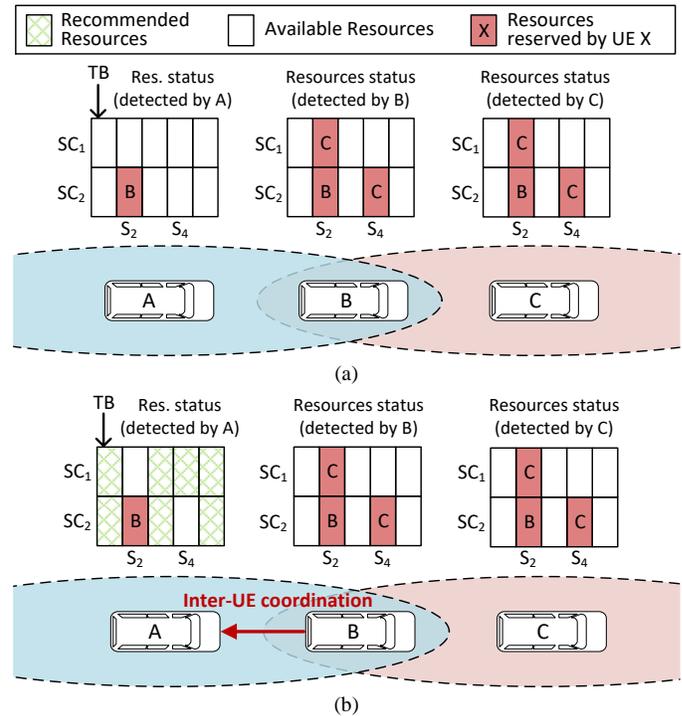


Fig. 37. Benefits of inter-UE coordination to mitigate the hidden terminal problem. (a) Without inter-UE coordination. (b) With inter-UE coordination.

reserved resources consisting of sub-channel SC_1 at slot S_2 and sub-channel SC_2 at slot S_4 . The resource reservation of UE B is detected by UE A and by UE C. However, reservations of UE C are only detected by UE B and not by UE A due to the hidden terminal problem. In slot S_1 , UE A generates a TB. Let's suppose that the selection window of UE A starts at S_2 and ends at S_5 . Since UE A has not detected the reservations of UE C, UE A would detect as available the resources reserved by UE C if there is no inter-UE coordination (Fig. 37(a)). In this case, UE A can select resources reserved by UE C with the consequent risk of packet collisions. Fig. 37(b) shows that inter-UE coordination reduces the hidden terminal problem. In this scenario, UE B acts as an assisting UE and sends a recommendation (type B) to UE A for its resource selection. As a result, the resource selection of UE A can consider the status of the resources detected by UE B (which include the reservations of UE C) and reduce the probability of collisions caused by the hidden terminal problem.

It is also interesting to highlight other enhancements that were analysed for NR V2X mode 2 during Rel. 16, but that were not adopted in Rel. 16 or Rel. 17. These enhancements could be of interest in future releases (Rel. 18 and beyond). One such enhancement is the possibility for a UE to schedule resources of other UEs [117]; it was referred to as mode 2(d) during the standardization process [62]. This enhancement requires establishing a group of member UEs, where a member UE acts as a scheduling UE⁵⁹ to schedule the transmissions of all the member UEs in the group. Mode 2(d) requires member UEs (or at least the scheduling UE) to operate in network coverage.

⁵⁹Different versions of the TR 38.885 refers to this UE as scheduling UE [118] or UE-A [62].

With mode 2(d), the scheduling UE acts as an intermediary between the gNB and other member UEs in the management of the SL radio resources. Different options were proposed to select the scheduling UE [118]. For example, the scheduling UE could be selected by the gNB or by the group members.

Another enhancement proposed is to perform the first transmission of a TB using NR V2X mode 2 and its retransmissions using NR V2X mode 1 in order to increase the reliability of NR V2X communications [120].

Another idea proposed for mode 2 is the concept of standalone reservation [121] where the selection window is divided into two windows W1 and W2. A UE that wants to transmit a TB selects one resource in each window, and announces in the resource reserved in W1 the resource it has selected in W2 to transmit the TB.

Another enhancement analysed (but finally not adopted) for mode 2 during Rel. 16 is the use of Time-Frequency Resource Patterns (TFRPs) [122]. Each TFRP contains a set of (pre-)configured resources in the resource pool. Each TFRP is adapted to meet specific QoS requirements, for example, by adjusting the number of resources in the TFRP, its periodicity and the number of resources reserved for retransmissions. TFRPs are repeated with certain periodicity. The concept of TFRP is associated with the resource allocation proposal referred to as mode 2(c) during the study item phase of Rel. 16 [62]. A TFRP can be assigned to a UE or it can be shared by a group of UEs. Several proposals were made for selecting the TFRP and avoid collisions between UEs. For example, a UE can select the TFRP using sensing information [123] or using information about the relative position of other UEs [122].

D. UE Relaying

3GPP has been developing the concept of UE relaying since Rel. 13. The initial scope focused mainly on public safety scenarios in order to extend the coverage range. First specifications referred to UE relaying as ProSe (Proximity Service) and UE-to-network relaying [81]. This was later extended in Rel. 14 and 15 for use cases where power-limited devices (e.g., wearables) use other UEs (e.g., smartphones) as relays to connect to the network [124]. Rel. 17 includes a work item on UE relaying with the objective to achieve maximum commonality between commercial (i.e., smartphones), critical (i.e., public safety) and V2X use cases [125]. There are several motivations to consider UE relaying for V2X communications in Rel. 17: extension of the network coverage and the SL transmission range, higher data rates, enhanced reliability, power saving, and spectral efficiency enhancements. UE relaying will be designed under Rel. 17 to operate in different network and spectrum scenarios. It shall be able to operate in network coverage, partial coverage, and out network coverage (see Fig. 38⁶⁰). It should also be able to operate using licensed, unlicensed, and ITS SL spectrum. It is expected that UE relaying will cover both UE-to-network relaying and UE-to-UE relaying (see Fig. 38), where relays forward the communication between a UE and the infrastructure or between two UEs, respectively. The work will primarily focus on utilizing a single

relay, which is referred to in 3GPP as single-hop NR SL-based relay [125]. However, forward compatibility for multi-hop relay support in a future release will be taken into account. The technical report entitled ‘Study on NR sidelink relay’ (3GPP TR 38.836) will collect the progresses related to UE relaying that will include (among others): relay discovery and selection procedures, service continuity during path switching between direct Uu connection and a connection via a UE-to-Network relay, and extending Rel-16 QoS management to UE-to-UE or UE-to-Network relaying.

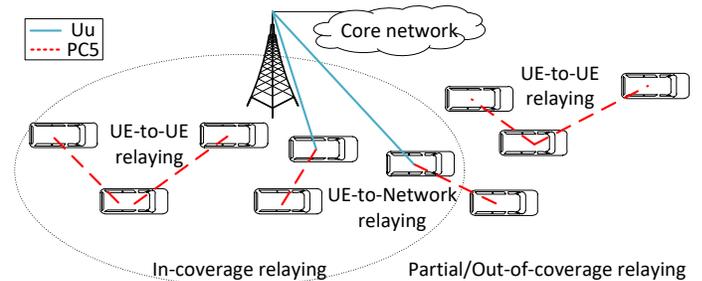


Fig. 38. UE relaying scenarios.

E. Uu multicast communications

Advanced eV2X use cases such as sensors monitoring, software updates, or update/download of high definition (HD) maps [13] could benefit from Uu multicasting or multicast broadcast services (MBS) to deliver the same information to different groups of UEs. In contrast to a broadcast transmission, a gNB sends content to a specific group of UEs in a multicast transmission. Multicast transmissions can support feedback to improve reliability. For MBS, the group of UEs can be determined based on UEs with a subscription to certain services or UEs which have indicated their interest in the content. MBS allows a more efficient use of the resources, since the same information can be sent to a group of UEs with fewer resources than in unicast transmissions. For instance, a vehicle may require real-time HD map updates. These updates may impact the navigation route of several vehicles in the same region. Thus, the updates can be delivered to a group of UEs in a multicast manner and avoid independent unicast transmissions to each UE over multiple resources. The group of UEs in a MBS transmission can be identified based on their location, but also on other factors like their heading. In fact, UEs near a traffic jam are interested in different content depending whether they are heading towards the traffic jam or in the opposite direction.

Besides V2X applications, there are other relevant applications and uses cases that can benefit from multicast and broadcast services (e.g, public safety and mission critical, IPTV, software updates, group communications and IoT) to improve user experience and service latency while efficiently utilizing the radio resources. Since Rel. 15 and Rel. 16 do not support MBS in NR, a new work item to support MBS has been agreed in Rel. 17 [126]. The objectives of this work item include the specification of RAN basic functions for MBS by providing a group scheduling mechanism to allow UEs to

⁶⁰Although UEs are represented in the figure by vehicles, the UE relaying scenario envisioned in Rel. 17 might also include smartphones, wearables, etc.

receive MBS, as well as the support for simultaneous operation with unicast reception. Within the work item, the support for basic mobility and dynamic change between multicast and unicast delivery with service continuity for a given UE should be specified. Necessary changes for addressing improved reliability of MBS, e.g., via feedback in the uplink, should also be specified. As a gNB may consist of a gNB-Control Unit (CU) and one or more gNB-Distributed Units (DUs), any required changes on the 5G-RAN architecture and interfaces between the gNB-CU and gNB-DUs should be defined assuming that the coordination functions are handled by the gNB-CU. The support for dynamic control (and its management) of the MBS transmission area within one gNB-DU should also be studied.

The work item on NR Multicast and Broadcast Services is limited to current Rel. 15 numerologies, physical channels and signals. Furthermore, lower priority will be given in this work item to any enhancements related to FR2. Although a flexible resource allocation between unicast services and MBS should be possible, allocating all the resources for MBS is not a mandatory requirement for this work item. The impact on UE implementation and UE complexity should be minimized to enable the deployment of the broadcast/multicast feature.

XII. CONCLUSIONS

This paper has presented an in-depth tutorial of the 3GPP Rel. 16 NR V2X standard, the first V2X standard based on the 5G NR air interface. This standard focuses mainly on the sidelink aspects that were not developed in Rel. 15 where the 5G NR air interface was first introduced. The 5G NR V2X standard has been developed to complement the capabilities of LTE V2X and be able to support more advanced eV2X use cases, including those related to connected and automated driving, where 5G NR is expected to play an important role. This tutorial will help the community to better understand the new standard and its advanced functionalities as a first step towards a thorough evaluation and further enhancement of its performance and capabilities. To this aim, the paper has also presented the evaluation methodology and the system and link level simulations that were defined in 3GPP for the development of Rel. 16. These can be used as a basis for the evaluation of 5G NR V2X and its future enhancements. Indeed, Rel. 16 has made significant contributions by developing the first V2X standard and the first NR SL standard based on the 5G NR air interface. However, improvements are possible, and some have already been identified as study items in Rel. 17. In this context, this paper also discusses some possible future enhancements related to beamforming, sidelink positioning, resource allocation (including the coexistence of different communication modes), relaying, and multicast communications. This discussion should help the research community to identify potential topics and contributions that could have a strong impact on Rel. 17 and beyond.

ACKNOWLEDGMENT

We would like to thank Konstantinos Manolakis and Lu Lei and for their useful inputs on the structure of the paper, the

physical layer section, and the QoS framework section. We would also like to thank Matthew Webb for his helpful comments and suggestions.

This work was supported in part by the Ministerio de Ciencia e Innovación (MCI), AEI and FEDER funds through the project TEC2017-88612-R, and the Ministerio de Universidades (FPU18/00691, IJC2018-036862-I).

The authors would also like to thank and provide full copyright acknowledgment to the following copyright holders: For Table IX in this paper (Table 5.4.4-1 in [51]):

"© 2020. 3GPP™ TSs and RSs are the property of ARIB, ATIS, CCSA, ETSI, TSDSI, TTA and TTC who jointly own the copyright in them. They are subject to further modifications and are therefore provided to you "as is" for information purposes only. Further use is strictly prohibited."

For Figs. 31 and 32 in this paper (Figure A-1 and A-2 in [103]): "© 2019. 3GPP™ TSs and RSs are the property of ARIB, ATIS, CCSA, ETSI, TSDSI, TTA and TTC who jointly own the copyright in them. They are subject to further modifications and are therefore provided to you "as is" for information purposes only. Further use is strictly prohibited."

REFERENCES

- [1] CAMP Vehicle Safety Communications Consortium, "Vehicle safety communications project: Task 3 final report: identify intelligent vehicle safety applications enabled by DSRC." National Highway Traffic Safety Administration, US Department of Transportation, Washington DC, 2005.
- [2] Car 2 Car Communication Consortium. "C2C-CC Manifesto." Overview of the C2C-CC System, 2007.
- [3] J. B. Kenney, "Dedicated Short-Range Communications (DSRC) Standards in the United States," *Proc. of the IEEE*, vol. 99, no. 7, pp. 1162-1182, July 2011.
- [4] A. Festag, "Cooperative intelligent transport systems standards in europe," *IEEE Communications Magazine*, vol. 52, no. 12, pp. 166-172, Dec. 2014.
- [5] F. Perry, "Overview of DSRC Messages and Performance Requirements," In UFTI DSRC and Other Communication Options for Transportation Connectivity Workshop, 2017.
- [6] Lin, Xingqin, J. G. Andrews, A. Ghosh and R. Ratasuk, "An overview of 3GPP device-to-device proximity services," *IEEE Communications Magazine*, vol. 52, no. 4, pp. 40-48, April 2014.
- [7] Chen *et al.*, "Vehicle-to-Everything (v2x) Services Supported by LTE-Based Systems and 5G," *IEEE Communications Standards Magazine*, vol. 1, no. 2, pp. 70-76, July 2017.
- [8] 3GPP, "RP-170837 New SI proposal: Study on evaluation methodology of new V2X use cases for LTE and NR," 3GPP TSG RAN Meeting #75, Dubrovnik, Croatia, March 2017.
- [9] 3GPP, "RP-181480. New SID: Study on NR V2X," 3GPP TSG RAN Meeting #80, La Jolla, USA, June 2018.
- [10] LG Electronics, Huawei, "RP-190766. New WID on 5G V2X with NR sidelink," 3GPP TSG RAN Meeting #83, Shenzhen, China, March 2019.
- [11] 3GPP, "TR 22.886: Study on enhancement of 3GPP Support for 5G V2X Services (v16.2.0, Release 16)," 3GPP, Tech. Rep., Dec. 2018.
- [12] 3GPP, "TS 22.186 Service requirements for enhanced V2X scenarios (v16.2.0, Release 16)," 3GPP, Tech. Spec., Jun. 2019.
- [13] 5GAA, "White Paper C-V2X Use Cases: Methodology, Examples and Service Level Requirements," Jun. 2019. [Online]. Available: https://5gaa.org/wp-content/uploads/2019/07/5GAA_191906_WP_CV2X_UCs_v1.pdf
- [14] R. Molina-Masegosa, J. Gozalvez, "LTE-V for Sidelink 5G V2X Vehicular Communications: A New 5G Technology for Short-Range Vehicle-to-Everything Communications," *IEEE Vehicular Technology Magazine*, vol. 12, no. 4, pp. 30-39, Dec. 2017.
- [15] 3GPP, "TR 36.885 Study on LTE-based V2X services (v14.0.0, Release 14)," 3GPP, Tech. Rep., Jul. 2016.

- [16] 3GPP, "TS 36.300 Evolved Universal Terrestrial Radio Access (E-UTRA) and Evolved Universal Terrestrial Radio Access Network (E-UTRAN); Overall description; Stage 2 (v15.7.0, Release 15)," 3GPP, Tech. Spec., Sept. 2019.
- [17] 3GPP, "TS 36.211 Evolved Universal Terrestrial Radio Access (E-UTRA); Physical channels and modulation (v14.3.0, Release 14)," 3GPP, Tech. Spec., June 2017.
- [18] 3GPP, "TS 36.213 Evolved Universal Terrestrial Radio Access (E-UTRA); Physical layer procedures (v14.8.0, Release 14)," 3GPP, Tech. Spec., October 2018.
- [19] 3GPP, "TS 36.331 Evolved Universal Terrestrial Radio Access (E-UTRA); Radio Resource Control (RRC); Protocol specification (v16.2.1, Release 16)," 3GPP, Tech. Spec., Sept. 2020.
- [20] A. Bazzi, G. Cecchini, A. Zanella, B.M. Masini, "Study of the Impact of PHY and MAC Parameters in 3GPP C-V2X Mode 4," *IEEE Access*, vol. 7, pp. 71685-71698, Nov. 2018.
- [21] M. Wang *et al.*, "Comparison of LTE and DSRC-Based Connectivity for Intelligent Transportation Systems," *Proc. of IEEE VTC2017-Spring*, Sydney, Australia, June 2017, pp. 1-5.
- [22] 3GPP, "TS 23.285 Architecture enhancements for V2X services (v14.3.0, Release 14)," 3GPP, Tech. Spec., June 2017.
- [23] 3GPP, "TS 36.213 Evolved Universal Terrestrial Radio Access (E-UTRA); Physical layer procedures (v15.7.0, Release 15)," 3GPP, Tech. Spec., Sept. 2019.
- [24] 3GPP, "TS 36.214 Evolved Universal Terrestrial Radio Access (E-UTRA); Physical layer; Measurements (v16.0.0, Release 16)," 3GPP, Tech. Spec., Dec. 2019.
- [25] R. Molina-Masegosa, J. Gozalvez, M. Sepulcre, "Configuration of the C-V2X Mode 4 Sidelink PC5 Interface for Vehicular Communications," *Proc. of 14th MSN 2018*, Shenyang, China, Dec. 2018.
- [26] M. Gonzalez-Martin, 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.
- [27] A. Mansouri, V. Martinez, J. Hürri, "A First Investigation of Congestion Control for LTE-V2X Mode 4," *Proc. of WONS*, Wengen, Switzerland, Jan. 2019, pp. 1-8.
- [28] B. Toghi, M. Saifuddin, Y.P. Fallah, MO Mughal. "Analysis of Distributed Congestion Control in Cellular Vehicle-to-everything Networks," *Proc. of IEEE VTC2019-Fall*, Honolulu, USA, Sept. 2019, pp. 1-7.
- [29] M. Boban *et al.*, "Connected roads of the future: Use cases, requirements, and design considerations for vehicle-to-everything communications," *IEEE Vehicular Technology Magazine*, vol. 13, no. 3, pp. 110-123, Sept. 2018.
- [30] SAE International, Standard J3016: "Taxonomy and Definitions for Terms Related to On-Road Motor Vehicle Automated Driving Systems", September 2016.
- [31] 3GPP, "TR 22.885 Study on LTE support for Vehicle-to-Everything (V2X) services (v14.0.0, Release 14)," 3GPP, Tech. Rep., Dec. 2015.
- [32] 3GPP, "TS 22.185 "Service requirements for V2X services; Stage 1 (v15.0.0, Release 15)," 3GPP, Tech. Spec., June 2018.
- [33] SA WG2, "SP-190625. Study on Architectural enhancements for 5G multicast-broadcast services," 3GPP TSG SA Meeting #85, Newport Beach, USA, Sept. 2019.
- [34] 3GPP, "TS 23.285 Architecture enhancements for V2X services (v16.2.0, Release 16)," 3GPP, Tech. Spec., Dec. 2019.
- [35] 3GPP, "TS 23.501 System architecture for the 5G System (5GS); Stage 2 (v16.4.0, Release 16)," 3GPP, Tech. Spec., March 2020.
- [36] 3GPP, "TS 38.211 NR; Physical channels and modulation (v16.0.0, Release 16)," 3GPP, Tech. Spec., March 2020.
- [37] 3GPP, "TS 38.101-1 NR; User Equipment (UE) radio transmission and reception; Part 1: Range 1 Standalone (v16.3.0, Release 16)," 3GPP, Tech. Spec., March 2020.
- [38] 3GPP, "TS 38.101-2 NR; User Equipment (UE) radio transmission and reception; Part 2: Range 2 Standalone (v16.3.1, Release 16)," 3GPP, Tech. Spec., March 2020.
- [39] Y. Zou *et al.*, "Impact of major RF impairments on mm-Wave communications using OFDM waveforms," *Proc. of the 2016 IEEE Global Communications Conference (GLOBECOM) Workshops*, Washington, USA, Dec. 2016, pp. 1-7.
- [40] J. Campos, "Understanding the 5G NR Physical Layer," Keysight Technologies, Nov. 2017.
- [41] E. Dahlman, S. Parkvall and J. Skold, *5G NR: The Next Generation Wireless Access Technology*. San Diego, CA, USA: Academic Press, 2018.
- [42] 3GPP, "TS 38.214 NR; Physical layer procedure for data (v16.3.0, Release 16)," 3GPP, Tech. Spec., September 2020.
- [43] 3GPP, "TS 38.211 NR; Physical channels and modulation (v15.3.0, Release 15)," 3GPP, Tech. Spec., September 2018.
- [44] LG Electronics, "R1-1913601. Summary of RAN1 Agreements/Working assumptions in WI 5G V2X with NR sidelink," 3GPP TSG RAN WG1 Meeting #99, Reno, USA, Nov. 2019.
- [45] 3GPP, "TR 37.985 Overall description of Radio Access Network (RAN) aspects for Vehicle-to-everything (V2X) based on LTE and NR (v16.0.0, Release 16)," 3GPP, Tech. Rep., Sept. 2020.
- [46] 3GPP MCC Support, "Final Report of 3GPP TSG RAN WG1 #101-e", 3GPP TSG RAN WG1 Meeting #102-e, August 2020
- [47] 3GPP MCC Support, "Final Report of 3GPP TSG RAN WG1 #95 v1.0.0", 3GPP TSG RAN WG1 Meeting #95, Spokane, USA, November 2019.
- [48] 3GPP MCC Support, "Final Report of 3GPP TSG RAN WG1 #AH_1901 v1.0.0", 3GPP TSG RAN WG1 Meeting #96, Taipei, Taiwan, March 2019.
- [49] 3GPP MCC Support, "Final Report of 3GPP TSG RAN WG1 #100-e v2.0.0", 3GPP TSG RAN WG1 Meeting #100bis-e, April 2020.
- [50] J. Schliez and A. Roessler, "Device to Device Communication in LTE," Whitepaper, Rhode & Schwarz, 2015.
- [51] 3GPP, "TS 23.287 Architecture enhancements for 5G System (5GS) to support Vehicle-to-Everything (V2X) services (v16.2.0, Release 16)," 3GPP, Tech. Spec., March 2020.
- [52] 3GPP, "TS 38.133 NR; Requirements for support of radio resource management (v16.5.0, Release 16)," 3GPP, Tech. Spec., September 2020.
- [53] 3GPP, "TS 38.212 NR; Multiplexing and channel coding (v16.3.0, Release 16)," 3GPP, Tech. Spec., August 2020.
- [54] 3GPP MCC Support, "Final Report of 3GPP TSG RAN WG1 #100bis-e v1.0.0", 3GPP TSG RAN WG1 Meeting #101-e, June 2020.
- [55] 3GPP MCC Support, "Draft Report of 3GPP TSG RAN WG1 #99 v0.1.0", 3GPP TSG RAN WG1 Meeting #99, Reno, USA, Nov. 2019.
- [56] 3GPP MCC Support, "Final Report of 3GPP TSG RAN WG1 #98bis v2.0.0", 3GPP TSG RAN WG1 Meeting #98bis, Chongqing, China, October 2019.
- [57] Huawei, HiSilicon, "R1-1910054. Sidelink physical layer structure for NR V2X," 3GPP TSG RAN WG1 Meeting #98bis, Chongqing, China, Oct. 2019.
- [58] P. Frenger, S. Parkvall and E. Dahlman, "Performance comparison of HARQ with Chase combining and incremental redundancy for HSDPA," *Proc. of IEEE VTC2001-Fall*, Atlantic City, USA, Oct. 2001, pp. 1829-1833.
- [59] Y. Qi, M. Hunukumbure, H. Nam, H. Yoo and S. Amuru, "On the Phase Tracking Reference Signal (PT-RS) Design for 5G New Radio (NR)," *Proc. of IEEE VTC2018-Fall*, Chicago, USA, Aug. 2018, pp. 1-5.
- [60] T. J. Roupheal, Ed., *Wireless Receiver Architectures and Design Antennas, RF, Synthesizers, Mixed Signal, and Digital Signal Processing*. Oxford, Great Britain: Elsevier Inc., 2015.
- [61] H.-J. Zepernick and A. Zinger, *Pseudo Random Signal Processing: Theory and Application*. West Sussex, England: John Wiley & Sons LTD, 2005.
- [62] 3GPP, "TR 38.885 NR; Study on Vehicle-to-Everything (v16.0.0, Release 16)," 3GPP, Tech. Rep., March 2019.
- [63] 3GPP, "TS 38.331 NR; Radio Resource Control (RRC); Protocol specification (v16.2.0, Release 16)," 3GPP, Tech. Spec., September 2020.
- [64] Huawei, HiSilicon, "R1-2001553. Remaining details of sidelink synchronization mechanisms," 3GPP TSG RAN WG1 Meeting #100bis-e, April 2020.
- [65] vivo, "R1-1911421. Sidelink synchronization mechanism," 3GPP TSG RAN WG1 Meeting #98bis, October 2019.
- [66] 3GPP, "TS 38.213 NR; Physical layer procedures for control (v16.3.0, Release 16)," 3GPP, Tech. Spec., September 2020.
- [67] Huawei, HiSilicon, "R1-1903068. Sidelink discovery for NR V2X," 3GPP TSG RAN WG1 Meeting #96, Athens, Greece, March 2019.
- [68] 3GPP MCC Support, "Final Report of 3GPP TSG RAN WG1 #96 v2.0.0", 3GPP TSG RAN WG1 Meeting #96, Athens, Greece, March 2019.
- [69] R. Comroe and D. Costello, "ARQ Schemes for Data Transmission in Mobile Radio Systems," in *IEEE Journal on Selected Areas in Communications*, vol. 2, no. 4, pp. 472-481, July 1984.

- [70] B. M. Popovic, "Generalized Chirp-Like polyphase sequences with optimum correlation properties," *IEEE Trans. Inf. Theory*, vol. 38, no. 4, pp. 1406-1409, July 1992.
- [71] D. C. Chu, "Polyphase Codes with Good Periodic Correlation Properties", *IEEE Trans. Inf. Theory*, vol. 18, no. 4, pp. 531-532, July 1972.
- [72] Huawei, HiSilicon, "R1-2001555. Remaining details of physical layer procedures for sidelink," 3GPP TSG RAN WG1 Meeting #100e-bis, April, 2020.
- [73] M. Nakamura, Y. Awad and S. Vadgama, "Adaptive control of link adaptation for high speed downlink packet access (HSDPA) in W-CDMA," *International Symposium on Wireless Personal Multimedia Communications*, Honolulu, USA, Dec. 2002, pp. 382-386.
- [74] 3GPP, "TS 38.321 NR; Medium Access Control (MAC) protocol specification (v16.2.1, Release 16)," 3GPP, Tech. Spec., September 2020.
- [75] 3GPP, "TS 36.213 Evolved Universal Terrestrial Radio Access (E-UTRA); Physical layer procedures (v14.3.0, Release 14)," 3GPP, Tech. Spec., Jun. 2017.
- [76] 3GPP, "TS 38.215 NR; Physical layer measurements (v16.3.0, Release 16)," 3GPP, Tech. Spec., September 2020.
- [77] MediaTek Inc., "R1-1906558. Discussion on V2X Physical layer procedure," 3GPP TSG RAN WG1 Meeting #97, Reno, USA, May 2019.
- [78] C. Ubeda Castellanos et al., "Performance of Uplink Fractional Power Control in UTRAN LTE," *Proc. of IEEE VTC2008-Spring*, Singapore, May 2008, pp. 2517-2521.
- [79] ETSI, "TS 136 331; LTE; Evolved Universal Terrestrial Radio Access (E-UTRA); Radio Resource Control (RRC); Protocol specification (v14.2.2)," ETSI, Tech. Spec., May 2017.
- [80] ETSI, "TS 103 574; Intelligent Transport Systems (ITS); Congestion Control Mechanisms for the C-V2X PC5 interface; Access layer part (v1.1.1)," ETSI, Tech. Spec., Nov. 2018.
- [81] 3GPP, "TS 23.303 Proximity-based services (ProSe); Stage 2 (v15.1.0, Release 15)," 3GPP, Tech. Spec., June 2018.
- [82] 3GPP, "TR 23.786 Study on architecture enhancements for the Evolved Packet System (EPS) and the 5G System (5GS) to support advanced V2X services (v16.1.0, Release 16)," 3GPP, Tech. Rep., June 2019.
- [83] 3GPP, "TR 21.905 Vocabulary for 3GPP Specifications (v16.0.0, Release 16)," 3GPP, Tech. Rep., June 2019.
- [84] ISO, "ISO 17419:2018 Intelligent transport systems - Cooperative systems - Globally unique identification," ISO, Tech. Spec., 2018.
- [85] 3GPP, "TS 38.300 NR; NR and NG-RAN Overall Description; Stage 2; (v16.1.0, Release 16)," 3GPP, Tech. Spec., March 2020.
- [86] OPPO, "R2-2000194 Summary of [108#50][V2X] Feature List and UE caps," 3GPP TSG RAN WG2 Meeting #109 electronic, 24 Feb - 6 Mar, 2020.
- [87] Huawei, HiSilicon, "R2-1915983. Running CR to TS 38.331 for 5G V2X with NR sidelink," 3GPP TSG RAN WG2 Meeting #108, Reno, USA, November 2019.
- [88] 3GPP, "TS 23.503 Technical Specification Group Services and System Aspects; System architecture for the 5G System (5GS) (v16.3.0, Release 16)," 3GPP, Tech. Spec., Dec. 2019.
- [89] 3GPP, "TS 23.288 Technical Specification Group Services and System Aspects; Architecture enhancements for 5G System (5GS) to support network data analytics services (v16.2.0, Release 16)," 3GPP, Tech. Spec., Dec. 2019.
- [90] 3GPP, "TS 36.201 Evolved Universal Terrestrial Radio Access (E-UTRA); LTE physical layer: General description (v15.2.0, Release 15)," 3GPP, Tech. Spec., Dec. 2018.
- [91] 3GPP, "TS 36.321 Evolved Universal Terrestrial Radio Access (E-UTRA); Medium Access Control (MAC) protocol specification (v8.10.0, Release 8)," 3GPP, Tech. Spec., Sept. 2011.
- [92] Intel Corporation, "R1-1810777. Analysis of NR Uu Support for Advanced V2X Use Cases," 3GPP TSG RAN WG1 Meeting #94bis, Chengdu, China, Oct. 2018.
- [93] 3GPP, "TS 36.212 Evolved Universal Terrestrial Radio Access (E-UTRA); Multiplexing and channel coding (v16.0.0, Release 16)," 3GPP, Tech. Spec., Dec. 2019.
- [94] Huawei, HiSilicon, "R1-1812212. NR Uu and LTE Uu support and enhancement for advanced V2X use cases," 3GPP TSG RAN WG1 Meeting #95, Spokane, USA, Nov. 2018.
- [95] Intel Corporation, "RP-192533. Status Report to TSG (WI NR Mobility Enhancements)," 3GPP TSG RAN Meeting #86, Sitges, Spain, Dec. 2019.
- [96] Intel Corporation, "RP-192278. Status Report to TSG (WI NR Mobility Enhancements)," 3GPP TSG RAN Meeting #85, Newport Beach, USA, Sept. 2019.
- [97] 3GPP, "TS 38.304 NR; User Equipment (UE) procedures in Idle mode and RRC Inactive state (v15.6.0, Release 15)," 3GPP, Tech. Spec., Dec. 2019.
- [98] Samsung Electronics and Qualcomm Incorporated, "R2-1911445. Introduction of T312 in NR," 3GPP TSG RAN WG2 Meeting #107, Prague, CZ, Aug. 2019.
- [99] 3GPP, "TS 36.133 Evolved Universal Terrestrial Radio Access (E-UTRA); Requirements for support of radio resource management (v16.3.0, Release 16)," 3GPP, Tech. Spec., Sept. 2019.
- [100] LG Electronics, "RP-200129. Revised WID on 5G V2X with NR sidelink," 3GPP TSG RAN Meeting #87e, March 2020..
- [101] 3GPP MCC Support, "Final Report of 3GPP TSG RAN WG1 #AH_1901 v1.0.0," 3GPP TSG RAN WG1 Meeting #96, Taipei, Taiwan, March 2019.
- [102] ETSI MCC, "Report of 3GPP TSG RAN WG2 #105bis," 3GPP TSG RAN WG2 Meeting #105bis, Xi'an, China, April 2019.
- [103] 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.
- [104] 3GPP, "TR 38.901 "Study on channel model for frequencies from 0.5 to 100 GHz (v16.1.0, Release 16)," 3GPP, Tech. Rep., Dec. 2019.
- [105] M. Boban *et al.*, "Multi-Band Vehicle-to-Vehicle Channel Characterization in the Presence of Vehicle Blockage," *IEEE Access*, vol. 7, pp. 9724-9735, Jan. 2019.
- [106] M. Boban, T. T. V. Vinhoza, M. Ferreira, J. Barros and O. K. Tonguz, "Impact of Vehicles as Obstacles in Vehicular Ad Hoc Networks," *IEEE Journal on Selected Areas in Communications*, vol. 29, no. 1, pp. 15-28, Jan. 2011.
- [107] ETSI, "TR 103 257-1 Intelligent Transport Systems (ITS); Access Layer; Part 1: Channel Models for the 5.9 GHz frequency band (v1.1.1), ETSI, Tech. Rep., May 2019.
- [108] M. Boban, X. Gong and W. Xu, "Modeling the Evolution of Line-of-Sight Blockage for V2V Channels," *Proc. of IEEE VTC2016-Fall*, Montreal, QC, Sept. 2016, pp. 1-7.
- [109] 3GPP TR 36.814: "Evolved Universal Terrestrial Radio Access (E-UTRA); Further advancements for E-UTRA physical layer aspects".
- [110] LG Electronics, "RP-193257. New WID on NR sidelink enhancement," 3GPP TSG RAN Meeting #86, Sitges, Spain, Dec. 2019.
- [111] Huawei, HiSilicon, "R1-1911885. Sidelink synchronization mechanisms for NR V2X," 3GPP TSG RAN WG1 Meeting #99, Reno, USA, Nov. 2019.
- [112] LG Electronics, FirstNet, "RP-201384. Study on scenarios and requirements of in-coverage, partial coverage, and out-of-coverage positioning use cases," 3GPP TSG RAN Meeting #88e, Jul. 2020.
- [113] COST Action CA15104, IRACON, "Whitepaper on New Localization Methods for 5G Wireless Systems and the Internet-of-Things," April 2018.
- [114] A. Kakkavas, M. H. Castañeda Garcia, R. A. Stirling-Gallacher and J. A. Nossek, "Multi-Array 5G V2V Relative Positioning: Performance Bounds," *Proc. of the 2018 IEEE Global Communications Conference (GLOBECOM)*, Abu Dhabi, United Arab Emirates, Dec. 2018, pp. 206-212.
- [115] 3GPP, "TS 36.321 Evolved Universal Terrestrial Radio Access (E-UTRA); Medium Access Control (MAC) protocol specification (v15.7.0, Release 15)," 3GPP, Tech. Spec., Sept. 2019.
- [116] LG Electronics, "R1-1912588. Discussion on resource allocation for Mode 2," 3GPP TSG RAN WG1 Meeting #99, Reno, USA, Nov. 2019.
- [117] LG Electronics, "RP-192745. Summary of email discussion on Rel-17 sidelink enhancement," 3GPP TSG RAN Meeting #86, Sitges, Spain, Dec. 2019.
- [118] 3GPP, "TR 38.885 NR; Study on Vehicle-to-Everything (v1.0.0, Release 16)," 3GPP, Tech. Rep., Nov. 2018.
- [119] Ericsson, "R1-1908914. Miscellaneous issues on resource allocation," 3GPP TSG RAN WG1 Meeting #98, Prague, CZ, Aug. 2019.
- [120] Fujitsu, "R1-1912080. Dynamic Sidelink Bi-mode Transmission under gNB Configuration," 3GPP TSG RAN WG1 Meeting #99, Reno, USA, May 2019.
- [121] Ericsson, "R1-1910534. Resource allocation for Mode-2 transmissions," 3GPP TSG RAN WG1 Meeting #98bis, Chongqing, China, Oct. 2019.
- [122] Fujitsu, "R1-1905375. Geographic Information based Dynamic TFRP Resource Selection Procedure in NR-V2X," 3GPP TSG RAN WG1 Meeting #96bis, Xian, China, April 2019.
- [123] Huawei, HiSilicon, "R1-1911884. Sidelink resource allocation mode 2 for NR V2X," 3GPP TSG RAN WG1 Meeting #99, Reno, USA, Nov. 2019.

- [124]3GPP, "TR 36.746 Study on further enhancements to LTE Device to Device (D2D), User Equipment (UE) to network relays for Internet of Things (IoT) and wearables (v15.1.1, Release 15)," 3GPP, Tech. Rep., April 2018.
- [125]3GPP, "RP-201474. Revised SID: Study on NR Sidelink Relay," 3GPP TSG RAN Meeting #89e, Sept. 2020.
- [126]Huawei, "RP-193248. New Work Item on NR support of Multicast and Broadcast Services," 3GPP TSG RAN Meeting #86, Sitges, Spain, Dec. 2019.
- [127]Fernandez, A. E., *et al.*, "5GCAR scenarios, use cases, requirements and KPIs." Fifth Generation Communication Automotive Research and innovation, Tech. Rep. D2.1 (2017).



Mario H. Castañeda Garcia received the B.Sc. degree in electrical engineering from the Universidad Nacional Autónoma de Honduras, Tegucigalpa, Honduras, in 2001, and the M.Sc. degree in communications engineering and the Dr.-Ing. degree from the Technische Universität München (TUM), Munich, Germany, in 2004 and 2014, respectively.

From 2005 to 2014, he was a research and teaching assistant with the Institute for Circuit Theory and Signal Processing, TUM, where he was involved in several industry and government projects in the areas of signal processing for wireless communication systems and for satellite-based positioning. In 2014, he joined the Munich Research Center, Huawei Technologies Düsseldorf GmbH, Munich, where he has been involved in several EU funded research projects as well as in 5G research and standardization activities. He was the recipient of the Best Overall Paper Award at IEEE VTC Fall 2018. His research interests comprise multi-antenna technologies, millimeter-wave systems, V2X communication and wireless positioning.



Alejandro Molina-Galan received a Telecommunications Engineering Degree in 2016 and a Master's Degree in Telecommunications Engineering in 2018, both from the Universidad Miguel Hernández (UMH) de Elche, Spain. He received the Best Student award in the Telecommunications Engineering Degree

and the Master's Degree by UMH. He joined the UWICORE research laboratory in October 2018, and he was awarded in September 2019 a highly competitive PhD fellowship by the Spanish government. He is currently pursuing his PhD on 5G and beyond V2X networks for connected and automated vehicles.



Mate Boban received the Diploma degree in informatics from the University of Zagreb, Croatia, and the Ph.D. degree in electrical and computer engineering from Carnegie Mellon University, Pittsburgh, PA, USA, in 2004 and 2012, respectively. He is a Principal Research Engineer with

Huawei Munich Research Center, Germany. Before Huawei, he was with NEC Labs Europe and Apple. He is an alumni of the

Fulbright Scholar program. He has co-chaired several IEEE workshops and conferences and has been involved in European Union funded projects (5G-CAR, DRIVE-C2X, and TEAM) as a Work Package Leader and editor of deliverables. He is actively involved in key industry and standardization bodies: 3GPP, 5GAA, and ETSI. His current research interests include resource allocation, machine learning applied to wireless communication systems (in particular, V2X and IIoT), and channel modeling. He coauthored three papers that received the Best Paper Award, at IEEE VTC Spring 2014, IEEE VNC 2014, and EuCAP 2019.



Javier Gozalvez received an electronics engineering degree from the Engineering School ENSEIRB (Bordeaux, France), and a PhD in mobile communications from the University of Strathclyde, Glasgow, U.K. Since October 2002, he is with the Universidad Miguel Hernández de Elche (UMH), Spain, where he is

currently Full Professor and Director of the UWICORE laboratory. At UWICORE, he leads research activities in the areas of vehicular networks, 5G and Beyond, and industrial wireless networks. He has received several awards at international and national conferences, the best research paper award from the Journal of Network and Computer Applications (Elsevier) in 2014, the Runner-up prize for the "Juan López de Peñalver" award of the Royal Academy of Engineering in Spain that recognizes the most notable Spanish engineers aged below 40, and the IEEE VTS Outstanding Service Award in 2019. He is an elected member to the Board of Governors of the IEEE Vehicular Technology Society (IEEE VTS) and was the 2016 and 2017 President of the IEEE VTS. He has served as IEEE Distinguished Speaker and IEEE Distinguished Lecturer for the IEEE VTS. He is the Editor in Chief of the IEEE Vehicular Technology Magazine, and the Founder and General Co-Chair for the IEEE Connected and Automated Vehicles Symposium. He was the General Co-Chair for the IEEE VTC-Spring 2015 conference in Glasgow (UK), ACM VANET 2013, ACM VANET 2012 and ISWCS 2006, and TPC Co-Chair for 2011 IEEE VTC-Fall and 2009 IEEE VTC-Spring. He is a scientific evaluator and reviewer for the European Commission and various national scientific and technical agencies.



Baldomero Coll-Perales received his M.Sc. and Ph.D. degrees in Telecommunications Engineering from the Universidad Miguel Hernández (UMH) de Elche, Spain. He is currently a Juan de la Cierva Research Fellow at the UWICORE laboratory (UMH). He has held visiting research positions at Hyundai Motor

Europe Technical Center (Germany), Italian National Research Council (Italy), WINLAB-Rutgers University (USA) and Institute of Telecommunications-King's College London (UK). His research interests lie in the field of advanced mobile and wireless communication systems, including the design of

device-centric technologies for future wireless Beyond 5G network and connected and automated vehicles. He serves/has served as Associate Editor for the International Journal of Sensor Networks and Springer's Telecommunication Systems. He has served as Track Co-Chair for IEEE VTC-Fall 2018, and as member of TPC in over 35 international conferences.



Taylan Şahin received his B.S. degree in electrical and electronics engineering from the Middle East Technical University, Ankara, Turkey, in 2014, and his M.Sc. degree in communications engineering from the Technical University of Munich, Germany, in 2016. Since February 2017, he has been a doctoral candidate at the Technical University of Berlin, Germany, and worked as a student in Huawei Germany Research Center, Munich, where he also conducted his master's work prior to that. Since December 2020, he is a radio access research specialist at Nokia Bell Labs, Munich. His research interests include radio resource management aspects of radio access networks, and V2X communications, targeting the next generations of wireless mobile networks.



Apostolos Kousaridas received his Ph.D. from the Department of Informatics & Telecommunications at the University of Athens, Athens. He holds B.Sc. degree in Informatics and M.Sc. degree in Information Systems from the Department of Informatics at Athens University of Economics and Business. Currently, he is a principal research engineer of the Huawei Research Center in Munich, contributing to the design of 5G and beyond 5G communication systems, by generating patents and standardization contributions to 3GPP. He has disseminated over 50 publications and he is also serving as delegate to the 5G Automotive Association (5GAA) and vice-chair of the 5G-PPP Automotive WG. He has participated in several impacting large and medium-scale European R&D projects, having technical and management roles. His research interests include vehicular communications, wireless networks and artificial intelligence.

Annex A.2 Publication



How does 5G NR V2X Mode 2 Handle Aperiodic Packets and Variable Packet Sizes?

Alejandro Molina-Galan¹, Baldomero Coll-Perales¹, Luca Lusvarghi^{2,3}, Javier Gozalvez¹, and Maria Luisa Merani^{2,3}

¹UWICORE laboratory, Universidad Miguel Hernández de Elche, Elche (Alicante), Spain.

²Dipartimento di Ingegneria Enzo Ferrari, Università degli Studi di Modena e Reggio Emilia, Italy.

³Consorzio Nazionale Interuniversitario per le Telecomunicazioni (CNIT), 43124 Parma, Italy

{alejandro.molinag, bcoll, j.gozalvez}@umh.es; {luca.lusvarghi5, marialuisa.merani}@unimore.it

Abstract—5G NR V2X complements LTE V2X to support advanced V2X services for connected and automated driving. 5G NR V2X introduces novel features at the MAC layer that are designed to cope with potential packet collisions, and that could help address the LTE V2X MAC inefficiencies observed under aperiodic traffic of variable size. This is the case of the re-evaluation mechanism that is a mandatory MAC feature of 5G NR V2X, and that seeks to avoid possible packet collisions detected before a vehicle transmits in selected resources. Evaluations conducted to date of 5G NR V2X do not consider the re-evaluation mechanism, and have focused on traffic patterns that do not fully account for the traffic variability of advanced V2X services. This paper extends the current state of the art with the first evaluation of a fully standard compliant 5G NR V2X implementation under the traffic patterns recommended by 3GPP for advanced V2X services. Our study shows that 5G NR V2X mode 2 still faces MAC challenges when using semi-persistent scheduling (SPS) to efficiently support aperiodic traffic of variable size.

Keywords—5G NR V2X, Aperiodic traffic, CAV, C-V2X, cellular V2X, connected automated vehicles, eV2X, LTE V2X, Mode 2, packet collisions, re-evaluation, variable packet size.

I. INTRODUCTION

3GPP has published in Release 16 the 5G NR V2X standard that represents the first 5G NR (New Radio) release that supports sidelink (SL) or direct V2V (Vehicle-to-Vehicle) communications [1]. 5G NR V2X (or NR V2X) is designed to complement, and not replace, its predecessor V2X technology based on the LTE air interface, a.k.a. LTE V2X. The aim of LTE V2X is to support basic safety applications, relying on the exchange of broadcast messages among neighboring vehicles. NR V2X also allows unicast and groupcast communications, and introduces novel features and functionalities to support advanced services with stringent requirements. To this aim, NR V2X includes two novel operating modes: mode 1, where the cellular infrastructure is in charge of selecting the communication resources for every V2V communication, and mode 2, where vehicles autonomously communicate with no infrastructure support.

In this paper we focus on analyzing the efficiency of 5G NR V2X mode 2 to support advanced V2X services. These services are expected to generate V2X traffic patterns characterized by variable packet sizes and aperiodic generation times according to 3GPP [2]. The study in [3] evaluated the performance of LTE V2X mode 4 (the counterpart of NR V2X mode 2) under periodic and aperiodic traffic with fixed and variable packet sizes. [3] demonstrated that the MAC of LTE V2X mode 4 faces certain inefficiencies when vehicles generate aperiodic traffic of variable packet size that increase packet collisions.

NR V2X mode 2 can also operate using SPS scheduling with similar procedures as those present in LTE V2X mode 4

[1]. A key mandatory feature introduced in the MAC of NR V2X mode 2 (and that is not present in LTE V2X mode 4) is the re-evaluation mechanism that is designed with the objective of detecting and avoiding imminent packet collisions [4]. Previous studies have evaluated the performance of NR V2X mode 2 under different traffic patterns [5]-[8]. However, these studies do not implement the re-evaluation mechanism and do not consider the variability in both time generation and size of V2X packets as recommended in the 3GPP guidelines [2] that present models for the generation of packets for advanced V2X services. This is for example the case of the Cooperative Perception Service (CPS) that generates aperiodic messages with different sizes depending on the number of objects perceived by the sensors of the vehicle [9]. In this context, this paper complements the state of the art with the first evaluation of a fully standard compliant implementation of NR V2X mode 2 using SPS scheduling and the re-evaluation mechanism. The evaluation considers different traffic patterns including periodic and aperiodic traffic with fixed and variable packet sizes based on the models reported in the 3GPP guidelines [2]. To this aim, we have implemented an ns-3 simulator that is fully compliant with the 5G NR V2X mode 2 standard [4][10]. Our results show that NR V2X mode 2 also faces challenges to efficiently handle aperiodic packets of variable size with SPS scheduling despite its new MAC features. The rest of the paper is organized as follows: Section II gives an overview of NR V2X mode 2. Section III presents the simulation environment and Section IV the performance evaluation of NR V2X mode 2. Finally, Section V reports the main conclusions of this paper.

II. 5G NR V2X MODE 2

A. Resource grid

The resource grid of NR V2X is organized into slots in the time domain and Resource Blocks (RBs) in the frequency domain (Fig. 1). The slot duration and RB bandwidth depend on the utilized OFDM numerology (μ) or subcarrier spacing (SCS). 5G NR V2X supports a scalable SCS given by $2^\mu \times 15$ kHz, where μ can be equal to 0, 1, 2, or 3. The slot duration is given by $2^{-\mu}$ ms and the RB consists of 12 consecutive subcarriers with the same SCS. All vehicles utilize the same SCS at a particular region. RBs are grouped into sub-channels (Fig. 1). A sub-channel is formed by RBs of the same slot. The sub-channel size (i.e., the number of RBs per sub-channel) can vary but it is common for all communicating vehicles. A sub-channel represents the smallest unit for a SL data transmission or reception. Data packets are transmitted in Transport Blocks (TB) that are carried in the Physical Sidelink Shared Channel (PSSCH). A TB can occupy one or several sub-channels depending on the packet size, the sub-channel size, and the utilized Modulation and Coding Scheme (MCS). In NR V2X, each TB is associated with a Sidelink Control Information (SCI) that indicates the resources used by the associated TB, as well as further information required for decoding the TB. The SCI in NR V2X is transmitted in two stages, and it is transmitted together with the TB in the same slot.

This 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.

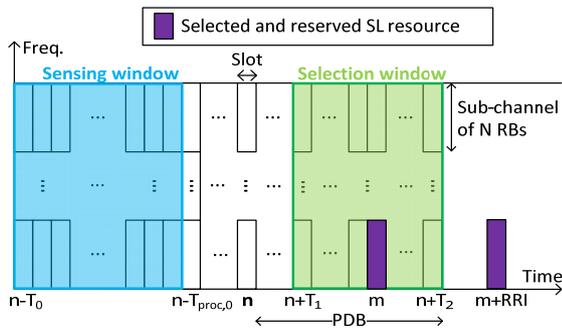


Fig. 1. NR V2X channelization and illustration of resource allocation mode 2 operation with SPS (when $T_2 = \text{PDB}$).

B. Resource Allocation

Mode 2 can operate using a dynamic scheme (DS) or a semi-persistent scheduling scheme (SPS). The DS selects new resources for each TB and can only reserve resources for the retransmissions of that TB. A vehicle can select and reserve resources for the transmission of several TBs when utilizing SPS. We should note that we distinguish in this work between a selected resource and a reserved resource. A reserved resource is a selected resource that a vehicle reserves for a future transmission by notifying neighboring vehicles using the SCI. Therefore, reserved resources provide stability to the operation of NR V2X mode 2. We focus on this paper on the SPS that, unlike DS, does not require the implementation of retransmissions to generate reserved resources.

SPS selects resources for a number of consecutive *Reselection Counter* TBs. The time period between the resources selected for the transmission of consecutive TBs is defined by the *Resource Reservation Interval (RRI)*. The possible values of the RRI are $\{0, [1:1:99], 100:100:1000\}$ ms. The *Reselection Counter* is randomly set within an interval that depends on the selected RRI. If $RRI \geq 100$ ms, this counter is randomly set within the interval $[5, 15]$. If $RRI < 100$ ms, the counter is randomly set within the interval¹ $[5 * C, 15 * C]$ where $C = 100 / \max(20, RRI)$ [4]. A vehicle initiates a new *Reselection Counter* when it selects new SL resources. A vehicle selects new SL resources when it generates a new TB and it does not have resources to perform the transmission that fits the size or latency requirements of the TB. The *Reselection Counter* is decremented by one after transmitting a TB. When *Reselection Counter* is depleted, the vehicle also selects new resources for the new TB with probability² $(1-P)$.

To select new SL resources, a vehicle first defines the selection window where it looks for candidate resources to transmit a TB. The selection window includes all resources within the range of slots $[n+T_1, n+T_2]$ (Fig. 1) [10], where n is the slot at which new resources must be selected. T_1 is the processing time (in slots) required by a vehicle to identify candidate resources and select new SL resources for transmission. The value of T_2 must be included within the range $T_{2min} \leq T_2 \leq \text{PDB}$, where PDB (Packet Delay Budget) is the latency deadline (in slots) by which the TB must be transmitted. Once the selection window is defined, the vehicle must identify the candidate resources within the selection window. A candidate resource is defined by a slot in time and L contiguous sub-channels in frequency. L must be selected such that the newly generated TB and its associated SCI fit in the candidate resource.

When a vehicle is not transmitting, it senses the SL resources during the sensing window that is defined in the

range of slots $[n-T_0, n-T_{proc,0})$ (Fig. 1). T_0 defines the length of the sensing window and it can take values in slots equivalent to 100 ms or 1100 ms. $T_{proc,0}$ represents the number of slots required to complete the sensing procedure. During the sensing process, the vehicle uses the SCI received from other vehicles to determine which candidate resources from the selection window should be excluded. The vehicle also measures the RSRP of the transmissions associated with the SCIs received from other vehicles.

Mode 2 defines a 2-step algorithm to select new resources [10][4]. Step 1 excludes candidate resources in the selection window. The first exclusions relate to the half-duplex operation since a vehicle cannot sense reservations from other vehicles in the slots where it is transmitting. In this case, the vehicle excludes candidate resources of the selection window to avoid packet collisions with other vehicles that could send reservations in the slots of the sensing window where the vehicle was transmitting. Then, the vehicle excludes the candidate resources based on the reservations received from other vehicles in the SCIs during the sensing window. In this case, the candidate resources are only excluded if the measured RSRP is higher than a threshold. After executing all exclusions in step 1, the vehicle checks whether the percentage of remaining available candidate resources in the selection window is equal or higher than $X\%$. If not, the RSRP threshold is increased by 3 dB, and step 1 is repeated again. In step 2, the vehicle randomly selects the SL resource for the transmission of the TB from the list of available candidate resources. A vehicle could optionally select up to 32 SL resources from the list of available candidate resources for the transmission and retransmissions of the TB.

The new re-evaluation mechanism introduced in NR V2X forces vehicles to check the availability of the selected resources. For example, a vehicle that has selected a SL resource at slot m must execute again step 1 at slot $m-T_3$ to check if the selected resource is still available (i.e., it is not excluded because another vehicle has reserved it). 3GPP defines this new execution of step 1 as re-evaluation check. The vehicle can optionally perform additional re-evaluation checks apart from the one executed at slot $m-T_3$. T_3 is the maximum time allowed for a vehicle (in slots) to complete the resource selection process. The re-evaluation check process works as follows. Consider n' in Fig. 2 as the slot at which a vehicle executes a re-evaluation check. The vehicle defines a new selection window SW' that starts at slot $n'+T_1$ and ends at slot $n'+T_2'$. T_2' must be within the range $T_{2min} \leq T_2' \leq \text{PDB} - (n'-n)$. The vehicle executes then step 1 over the candidate resources in SW' in order to evaluate the currently available and excluded resources. If the selected resource at slot m is now excluded, then the vehicle has detected what is called in 3GPP standards a re-evaluation [10]. This re-evaluation detection triggers the execution of step 2 to select a new SL resource among the currently available resources in SW' [4] (Fig. 2)³. The vehicle does not execute step 2 if the initially selected resource remains available.

It is important to distinguish two cases where the vehicle executes a re-evaluation check. The first case happens when the vehicle has selected new SL resources in the selection window due to a reselection of resources (Fig. 2). We should note that the execution of a re-evaluation check is mandatory in this case according to the standard [4]. The second case occurs when the vehicle has not utilized a reservation announced in the SCI and it generates a new TB. Let us consider that the vehicle has not utilized a reservation at slot

¹ C is a constant to compute the *Reselection Counter* interval for RRI values lower than 100 ms.

² P is the probability to keep the same resources.

³ The selection of new SL resources with re-evaluation does not imply that the vehicle initiates a new *Reselection Counter*.

$m+RRI$ and that it generates a TB at slot n_2 , with $n_2 > m+RRI$. In this case, the vehicle could use a selected resource located at the slot $m+Y*RRI$, where Y is the minimum integer that fulfills $m+Y*RRI > n_2$, if $(m+Y*RRI - n_2) \leq PDB$. Since the resource located at $m+Y*RRI$ has not been reserved by the vehicle with a previous SCI, this resource is then considered a selected resource and it is up to UE implementation whether the vehicle also executes a re-evaluation check over it.

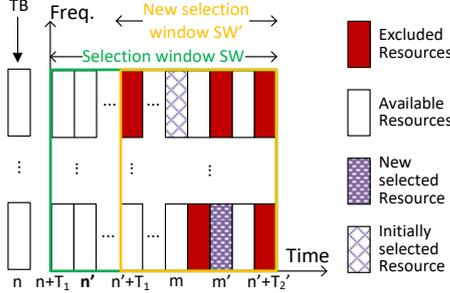


Fig. 2. Re-evaluation mechanism of NR V2X mode 2.

III. SIMULATION ENVIRONMENT

The evaluation of NR V2X mode 2 has been conducted using an in-house developed ns-3 simulator that implements NR V2X mode 2 following the 3GPP standard [10][4]. The evaluation considers the reference 3GPP highway scenario of 5 Km and 3 lanes in each driving direction [2]. The density of vehicles in the scenario varies between 25 and 100 veh/km, and we consider that the vehicle speed is 70 km/h. The vehicles transmit V2X packets following the periodic and aperiodic traffic models included in the 3GPP evaluation guidelines [2]. The periodic traffic model considers that 190-byte packets are generated with an inter-packet arrival time and latency requirement equal to $\{20, 100\}$ ms. The 100 ms and 20 ms inter-packet arrival times are referred to as low and high traffic intensity scenarios by 3GPP. The aperiodic traffic model considers that the packets are generated with an inter-packet arrival time $\tau=c+r$, where c is a constant and r is an exponentially distributed random variable. The latency requirement of aperiodic traffic is set to c [2]. The low and high traffic intensity scenarios are modeled for aperiodic traffic considering that $c=\bar{r}=50$ ms and $c=\bar{r}=10$ ms, respectively. \bar{r} is the mean of the exponential random variable. For the aperiodic traffic, the variable packet size is uniformly distributed in the $[200, 1200]$ byte range with a 200-byte step, while the fixed packet size is set to 200 bytes.

NR V2X mode 2 is configured to operate over a 20 MHz channel with a subcarrier spacing of 30 kHz in the 5.9 GHz frequency band. This channel bandwidth corresponds to 51 RBs. The sub-channel size is set to 12 RBs, which results in 4 sub-channels available in the channel. The transmission power has been set to 23dBm and the sensitivity has been set to -103.5dBm following [11]. The pathloss is modeled using the reference 3GPP models available in [2], and the shadowing effects are modeled using a log-normal distribution with zero mean and a standard deviation of 3dB [2]. Spatial shadowing correlation is also modeled following [2]. We consider that TBs are transmitted using the MCS index 13 with 16QAM modulation and coding rate equal to 0.5. With this MCS, packets of 190, 200, 400, 600, 800, 1000 and 1200 bytes are transmitted using 1, 1, 2, 3, 3, 4 and 4 sub-channels, respectively. Lookup tables from 3GPP working documents that relate the BLER (Block Error Rate)-SINR (Signal to Interference plus Noise Ratio) are used to model the transmissions of the TB [12] and SCI [13]. At the MAC layer, $T_{proc,0}$, T_0 and T_3 are set equal to 1 slot, 2200 slots and 5 slots respectively. The limits of the selection window T_1 , T_{2min} and

T_2 are set equal to 2 slots, 2 slots and PDB respectively. The PDB is set equal to the latency requirement of the traffic in number of slots. The percentage X of resources that must be available after the execution of step 1 in mode 2 is set to 20%, and the RSRP threshold that determines whether a resource is excluded or not is set to its minimum value (i.e. -128 dBm) following the results obtained in [6]. In addition, the probability P is set to 0 which indicates that a vehicle will always execute the 2-step SPS process when the *Reselection Counter* is depleted.

A key parameter for the operation of NR V2X mode 2 is the *RRI*. We evaluate two strategies for the *RRI* selection: average *RRI* and minimum *RRI*. These strategies set the *RRI* to the average and minimum inter-packet arrival time, respectively. Note that for periodic traffic the two *RRI* strategies result in the same *RRI* value. However, for aperiodic traffic, the average *RRI* strategy sets the *RRI* value to $c + \bar{r}$, while the minimum *RRI* strategy sets it to c .

Finally, we consider that vehicles execute a re-evaluation check at the slot that is mandatory by the standard (i.e. at slot $m-T_3$, where m is the slot where the resource was initially selected). We also consider that the vehicle executes a re-evaluation check in all selected resources, including the selected resource available after an unutilized reservation.

IV. RESULTS

A. Metrics

We use the following metrics to evaluate the performance of NR V2X mode 2:

- Packet Delivery Ratio (PDR): fraction of correctly received TBs over the total number of transmitted TBs.
- 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.
- Size Reselection Ratio (SRR): fraction of TBs that produce a reselection due to the size of the TBs over the total number of transmitted TBs [3].
- Latency Reselection Ratio (LRR): fraction of TBs that produce a reselection due to the latency requirement of the TBs over the total number of transmitted TBs [3].
- Unutilized Reservation Ratio (URR): fraction of previously reserved resources that are not utilized for transmitting a TB over the total number of resource reservations. This metric does not account for unutilized reservations that are considered in the size and latency reselection ratios [3].
- Unused Sub-channels Ratio (USR): fraction of unused sub-channels in the resources used to transmit a TB over the total number of sub-channels that are in the resources used to transmit a TB [3].
- Reselection Counter Depletion Ratio (RCDR): fraction of Reselection Counters that deplete over the total number of initiated Reselection Counters.
- Re-evaluation Check Ratio (ReCR): fraction of TBs that are 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 one re-evaluation detection and trigger the selection of a new resource over the total number of transmitted TBs.
- Channel Busy Ratio (CBR): fraction of sub-channels that experience an RSSI higher than a threshold within an observation window of $100 \cdot 2^u$ slots.

B. Single traffic scenario

We first analyze the impact of the V2X traffic characteristics on the operation and performance of the SPS scheduling scheme of NR V2X mode 2 considering that all vehicles in the scenario implement the same traffic pattern. Fig. 3 compares the PDR achieved by NR V2X mode 2 when vehicles transmit periodic traffic with fixed packet size and aperiodic traffic with fixed or variable packet size. Results are also reported in Fig. 3 for the two implemented *RRI* strategies that set the *RRI* value to the average and minimum inter-packet arrival time, and for the low (Fig. 3-left) and high (Fig. 3-right) traffic intensity scenarios that are characterized by average inter-packet arrival times of 100 ms and 20 ms, respectively. The CBR or channel load generated in each of these scenarios is shown in Table I. The results reported in Fig. 3 show that NR V2X mode 2 achieves a better PDR performance when vehicles generate packets periodically compared to the case when they generate packets aperiodically. The obtained results show that the differences in the PDR achieved by NR V2X mode 2 with periodic and aperiodic traffic increase with the vehicle density and traffic intensity (i.e., with the increasing CBR., Table I). For example, Fig. 3.a-left shows small PDR differences (below 2% when the distance between the transmitter and receiver is 300 m) between periodic traffic and aperiodic traffic with fixed packet size under the lowest evaluated vehicle density and traffic intensity. Under the highest evaluated vehicle density and traffic intensity, Fig. 3.b-right shows that the PDR of NR V2X mode 2 decreases by 36.3% for aperiodic traffic with fixed packet size with respect to periodic traffic when the distance between the transmitter and receiver is 300 m. It is important to note that this is the case despite the lower CBR values experienced for aperiodic traffic with fixed packet size with respect to periodic traffic⁴ (see Table I).

TABLE I. CBR (IN %) IN THE SINGLE TRAFFIC SCENARIO

Veh. Density (veh/km)	Periodic		Aperiodic Fixed size		Aperiodic Variable Size	
	Low	High	Low	High	Low	High
25	6	28	6	26	14	46
100	22	82	22	68	40	89

TABLE II. NR V2X OPERATION IN THE SINGLE TRAFFIC SCENARIO

RRI strategy	Periodic		Aperiodic Fixed size		Aperiodic Variable Size	
	Low	High	Low	High	Low	High
a) Reselection Counter Depletion Ratio (RCDR) in %						
Min RRI	100	100	83.9	10.8	44.2	5.5
Avg RRI			0.6	0	0.3	0
b) Latency Reselection Ratio (LRR) in %						
Min RRI	0	0	1	5	1	5
Avg RRI			55	58.1	55	57
c) Size Reselection Ratio (SRR) in %						
Min RRI	0	0	0	0	5	5
Avg RRI			0	0	26	27
d) Unutilized Reservation Ratio (URR) in %						
Min RRI	0	0	60	59	56	57
Avg RRI			6	6	4	4
e) Unused Sub-channels Ratio (USR) in %						
Min RRI	0	0	0	0	29	29
Avg RRI			0	0	27	27

Table II reports a set of key metrics that show that the SPS scheduling scheme of NR V2X mode 2 achieves better performance under the presence of periodic traffic than aperiodic traffic. This is for example shown by means of the RCDR (Table II.a) that is 100% when vehicles generate packets periodically. Achieving an RCDR equal to 100% indicates that vehicles always deplete their *Reselection*

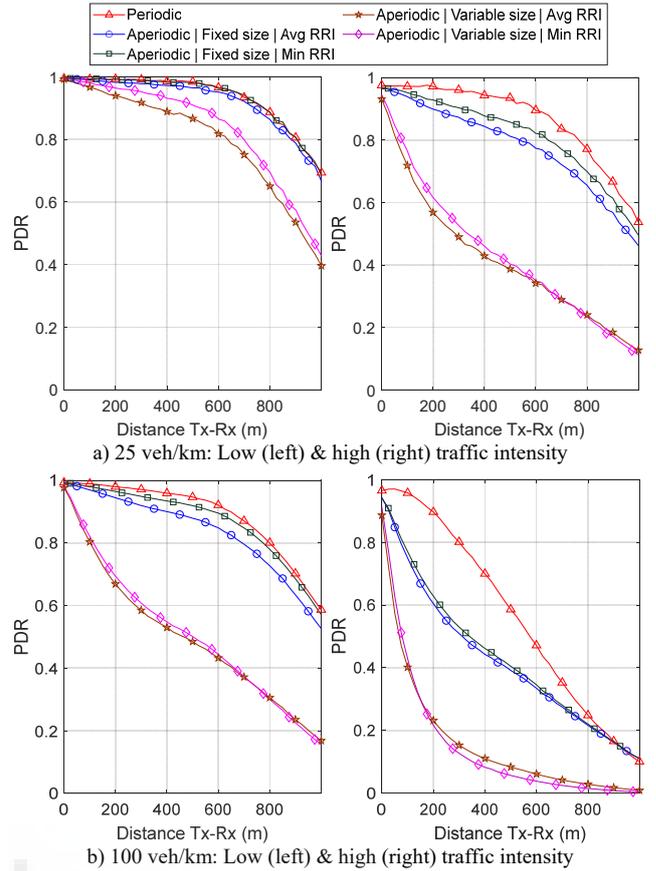


Fig. 3. Packet Delivery Ratio (PDR) in the single traffic scenario.

Counter reservations to perform the transmissions of their packets (or TBs) without executing additional resource reselections. This behavior brings significant stability to the operation of NR V2X mode 2 with SPS that results in the better PDR performance reported in Fig. 3. Table II.a) also shows that the RCDR is below 100% for the aperiodic traffic patterns. For example, for the case of aperiodic traffic with fixed packet size, the variations in the time between packets can cause additional resource reselections before the *Reselection Counter* depletes. These are referred to as latency reselections and are computed by means of the LRR in Table II.b). As it is shown in Table II.b), the chosen *RRI* strategy highly impacts the LRR generated by aperiodic traffic. When the average *RRI* strategy is implemented, more than 50% of the packets trigger a resource reselection, and therefore the *Reselection Counter* is very seldomly depleted (see Table II.a)). The higher number of resource reselections makes the SPS scheduling scheme of NR V2X mode 2 more unstable and likely to provoke packet collisions. This is the case because with the increasing number of resource reselections it is more likely that there are several vehicles selecting new resources around the same time (i.e., with their selection windows overlapping), and that they end up choosing the same ones. The minimum *RRI* strategy significantly reduces the LRR compared to the average *RRI* strategy (see Table II.b)⁵, which results in the higher RCDR levels reported in Table II.a). The different RCDR levels experienced by the minimum *RRI* strategy under aperiodic traffic with fixed packet size in the scenarios with low and high traffic intensity (i.e., 84% vs 11%) are due to the different ranges for selecting the *Reselection Counter* when the *RRI* is set to 50 ms (low

⁴ Periodic and aperiodic traffic with fixed packet size generate on average the same number of packets. The lower CBR values measured for aperiodic traffic in Table I are due to its higher packet collisions (see Fig. 4).

⁵ The latency reselections with minimum *RRI* only occur in a corner case where the new TB to be transmitted is generated in the middle of a slot where a reservation is located.

traffic intensity) and 10 ms (high traffic intensity). As it is described in Section II, the resulting ranges are [10, 30] and [25, 75], respectively, which makes more likely to deplete the *Reselection Counter* when the *RRI* is set to 50 ms. However, the side effect of utilizing the minimum *RRI* strategy are the higher values of unutilized reservations measured by the *URR* (Table II.d)). The unutilized reservations can also negatively impact the operation of the SPS scheduling scheme of NR V2X mode 2. This is the case because unutilized reservations are discarded by the vehicles as part of step 1 of the SPS scheduling scheme (see Section II), although at the end they are not going to hold the transmission of a TB. This reduces the number of available candidate resources over which vehicles randomly select from during the execution of step 2 of SPS (see Section II). Therefore, the likelihood that two vehicles end up selecting the same resources increases. The operation of NR V2X mode 2 under aperiodic traffic with fixed packet size when implementing the average and minimum *RRI* strategies has shown some trade-offs in terms of *RCDR*, *LRR*, and *URR*. The results reported in Fig. 3 show that these trade-offs compensate each other and only a slightly higher *PDR* is achieved when NR V2X mode 2 is configured with the minimum *RRI* strategy.

Fig. 3 shows that NR V2X mode 2 sees its *PDR* degrade when vehicles generate aperiodic traffic with variable packet size. Aperiodic traffic with variable packet sizes challenges the operation of SPS of NR V2X mode 2. In particular, aperiodic traffic with variable packet size experiences not only the latency reselections and unutilized reservations that aperiodic traffic is prone to (see Table II). The variable packet sizes also produce additional resource reselections when the new generated TB does not fit in the previously reserved or selected resource. These reselections are referred to as size reselections and are computed in Table II.c) by means of the *SRR*. Size reselections represent an additional source of instability for the operation of NR V2X mode 2 that is reflected in the reduction (~50%) of the *RCDR* (Table II.a)) and *PDR* (Fig. 3) with respect to aperiodic traffic with fixed packet size. Table II.c) shows that the implemented *RRI* strategy also impacts the *SRR*. In particular, the obtained results show a significant reduction (~80%) in the *SRR* when NR V2X mode 2 is configured with the minimum *RRI* strategy with respect to the average *RRI* strategy. This is the case because with the minimum *RRI* strategy, the resource reselections not caused by the depletion of the *Reselection Counter* are almost limited to size reselections. Size reselection triggered to accommodate a TB that occupies more sub-channels than the reserved ones can also be used to transmit upcoming TBs of equal or smaller sizes. This is at the cost of leaving some sub-channels unused (see Table II.e)) that other vehicles cannot utilize since they are reserved. The increasing *USR* has also a negative impact on the operation of SPS of NR V2X mode 2 because it reduces the availability of candidate resources and therefore increases the risks of packet collisions. With the average *RRI* strategy, the additional latency reselections result in that the selected/reserved resources are fitted more often to the size of the new generated packets. Then, size reselections need to be triggered when forthcoming packets request additional sub-channels than the ones previously reserved in the latency reselection. Like it has been shown above for aperiodic traffic with fixed packet size, there exist different trade-offs with aperiodic traffic with variable packet size that impact the operation of the NR V2X mode 2 when it is configured with the minimum or average *RRI* strategies. In this case, results reported in Fig. 3 show that these trade-offs also compensate each other, and only slight differences in the *PDR* are achieved when NR V2X mode 2 is configured with the minimum *RRI* or average *RRI* strategy.

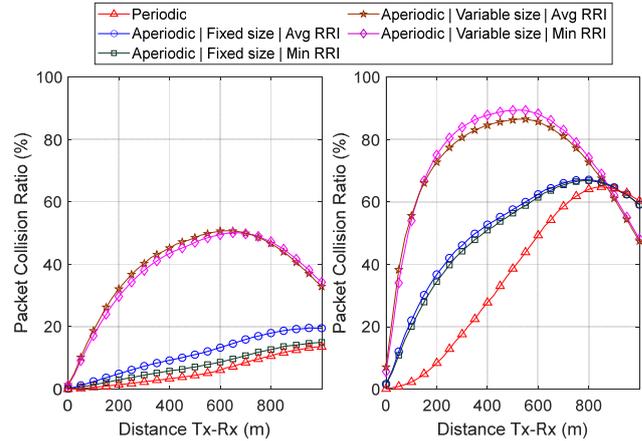


Fig. 4. Packet Collision Ratio (PCR) under low (left) & high (right) traffic intensity; 100 veh/km – similar trends observed for other vehicle densities.

The analysis performed above has highlighted the challenges that NR V2X mode 2 faces to efficiently support aperiodic traffic with fixed and variable packet sizes. The analysis is revealing that these challenges end up making the SPS scheduling scheme of NR V2X mode 2 more unstable with the potential risk of causing additional packet collisions. Fig. 4 reports the Packet Collision Ratio (PCR) as a function of the distance between the transmitter and the receiver for all the evaluated traffic patterns and *RRI* strategies when the vehicle density is 100 veh/km. The results reported in Fig. 4 help identifying that the degradation in *PDR* between the scenarios with periodic and aperiodic traffic (Fig. 3) is due to higher packet collisions experienced with aperiodic traffic. For example, Fig. 4 shows that aperiodic traffic with variable packet size increases the PCR from 2.3% to 38% and from 17.6% to 83.9% with respect to periodic traffic under low and high traffic intensity when the distance between the transmitter and the receiver is 300 m.

As it was demonstrated in [3], LTE V2X mode 4 shows similar inefficiencies and high PCR figures when transmitting aperiodic traffic with variable packet size. NR V2X mode 2 faces similar challenges despite introducing the re-evaluation mechanism originally designed with the objective to reduce the packet collisions. We measure the operation of re-evaluation in terms of the re-evaluation check (ReCR) and detection (ReDR) ratios that indicate the fraction of TBs that are eligible to be re-evaluated and those over which the re-evaluation is actually detected because a potential collision is identified, respectively. It is important to recall that only TBs that are to be transmitted in a selected resource (i.e., not reserved) are eligible for the re-evaluation check. The reselections that occur when transmitting aperiodic packets (because of the size and latency reselections, see Table II) result in ReCR levels between 50% and 70% independently of the vehicle density. The ReDR experiences levels up to 30%. Higher ReDR levels are measured with increasing vehicle densities because of the higher CBR and thus likelihood that two vehicles select the same resources. The observed ReDR levels show that re-evaluation is active. However, the results reported in Fig. 4 show that it is not effective as NR V2X mode 2 still experiences high PCR levels when transmitting aperiodic traffic. Fig. 4 questions the impact of re-evaluation on the performance of NR V2X mode 2 when transmitting aperiodic packets. For example, when aperiodic traffic with variable packet size is transmitted, NR V2X mode 2 configured with the minimum *RRI* strategy experiences an ReDR level of ~30%, while this level reduces to 2.5% (low traffic intensity) and to 7.5% (high traffic intensity) for the average *RRI* strategy. In spite of these ReDR differences, both strategies experience similar PCR levels. The obtained results

show that even although re-evaluation is avoiding the detected collisions for a non-negligible percentage of TBs, it is not capable of counteracting the high packet collisions caused by the instability that aperiodic traffic is introducing in the operation of the SPS scheduling scheme of NR V2X mode 2.

C. Mixed traffic scenario

Finally, we analyze the operation of NR V2X mode 2 in a more realistic scenario where vehicles generate traffic with different patterns. In particular, we consider that 80% and 20% of the vehicles in the scenario implement the low and high traffic intensity models characterized by 100 ms and 20 ms average inter-packet arrival times, respectively. These mixed traffic intensities result in the CBR levels reported in Table III which are in line with the weighted values of the low and high traffic intensity ones reported in Table I. The PDR differences observed in the mixed traffic scenario between periodic and aperiodic traffic (Fig. 5-left) are again due to the higher PCR values experienced by aperiodic traffic (Fig. 5-right).

It should be noted that these high PCR values are experienced by aperiodic traffic in the mixed traffic scenario in spite of the increased ReDR levels with respect to the single traffic scenario. For example, the ReDR levels increase from 30% (low and high traffic intensity) in the single traffic scenario to 44.6% in the mixed scenario when transmitting aperiodic packets with variable size and using the minimum *RRI* strategy. The ReDR levels increase from 2.5% (low traffic intensity) or 7.5% (high traffic intensity) in the single traffic scenario to 16% in the mixed traffic scenario when transmitting aperiodic packets with variable size and using the average *RRI* strategy. However, despite the increase in the ReDR, we can observe similar trends in the PCR reported in Fig. 5-right to those reported in Fig. 4.

TABLE III. CBR (IN %) IN THE MIXED TRAFFIC SCENARIO

Veh. Density (veh/km)	Periodic	Aperiodic Fixed size	Aperiodic Variable Size
25	11	10	23
100	38	35	62

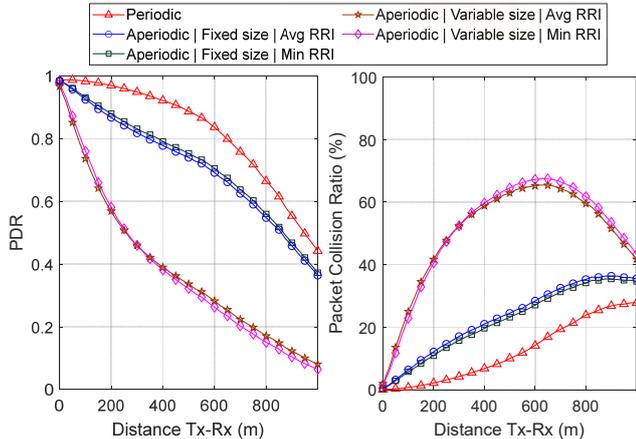


Fig. 5. Packet Delivery Ratio (PDR) -left- and Packet Collision Ratio (PCR) -right- for the mixed scenario (100 veh/km - similar trends observed for other vehicle densities).

V. CONCLUSIONS

This paper has presented the first fully standard compliant evaluation of 5G NR V2X mode 2 with SPS scheduling under periodic and aperiodic traffic of fixed and variable packet size in accordance with the 3GPP traffic model recommendations. The configuration of NR V2X mode 2 has considered two strategies that set the *RRI* to the minimum or to the average inter-packet arrival times. Both strategies result in multiple

trade-offs, but none reduces the resource management instability of 5G NR V2X mode 2 observed when transmitting aperiodic traffic with variable packet size using SPS. These instabilities reduce the PDR and increase packet collisions with respect to scenarios where vehicles generate packets periodically. The obtained results present similar trends as those observed with LTE V2X mode 4 when transmitting aperiodic traffic, despite the new and mandatory re-evaluation mechanism of 5G NR V2X mode 2 that identifies and avoids possible packet collisions. These results call for further improvements to the MAC of 5G NR V2X mode 2 in order to efficiently handle variability in packet generation and size. An option that has been discussed is adapting the MCS to reduce size reselections triggered when the generated TB does not fit in the reserved resources. However, this is not exempt from challenges since the objective of adapting the MCS (a.k.a. link adaptation) is to adapt the error protection (and data in a packet) as a function of the link quality. Adapting the MCS to avoid size reselections could compromise the correct reception of the TB based on the link quality, and requires a careful study and design.

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] 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.
- [4] 3GPP, "TS 38.321 NR; Medium Access Control (MAC) protocol specification (v16.7.0, Release 16)," 3GPP, Tech. Spec., Jan. 2022.
- [5] 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.
- [6] 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.
- [7] 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.
- [8] 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.
- [9] G. Thandavarayan, M. Sepulcre and J. Gozalvez, "Generation of Cooperative Perception Messages for Connected and Automated Vehicles," *IEEE Transactions on Vehicular Technology*, vol. 69, no. 12, pp. 16336-16341, Dec. 2020.
- [10] 3GPP, "TS 38.214 NR; Physical layer procedure for data (v16.8.0, Release 16)," 3GPP, Tech. Spec., Jan. 2022.
- [11] 5GAA, "V2X Functional and Performance Test Report; Test Procedures and Results," 5GAA Report, Apr. 2019.
- [12] Huawei, HiSilicon, "R1-1900852. Link level evaluations on sidelink for NR V2X," 3GPP TSG RAN WG1 Ad-Hoc Meeting 1901, Taipei, Jan. 2019.
- [13] Ericsson, "R1-1903180. Link level evaluations of NR PSCCH," 3GPP TSG-RAN WG1 Meeting #96, Athens, Greece, March 2019.

Annex A.3 Publication



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 including, for example, the impact of retransmissions. The studies reported in [8], [9] and [10] evaluate NR V2X mode 2 considering also aperiodic traffic of fixed size. In [8], the authors analyze the performance of NR V2X mode 2 under different configurations. The work reported in [9] compares the performance of the two scheduling schemes of NR V2X mode 2 under different message generation patterns, and [10] compares the performance of NR V2X mode 2 with the performance of LTE V2X mode 4. Despite their relevant contributions, the studies reported in [7]-[10] did not implement the re-evaluation mechanism despite being a mandatory MAC feature in 3GPP standards. In addition, these studies only consider periodic or aperiodic traffic of fixed size. However, the 3GPP evaluation methodology guidelines for NR V2X reported in [2] recommend traffic generation models for advanced V2X services that also include traffic of variable size in line with the message patterns characteristic of Day 2 or Day 3 V2X services such as cooperative perception [11] or

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.

maneuver coordination [12]. In this context, this paper extends the current state of the art by analyzing the operation and performance of NR V2X mode 2 considering periodic or aperiodic traffic of fixed or variable size. In particular, this study provides an in-depth analysis of the operation of the re-evaluation mechanism introduced in NR V2X mode 2 to clearly understand the conditions under which re-evaluation can be effective in avoiding packet collisions. We should note that the first study that evaluated the system-level performance of NR V2X mode 2 with re-evaluation was reported by the authors in [13]. This study analyzed the performance of NR V2X mode 2 under different data traffic patterns when vehicles use the Semi-Persistent Scheduling (SPS) scheme and there are no retransmissions. The study showed that the performance of NR V2X mode 2 degrades when vehicles generate aperiodic traffic of variable size compared to when the traffic is periodic and of fixed size. This degradation is due to an increment of packet collisions despite the fact that re-evaluation has been specifically designed and introduced to detect and avoid packet collisions. The study in [13] provides a first system level evaluation of NR V2X mode 2, but does not explain why re-evaluation is not effective in avoiding packet collisions under certain scenarios. Understanding why this is the case is critical to optimize NR V2X mode 2. In addition, it is necessary to analyze and understand if the reported observations hold for other scenarios, for example, using a different scheduling scheme as well as when utilizing retransmissions. In this context, 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 NR V2X mode 2. To the authors' knowledge, this is the first study that analyzes when (NR V2X mode 2 configuration and scenario) and why re-evaluation is effective or not to detect and avoid packet collisions. In particular, this study analyzes the effectiveness of re-evaluation to detect and avoid packet collisions when NR V2X mode 2 operates with the 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 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. The analysis is done 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 NR

V2X 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. For the sake of brevity, we refer to NR V2X mode 2 as mode 2 in the rest of the paper.

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 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 slot are grouped into sub-channels that represent the smallest unit for SL data transmission or reception (see Slot and Sub-channel (n PRBs) in Fig. 1¹). 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.

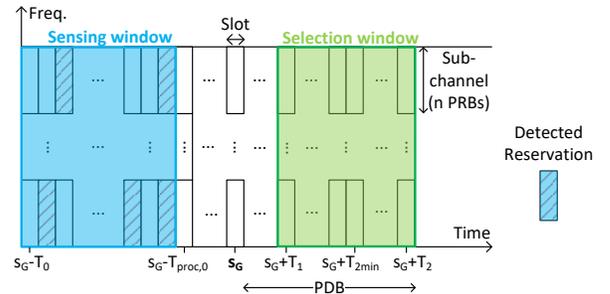


Fig. 1. NR V2X channelization and illustration of resource allocation in mode 2 (when $T_2 = PDB$).

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-

¹ The remaining variables and processes represented in this figure are explained below in Section II.A since they are related to resource allocation.

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.

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 [14]. 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*. According to 3GPP standards [6], the vehicle randomly sets the *Reselection Counter* within the interval $[5, 15]$ when $RRI \geq 100$ ms, and within the interval $[5 * C, 15 * C]$, where $C = 100 / \max(20, RRI)$, when $RRI < 100$ ms.

New resources (for both the DS and SPS schemes) are selected in the so-called selection window portrayed in Fig. 1 [14]. 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, and $T_1 \leq T_{proc,1}$, 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 transmitted². According to 3GPP standards [14], 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 $[s_G - T_0, s_G - T_{proc,0}]$. According to 3GPP standards [14], T_0 can be equal to 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][14]. 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 ³. 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) [14]. 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].

² The constraint $T_{2min} \leq T_2 \leq PDB$ prevents the vehicle from violating the PDB of the TB that must be transmitted.

³ The resource exclusions due to half-duplex operation have to consider all possible *RRI* values of the *RRI* list.

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 excluded⁴, where T_3 is equal⁵ 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 $[s_G' + T_1, s_G' + T_2']$. 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 violate the PDB of the TB to be transmitted. Step 1 is executed over the candidate resources in 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 [14]. 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.

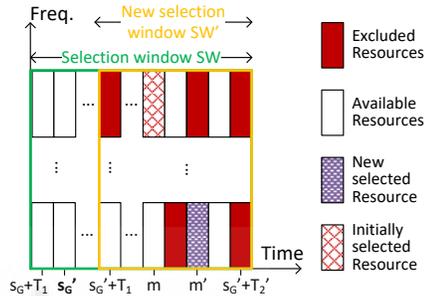


Fig. 2. Operation of re-evaluation under NR V2X mode 2.

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

⁴ 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].

⁵ T_3 is equivalent to 3, 5, 9 or 17 slots for a SCS of 15, 30, 60 or 120 kHz, respectively.

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_L . The current reservation at s_{R4} does not respect the latency requirements of the TB since $s_{R4} > s_L$. 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).

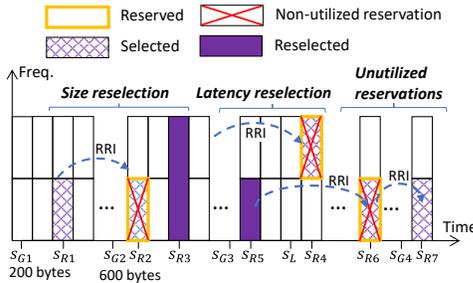


Fig. 3. Impact of packet variability on SPS under NR V2X mode 2.

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. If the re-evaluation check cannot be performed due to insufficient processing capabilities, then the vehicle uses the same selected resource for 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 $T_2 = \text{PDB} = \text{RRI}$, and that $\mu = 0$, $T_1 = 1$ slot and $T_3 = 5$ slots. In this case, the percentage of selection window resources that are not eligible for a re-evaluation check is equal to $\{5, 25, 50\}$ % when $\text{RRI} = \{100, 20, 10\}$ 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.

As demonstrated in the remainder of this section, only collisions between selected and reserved resources can be identified by the re-evaluation detection.

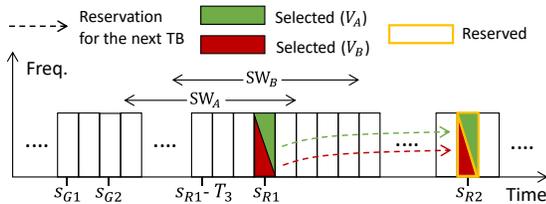


Fig. 4. Persistent collisions not detected by re-evaluation when using SPS.

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. Note that, in Fig. 5(a), V_A is using the SPS scheme since it scheduled its next transmission on reserved resources, whereas V_B might be employing either the SPS or the DS scheme since it is transmitting on selected resources. Accordingly, this type of collision can be detected when both vehicles employ the SPS scheme but also when SPS (V_A) and DS (V_B) coexist.

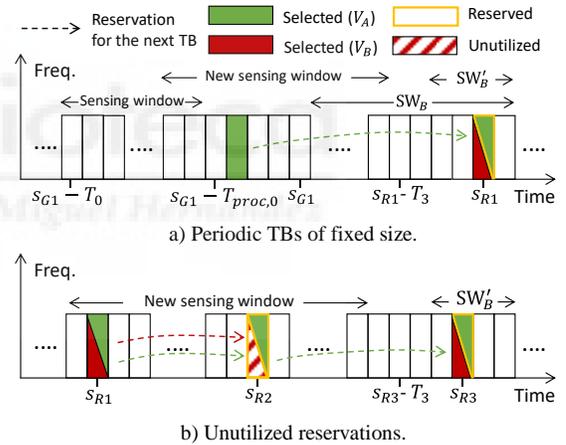


Fig. 5. Re-evaluation detection when vehicles use the same RRI .

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 $s_{R2} = s_{R1} + \text{RRI}$ unutilized. In this case, V_B will not be able to announce the reserved resources at slot $s_{R3} = s_{R2} + \text{RRI}$, and will transmit its next TB in selected, rather than reserved, resources. As a result, V_B runs a re-evaluation check at slot $s_{R3} - T_3$, 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 $RRIs$. 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.

Like in Fig. 5(a), note that V_A is using the SPS scheme since it scheduled its next transmission on reserved resources, whereas V_B might be employing either the SPS or the DS scheme since it is transmitting on selected resources in Fig. 6(a). Accordingly, this type of collision can be detected when both vehicles employ the SPS scheme but also when SPS (V_A) and DS (V_B) coexist.

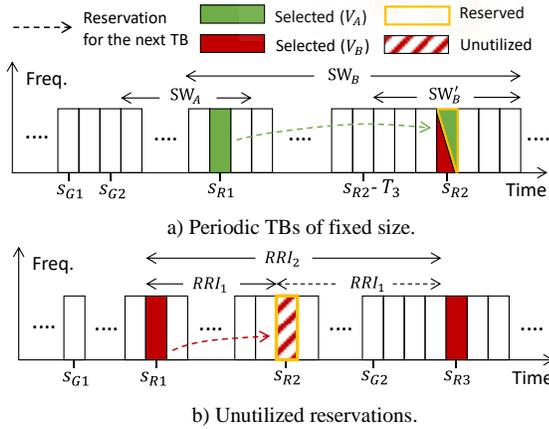


Fig. 6. Re-evaluation detection when vehicles use different $RRIs$.

The fourth case where re-evaluation detection successfully detects a collision occurs when the two RRI values are multiples of each other (e.g., $RRI_2 = 2 \cdot RRI_1$) 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.

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. Then, V_B triggers 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 collision detected at slot s_{R3} by the re-evaluation mechanism involves the resources reserved by V_A for its retransmission and the resources selected by V_B for the retransmission of its TB. Since $N_{SCI} = 2$, both the SPS and DS

schemes can accommodate the retransmission of TBs on reserved resources. As a result, the collision illustrated in Fig. 7(a) can be detected when: (i) V_A and V_B employ the SPS scheme; (ii) V_A and V_B employ the DS scheme; (iii) V_A uses the SPS and V_B uses the DS scheme, or vice versa.

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. In this case, the collision detected at slot s_{R2} by the re-evaluation mechanism involves the resource reserved by V_A for its retransmission and the resources selected by V_B for its initial transmission. Like in Fig. 7(a), also the collision illustrated in Fig. 7(b) can be detected in three different circumstances if $N_{SCI} = 2$, namely: (i) if V_A and V_B employ the SPS scheme; (ii) if V_A and V_B employ the DS scheme; (iii) if V_A uses the SPS and V_B uses the DS scheme, or vice versa.

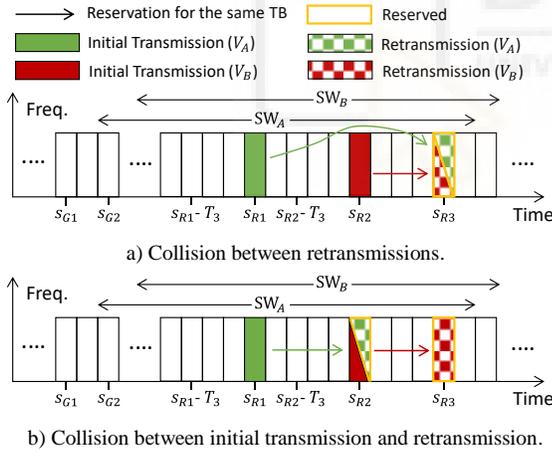


Fig. 7. Impact of retransmissions on the re-evaluation detection.

We should note that the discussion and analysis of the re-evaluation mechanism presented in this Section exclusively relies on the notions of selected and reserved resources, and it does not depend on the scheduling scheme employed by the vehicles. A selected resource is a resource that a vehicle selects during a resource reselection to transmit a TB. A reserved resource is a selected resource that the vehicle reserves for a future transmission by notifying neighbouring vehicles using the 1st-stage SCI.

When a single transmission per TB is considered ($N=1$), only vehicles employing the SPS scheme can transmit on reserved resources, since the DS scheme forces the selection of new resources for every TB. On the other hand, both the DS and the SPS scheme allow vehicles to accommodate their transmissions

on selected resources. Therefore, the re-evaluation mechanism can identify collisions (between selected and reserved resources) in two different cases if $N=1$: (i) when all vehicles utilize the SPS scheme; (ii) when the SPS and DS scheme coexist. If all vehicles use the DS scheme, the re-evaluation mechanism is not able to detect and avoid any collision, since collisions only occur between selected resources, and collisions between selected resources cannot be detected by re-evaluations.

If a TB is transmitted more than once ($N=2$ in this study), and the 1st-stage SCI associated with the TB's initial transmission can reserve the resources used for the retransmission of the same TB (i.e., $N_{SCI} = 2$), then both the SPS and DS schemes can accommodate the retransmission of TBs on reserved resources. In this case, the re-evaluation mechanism can identify collisions (between selected and reserved resources) in three different circumstances: (i) when all vehicles utilize the SPS scheme; (ii) when all vehicles utilize the DS scheme; (iii) when the SPS and DS scheme coexist.

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 reselection 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 standard-compliant 5G NR V2X mode 2 simulator⁶ implemented by the authors in ns-3. The implementation of our simulator adheres to 3GPP MAC and PHY layer specifications introduced in Release 16 [6][14], and follows the 5G NR V2X mode 2 evaluation guidelines defined by 3GPP in [2]. 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 [15]. 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 assume that each TB is transmitted using 16QAM and a coding rate equal to 0.5. In all simulations, we consider broadcast transmissions. 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 [16] for the transmission of TBs and the ones from [17] 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 [13]:

- 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. Table I summarizes the key parameters used in the simulations.

TABLE I. KEY SIMULATION PARAMETERS

Parameter	Values evaluated
Channel bandwidth	20 MHz
Subcarrier spacing	30 kHz
Sub-channels per slot	4
Transmission power	23 dBm
Modulation and coding scheme	16QAM 0.5
Highway length	5 km
Number of lanes	6 (3 per direction)
Traffic density	25, 50, 100 veh/km
Size of the TB (periodic traffic)	190 bytes
Size of the TB (aperiodic traffic)	[200, 1200] bytes (200-byte step)
Avg. inter-packet arrival time	20, 100 ms
PDB	10, 20, 50, 100 ms
Processing delay time T_3	5 slots
Transmissions per TB (N)	1, 2
RRI	10, 20, 50, 100 ms

We define the following evaluation metrics:

- Packet Delivery Ratio (PDR) [2]: fraction of correctly received TBs over the total number of transmitted TBs. In the

⁶ The simulator is available at: <https://github.com/LLusvarghi/MoReV2X>

case of retransmissions, a TB is labeled as correctly received if at least 1 out of the N transmissions is correctly received. According to the 3GPP evaluation guidelines reported in [2], the PDR is computed relying on the notion of distance interval. The i -th distance interval is defined as the set of transmitter-receiver distances that fall within the $(a_i, b_i]$ range, $a_i = i \cdot 25$ m and $b_i = (i + 1) \cdot 25$ m. For the i -th interval, the PDR is computed as:

$$PDR = \frac{\sum_{j=1}^M X_i^j}{\sum_{j=1}^M Y_i^j} \quad (1)$$

where X_i^j indicates the number of vehicles within the i -th interval that correctly decoded the j -th TB, Y_i^j represents the number of vehicles within the i -th interval when the j -th TB was transmitted, and M denotes the total number of transmitted TBs.

- **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.
- **Half-Duplex Losses Ratio (HDLR):** fraction of TBs that are incorrectly received because of the half-duplex limitation over the total number of transmitted TBs. This error occurs when the TB cannot be received because the receiver was transmitting in the same slot. The HDLR is computed per distance interval.
- **Propagation Losses Ratio (PLR):** fraction of TBs that cannot be correctly decoded because the received power level is below the sensitivity level or the Signal to Noise Ratio (SNR) is not sufficiently high over the total number of transmitted TBs. Propagation errors exclude half-duplex errors. The PLR metric is also computed per distance interval.
- **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. Collision errors exclude propagation and half-duplex errors. The PCR metric is also computed per distance interval.
- **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.
- **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 II.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 II.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 II.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 II 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 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

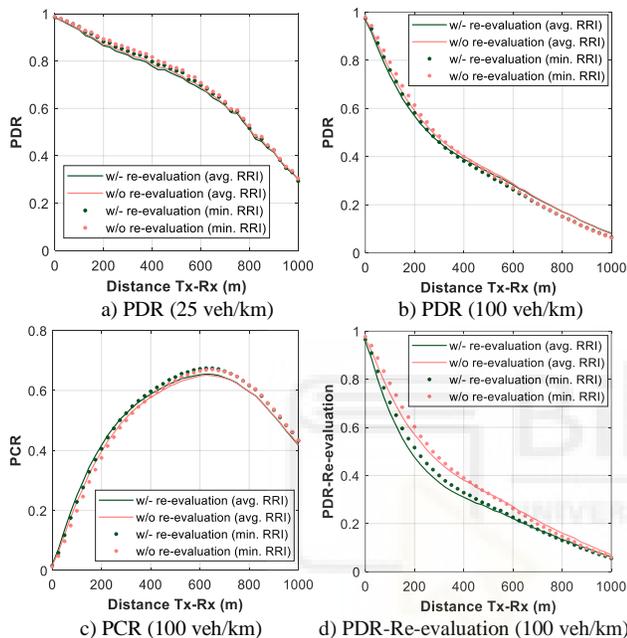
TABLE II. PERFORMANCE METRICS (IN %) OF SPS WHEN $N=1$

A) APERIODIC TRAFFIC OF VARIABLE SIZE AND MIXED TRAFFIC SCENARIO

<i>RRI</i> strategy	ReCR	SRR	LRR	URR	25 veh/km		50 veh/km		100 veh/km	
					ReDR	IReDR	ReDR	IReDR	ReDR	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

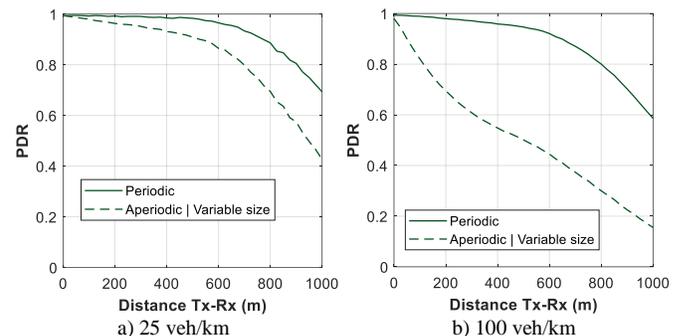
Scenario	ReCR	25 veh/km		50 veh/km		100 veh/km	
		ReDR	IReDR	ReDR	IReDR	ReDR	IReDR
Single traffic	10.3	0.006	0	0.01	0	0.03	0
Mixed traffic	5.5	0.1	0	0.2	0	0.5	0

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

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 II.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 windows⁷. 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 II.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 II.b and the small impact of re-evaluation on the PCR in Fig. 10(a), where the impact of re-evaluation on SPS is reported in the mixed traffic scenario for periodic traffic. 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

⁷ Using [18], we can estimate that around 30% of packets that trigger a resource reselection would experience persistent collisions with 100 veh/km.

not implemented. In this case, re-evaluations were effective to avoid the limited set of packet collisions detected with periodic traffic of fixed size.

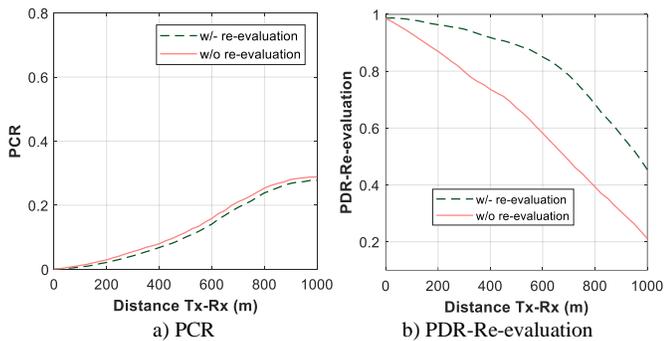


Fig. 10. SPS performance in mixed traffic scenario for periodic traffic, 100 veh/km, $N = 1$. Similar trends are observed under the single traffic scenario.

We should note that the differences observed in Fig. 8 and Fig. 10 when comparing the PCR, PDR, and PDR-Re-evaluation with and without re-evaluation are exclusively due to the impact of the re-evaluation mechanism⁸. This is the case because the other types of errors (half-duplex and propagation errors) do not depend on the re-evaluation mechanism as visible in Fig. 11. The figure reports the HDLR (Fig. 11(a)) and PLR (Fig. 11(b)) metrics as a function of the transmitter-receiver distance under the same conditions as Fig. 8(c)⁹. Fig. 11 clearly shows that the same HDLR and PLR performance is experienced whether re-evaluation is used or not. On the other hand, re-evaluation impacts the probability of packet collision, and hence the PCR and PDR. As a result, only the re-evaluation mechanism is responsible for the differences observed when comparing the performance of 5G NR V2X mode 2 with and without re-evaluation.

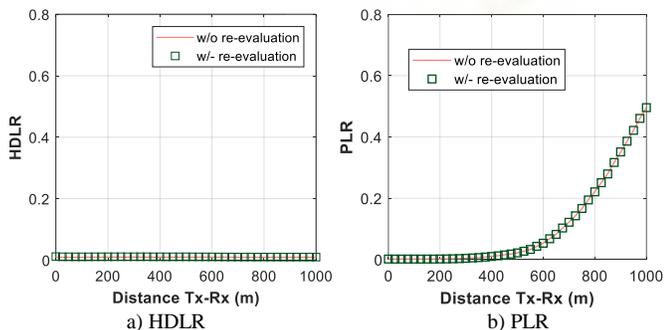


Fig. 11. HDLR (a) and PLR (b) for aperiodic traffic of variable size in the mixed traffic scenario, 100 veh/km, $N = 1$.

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 III 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 II.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, the PDR for the packets that detected a re-evaluation (PDR-Re-evaluation) significantly outperforms the PDR without re-evaluation (Fig. 12(a)); this was not the case without retransmissions ($N=1$) as shown in Fig. 8(d). Fig. 12(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. 12(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 III) as well as the lower positive impact of re-evaluation in Fig. 12(b) compared to Fig. 12(a).

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

⁸ We should note that the comparison with and without re-evaluation is always done considering the same RRI selection strategy, number of retransmissions, vehicular density, traffic type, and scheduling scheme.

⁹ Fig. 11(b) shows that, as expected, the PLR increases with the distance since the higher the distance the lower the received power levels. This trend

explains the shape of the PCR curve in Fig. 8(c) given that collision errors exclude propagation and half-duplex errors. As a result, the higher the PLR, the higher the number of TBs that are excluded in the PCR metric. In this case, the PCR starts decreasing from the distance at which propagation errors become the dominant source of errors.

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. 13(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.

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%).

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

Scenario	ReCR	SRRL	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

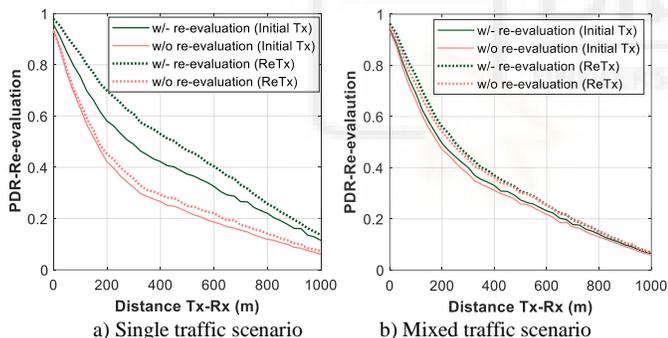


Fig. 12. PDR-Re-evaluation experienced by SPS for aperiodic traffic of variable size when $N = 2$ (50 veh/km, average RRI strategy). Similar trends have been obtained for other densities.

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.

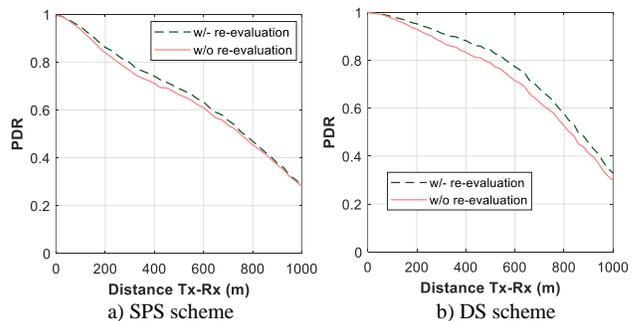


Fig. 13. PDR experienced by mode 2 for aperiodic traffic of variable size when $N = 2$, 50 veh/km (average RRI strategy for SPS scheme) in the single traffic scenario. Similar trends have been obtained for other densities.

Table IV reports the ratio of re-evaluation checks (ReCR) and detections (ReDR) that characterize the DS in the different settings considered in this work. Table IV 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 and Table III. 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 IV 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 IV). Such ReDR reduction is more evident at larger densities.

Fig. 14 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. 14(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. 14(b) which corresponds to the highest vehicular 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. Fig. 14 shows that the vehicular density has an impact on the operation and effectiveness of the re-evaluation mechanism, hence affecting the system performance. For low vehicular densities, the re-evaluation mechanism is able to select new collision-free resources during the resource replacement phase, therefore

avoiding the detected collisions and reducing the total number of collisions. As the vehicular density increases, the number of detected collisions augments, and so does the number of resource reselections. This reduces the probability of selecting collision-free resources after a re-evaluation and deteriorates the effectiveness of the re-evaluation mechanism.

TABLE IV. 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

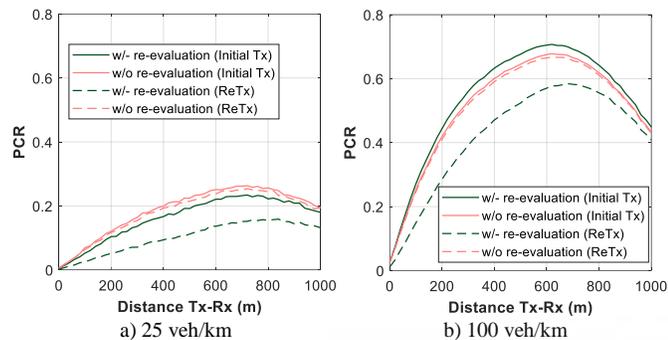


Fig. 14. PCR experienced by DS with aperiodic traffic of variable size in the single traffic scenario when $N=2$.

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. 13(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 IV, 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. 14, 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.

VIII. CONCLUSIONS

This paper has presented a comprehensive analysis and evaluation of the impact of the re-evaluation mechanism on the operation and performance of NR V2X mode 2 sidelink communications. The re-evaluation mechanism has been introduced in 3GPP Release 16 standards to reduce packet collisions. This study shows that the effectiveness of re-evaluation to avoid collisions depends on the data traffic patterns and mode 2 configurations. In particular, the study shows that re-evaluation is effective in detecting collisions when vehicles transmit periodic traffic of fixed size. However, the impact on the performance of NR V2X mode 2 is small

since the number of packet collisions detected by re-evaluation is low under periodic traffic of fixed size. The effectiveness of re-evaluation can decrease under the presence of aperiodic traffic of variable size because traffic variability increases the probability of selecting new resources, and re-evaluation cannot detect collisions on new selected resources. This is particularly the case when there are no retransmissions. Without retransmissions, re-evaluation can only detect collisions with the SPS scheduling scheme. Vehicles using DS select new resources for every TB, and re-evaluation cannot detect packet collisions on new selected resources. With retransmissions, re-evaluation can detect collisions for both SPS and DS scheduling schemes since the retransmissions always take place on reserved resources. Our study shows that re-evaluation is more effective in detecting packet collisions with retransmissions, even with aperiodic traffic of variable size. However, the impact of re-evaluation on the performance of SPS and DS with retransmissions is low since, without re-evaluation, a TB is correctly received if just one of the two transmissions is correctly received.

We performed additional simulations to explore mixed scenarios in which some vehicles employ SPS while others use DS. The outcomes and trends observed in these mixed scenarios regarding the effectiveness of the re-evaluation mechanism closely align with those discussed in Section III and quantitatively analyzed for SPS and DS. In fact, our comprehensive analysis in Section III primarily focuses on the concepts of selected and reserved resources, and this remains independent of the scheduling scheme employed by the vehicles.

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, we should note that this study has demonstrated that re-evaluation does not ultimately provide significant benefits for NR V2X mode 2, and is not that effective in avoiding packet collisions. The implementation of re-evaluation (currently mandatory according to 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 NR V2X mode 2 sidelink communications.

REFERENCES

- [1] M. H. C. Garcia *et al.*, "A Tutorial on 5G NR V2X Communications," in *IEEE Communications Surveys & Tutorials*, vol. 23, no. 3, pp. 1972-2026, thirdquarter 2021.
- [2] 3GPP TR 37.885, "Study on evaluation methodology of new Vehicle-to-Everything (V2X) use cases for LTE and NR," Release 15 V15.3.0, June 2019.
- [3] L. Lusvarghi and M. L. Merani, "On the Coexistence of Aperiodic and Periodic Traffic in Cellular Vehicle-to-Everything," in *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," in *IEEE Access*, vol. 8, pp. 121526-121548, 2020.
- [5] L. Lusvarghi and M. L. Merani, "Machine Learning for Disseminating Cooperative Awareness Messages in Cellular V2V Communications," in *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," Release 16 V16.7.0, Jan. 2022.

- [7] Z. Ali, S. Lagén, L. Giupponi and R. Rouil, "3GPP NR V2X Mode 2: Overview, Models and System-Level Evaluation," in *IEEE Access*, vol. 9, pp. 89554-89579, 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," in *IEEE Access*, vol. 9, pp. 145648-145661, 2021.
- [9] C. Campolo, V. Todisco, A. Molinaro, A. Berthet, S. Bartoletti and A. Bazzi, "Improving Resource Allocation for beyond 5G V2X Sidelink Connectivity," *2021 55th Asilomar Conference on Signals, Systems, and Computers*, Pacific Grove, CA, USA, 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," *2022 International Conference on Artificial Intelligence in Information and Communication (ICAIIIC)*, Jeju Island, Korea, Republic of, 2022, pp. 171-175.
- [11] G. Thandavarayan, M. Sepulcre and J. Gozalvez, "Analysis of Message Generation Rules for Collective Perception in Connected and Automated Driving," *2019 IEEE Intelligent Vehicles Symposium (IV)*, Paris, France, 2019, pp. 134-139.
- [12] A. Correa *et al.*, "Infrastructure Support for Cooperative Maneuvers in Connected and Automated Driving," *2019 IEEE Intelligent Vehicles Symposium (IV)*, Paris, France, 2019, pp. 20-25.
- [13] A. Molina-Galan, B. Coll-Perales, L. Lusvarghi, J. Gozalvez and M. L. Merani, "How does 5G NR V2X Mode 2 Handle Aperiodic Packets and Variable Packet Sizes?," *2022 IEEE 23rd International Conference on High Performance Switching and Routing (HPSR)*, Taicang, Jiangsu, China, 2022, pp. 183-188.
- [14] 3GPP TS 38.214, "NR; Physical layer procedure for data," Release 16 V16.8.0, Jan. 2022.
- [15] 5GAA, "V2X Functional and Performance Test Report; Test Procedures and Results," 5GAA Report, Apr. 2019.
- [16] R1-1900852, "Link level evaluations on sidelink for NR V2X," Huawei, HiSilicon, 3GPP TSG RAN WG1 Ad-Hoc Meeting 1901, Taipei, Jan. 2019.
- [17] R1-1903180, "Link level evaluations of NR PSCCH," Ericsson, 3GPP TSG RAN WG1 Meeting #96, Athens, Greece, March 2019.
- [18] M. Gonzalez-Martín, M. Sepulcre, R. Molina-Masegosa and J. Gozalvez, "Analytical Models of the Performance of C-V2X Mode 4 Vehicular Communications," in *IEEE Transactions on Vehicular Technology*, vol. 68, no. 2, pp. 1155-1166, Feb. 2019.

Annex A.4 Publication



Re-Evaluation Strategies for 5G NR V2X Communications

Alejandro Molina-Galan, Baldomero Coll-Perales and Javier Gozalvez
UWICORE laboratory, Universidad Miguel Hernández de Elche, Elche (Alicante), Spain.
{alejandro.molinag, bcoll, j.gozalvez}@umh.es

Abstract—5G NR V2X introduces novel MAC features to support advanced V2X services for connected and automated driving. One of these features is the re-evaluation mechanism designed to detect and avoid packet collisions. The re-evaluation mechanism checks before every transmission if the selected resources are still available. If another vehicle has selected the same resources, the re-evaluation mechanism triggers a reselection of resources. The 3GPP standard defines different strategies for vehicles to decide when to perform a re-evaluation check, and how many times a re-evaluation check should be executed before the transmission of a packet. This study evaluates for the first time the impact of two standard strategies on the operation and performance of NR V2X mode 2. The evaluation demonstrates that the two strategies offer a balance between transmission latency and computational cost. We then propose an alternative re-evaluation check strategy that can simultaneously reduce the transmission latency and computational cost without actually degrading the reliability of NR V2X mode 2 transmissions.

Keywords—5G NR V2X, re-evaluation, CAV, C-V2X, cellular V2X, 5G, Mode 2, LTE V2X, 3GPP, latency.

I. INTRODUCTION

3GPP published in Release 16 the 5G NR V2X (or NR V2X) standard to support advanced V2X (Vehicle to Everything) services. This is the first standard based on the 5G NR (New Radio) radio interface that supports sidelink (SL) or direct Vehicle-to-Vehicle (V2V) communications [1]. NR V2X is designed to complement LTE V2X that supports basic awareness, safety and traffic services. NR V2X is expected to support more advanced V2X services with stringent requirements (e.g. in terms of reliability and latency) owed to the capabilities of 5G NR and new features introduced in NR V2X at the MAC and PHY layers. NR V2X defines two modes of operation. In mode 1, the cellular infrastructure selects and manages the radio resources for direct or SL V2V communications. In mode 2, vehicles autonomously select and manage the resources without the support of the cellular infrastructure. This paper focuses on NR V2X mode 2 since it provides vehicles the capacity to communicate even without cellular coverage.

NR V2X introduces several new features compared to LTE V2X. One of the most relevant ones at the MAC level is the re-evaluation mechanism. This mechanism is introduced to detect and avoid possible packet collisions. In NR V2X mode 2, vehicles autonomously select resources to transmit several consecutive packets. The re-evaluation mechanism checks for every transmission whether the selected resources are still available before the vehicle transmits its next packet. If the vehicle detects that another vehicle has selected the same resources, the re-evaluation mechanism triggers a reselection of resources to avoid a possible packet collision. The 3GPP standard defines different strategies for vehicles to decide when to perform the re-evaluation check and how many times before the transmission of a packet [2]. One strategy only executes re-evaluation once, and it does so just prior to the

transmission of a packet. Other strategies execute multiple times the re-evaluation mechanism. The standard leaves up to UE implementation which re-evaluation check strategy to utilize. To the authors' knowledge, no study has yet evaluated the impact of these strategies on the operation and performance of NR V2X mode 2. In fact, most of the NR V2X mode 2 studies to date do not consider the re-evaluation mechanism even though the standard specifies that its use is mandatory [2]. The study presented by the authors in [3] is one of the first evaluating NR V2X mode 2 under different traffic patterns, including messages generated aperiodically and with variable packet sizes following 3GPP guidelines [4]. However, the study only considered a single strategy for the execution of the re-evaluation mechanism. In this context, this paper advances the current state-of-the-art with the first evaluation of two different re-evaluation check strategies proposed by the 3GPP standard for NR V2X mode 2. The evaluation is conducted following 3GPP evaluation guidelines and using a simulator developed by the authors that faithfully implements NR V2X mode 2 in ns3. The evaluation demonstrates that the two strategies challenges NR V2X mode 2 since they offer a balance between transmission latency and computational cost of the re-evaluation mechanism. The study proposes then an alternative re-evaluation check strategy that seeks to minimize both the latency and the computational cost. We demonstrate that the proposed re-evaluation check strategy can simultaneously reduce the transmission latency and computational cost without actually degrading the reliability of NR V2X mode 2 transmissions.

II. 5G NR V2X MODE 2

A. Resource grid

NR V2X defines a resource grid structured into Resource Blocks (RBs) and slots in the frequency and time domain, respectively (Fig. 1). The OFDM numerology (μ) or subcarrier spacing (SCS) determines the RB bandwidth and the slot duration. NR V2X supports a flexible SCS equal to $2^\mu \times 15$ kHz, where μ can take values equal to 0, 1, 2 or 3. The RB bandwidth is equal to the space occupied by 12 consecutive subcarriers with the same SCS, and the slot duration is equal to $2^{-\mu}$ ms. The same SCS is utilized by all vehicles within a specific region. Several RBs of the same slot are grouped into a sub-channel, which is the smallest unit for a SL data transmission or reception (Fig. 1). The number of RBs per sub-channel (or sub-channel size) is configurable, but all communicating vehicles use the same one. Data packets are transmitted in Transport Blocks (TB). A TB can be transmitted in one or more sub-channels depending on the packet size, the Modulation and Coding Scheme (MCS) and the sub-channel size. In NR V2X, a Sidelink Control Information (SCI) is associated to each TB. The SCI and the associated TB are transmitted in the same slot. The SCI indicates the resources utilized by the associated TB and more information necessary for decoding the TB. The SCI can also announce the reservation of future resources.

This 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), by Generalitat Valenciana (CIGE/2021/096) and by UMH's Vice Rector for Research grants.

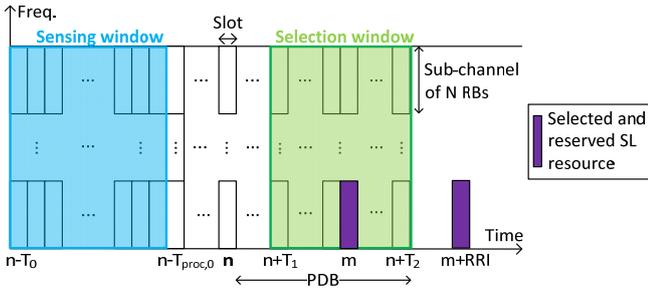


Fig. 1. NR V2X resource grid and mode 2 SPS (with $T_2 = PDB$).

B. Resource allocation

NR V2X mode 2 defines two resource allocation schemes: a semi-persistent scheduling scheme (SPS) and a dynamic scheme (DS) [2]. With SPS, a vehicle selects and reserves resources for several consecutive TB transmissions. With DS, a vehicle only selects resources for the current TB, and new resources must be selected for each TB. This study focuses on SPS since DS does not utilize the re-evaluation mechanism.

With SPS, vehicles select SL resources for consecutive *Reselection Counter* TBs. The vehicle uses the SCI associated with a TB to reserve the resources for a future TB transmission. It is important to differentiate between a selected and a reserved resource. A selected resource is labeled as reserved when it is announced by a SCI. After transmitting a TB, the *Reselection Counter* is decremented by one. The vehicle selects new resources with probability $(1-P)$ when the *Reselection Counter* is equal to zero. A vehicle must also select new resources if a new TB does not fit in the selected/reserved resources or if the new TB needs to be transmitted before the selected/reserved resources. The *Resource Reservation Interval (RRI)* defines the time period between the resources selected for consecutive TBs.

A vehicle that needs to select new SL resources first configures the selection window to identify candidate resources to transmit a TB. All resources that are within the range of slots $[n+T_1, n+T_2]$ (Fig. 1) are included in the selection window. n is the slot at which the packet is generated and the vehicle selects new resources. T_1 represents the time (in slots) needed by a vehicle for identifying candidate resources and selecting new resources. The vehicle must configure T_2 within the range $T_{2min} \leq T_2 \leq PDB$. PDB (Packet Delay Budget) is the latency deadline to transmit a TB. After defining the selection window, the vehicle identifies the candidate resources within the selection window. A candidate resource is defined by a slot and N_{SCH} contiguous sub-channels. The vehicle must select N_{SCH} so that the new TB and its associated SCI fit in the candidate resource. The vehicle also defines the sensing window in the range of slots $[n-T_0, n-T_{proc,0})$ to sense the transmissions from neighboring vehicles when it is not transmitting (Fig. 1)[5]. During the sensing window, the vehicle receives the SCIs transmitted by other vehicles and that are used to announce the resources they are reserving for their future transmissions. The vehicle stores the information about reservations from other vehicles, along with the measured Reference Signal Received Power (RSRP) of the TBs associated with the received SCIs.

Mode 2 defines a 2-step algorithm to select new resources. First, the vehicle excludes in step 1 candidate resources from the selection window that have been reserved by other vehicles using their SCIs during the sensing window. The candidate resources are only excluded if the measured RSRP is higher than a threshold. After step 1, the vehicle verifies the

percentage of remaining available candidate resources in the selection window. If this percentage is not equal or higher than a threshold (20%, 35% or 50% [5]), the vehicle increases by 3 dB the RSRP threshold and executes again step 1. After the percentage threshold is satisfied, the vehicle randomly selects in step 2 one of the available candidate resources.

C. Re-evaluation mechanism

A vehicle can experience a packet collision in SPS if other vehicles select the same resources between the time the vehicle selected its resources and the time it transmits a TB. NR V2X mode 2 introduces the re-evaluation mechanism to avoid these possible collisions. To do so, the vehicle must check if the selected resources remain available before transmitting the TB by executing again step 1. This new execution of step 1 before transmitting a TB is defined as *re-evaluation check* in 3GPP standards [2].

Fig. 2 illustrates the re-evaluation process in NR V2X mode 2. The figure represents a scenario where a vehicle initially selects a resource at slot m , and executes a re-evaluation check at slot n' . In this case, the vehicle defines a new selection window SW' within the range of slots $[n'+T_1, n'+T_2]$, where T_2' must take values within $T_{2min} \leq T_2' \leq PDB - (n'-n)$. We consider that T_2' is always chosen so that the upper limits of SW and SW' coincide. The vehicle executes step 1 at slot n' over SW' to check which resources are currently available. The 3GPP standard indicates that a vehicle *detects* a re-evaluation if the initially selected resource at slot m is not available anymore when executing again step 1 at slot n' [5]. In this case, the vehicle executes step 2 after the *re-evaluation detection* to select new resources among the available resources in SW' [2] (new selected resources at m' in Fig. 2). The new execution of step 2 as part of the re-evaluation mechanism is referred to as *resource replacement* in 3GPP standards [2]. It is also important to highlight that the resources selected after a resource replacement (at slot m' in Fig. 2) can be located earlier or later than the slot at which the initially selected resources are located (at slot m in Fig. 2). This is the case because SW and SW' overlap, and resources are selected randomly after a resource replacement. The likelihood that a TB is transmitted earlier or later compared to the initially selected resources after a resource replacement depends on the size of SW' , and hence on when the re-evaluation check is performed.

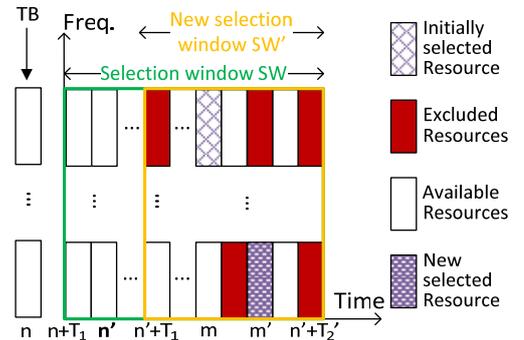


Fig. 2. Re-evaluation mechanism in NR V2X mode 2.

The 3GPP standard defines that vehicles can execute the re-evaluation check process over selected SL resources in two cases. The first case is when a vehicle selects new SL resources to transmit a new TB. In this case, vehicles are forced to execute the re-evaluation check process [2]. The second case is when a vehicle leaves unutilized a reserved resource at time t , and it utilizes the next selected resource (at

$t+RRI$) to transmit the TB; we should note that the selected resource at $t+RRI$ was not reserved since there was no transmission announcing the reservation at t . In this case, the standard defines that the decision to execute the re-evaluation check process over these selected resources is up to UE implementation. We consider in this study that vehicles execute the re-evaluation check in both cases since unutilized reservations can generate packet collisions as shown in [3].

III. RE-EVALUATION CHECK STRATEGIES

The re-evaluation check must be executed before the slot where the initially selected resources are located (i.e. before slot m in Fig. 3). In particular, the standard indicates that the re-evaluation check can be executed at any slot from the slot n at which the packet is generated until the slot $m-T_3$ [2]. T_3 is the maximum time permitted for a vehicle (in slots) to perform the process to select resources. The re-evaluation check can then be performed at any slot in the range $[n+1, m-T_3]$ in Fig. 3. This range defines the different re-evaluation check strategies allowed by 3GPP standard. We evaluate in this paper the two strategies analyzed by 3GPP for 5G NR V2X. In the first one, the vehicle only executes once a re-evaluation check at $m-T_3$. In the second strategy, the vehicle executes a re-evaluation check in all slots in the range $[n+1, m-T_3]$. The standard leaves up to UE implementation which re-evaluation check strategy to utilize. In the rest of the paper, we refer to these strategies as one-slot and all-slots, respectively.

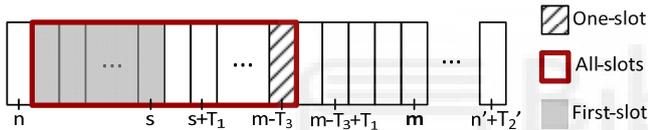


Fig. 3. Illustration of re-evaluation check strategies (when T_3 and T_1 are equal to 5 and 2 slots, respectively).

A. One-slot strategy

The one-slot strategy executes a single re-evaluation check before the initially selected resource. The check is executed at slot $n'=m-T_3$ (Fig. 3), i.e. at the last possible moment within the range $[n+1, m-T_3]$. Delaying the re-evaluation check until $m-T_3$ increases the likelihood that the new selected resource after a resource replacement is located at a slot later than the initially selected resource, and this delays the transmission of the TB. This likelihood increases because the new selected resource must be selected in SW' that is defined by the range of slots $[m-T_3+T_1, m-T_3+T_2]$. Let us denote as m' the slot where the new selected resource is located. The new selected resource would anticipate the transmission of the TB with respect to the initially selected resource only if m' falls within the range $[m-T_3+T_1, m]$. This range only includes T_3-T_1 slots, which is a small number of slots. T_1 is defined in the range $[0, T_3]$ slots. For example, if $T_1=2$ slots and $T_3=5$ slots (with $SCS=30$ kHz), T_3-T_1 is equal to 3 slots (or 1.5 ms when $SCS=30$ kHz).

If a vehicle selects a new resource at slot m' using the one-slot strategy, the standard forces the vehicle to execute a new re-evaluation check at slot $m'-T_3$. This process is repeated until it is not possible to define a new selection window (e.g., because $T_{2min}>T_2$), or when the slot at which the vehicle should execute the re-evaluation has passed.

B. All-slots strategy

The standard also allows for the possibility to conduct re-evaluation checks in all slots in the range $[n+1, m-T_3]$. We

refer to this strategy as all-slots (Fig. 3). This strategy increases the computational cost compared to the one-slot strategy since it significantly augments the number of re-evaluation checks per TB. However, the all-slots strategy has the potential to reduce the transmission latency compared to the one-slot strategy. This is the case because it is more likely with the all-slots strategy that a new selected resource (in the event of a resource replacement) anticipates the transmission of the TB with respect to the initially selected resource. Fig. 3 shows that the all-slots strategy performs re-evaluation checks from the slot $n+1$ to the slot $m-T_3$. Let us assume that a re-evaluation check at slot $n'=s$ within the range $[n+1, m-T_3]$ detects a potential collision in the initially selected resource allocated at the slot m . When a resource replacement process is triggered, a new selection window SW' is defined in the range of slots $[s+T_1, s+T_2]$ to select new resources. If s is any slot before $m-T_3$ (as represented in Fig. 3), the new selection window SW' defined by the all-slots strategy starts earlier than the one defined with the one-slot strategy since $(s+T_1) < (m-T_3+T_1)$. This increases the likelihood that the new selected resources by the all-slots strategy fall within the range of slots that advances the TB transmission. This is the case because the size of the range of slots that advances the TB transmission is higher in all-slots ($[s+T_1, m]$) than the one in one-slot ($[m-T_3+T_1, m]$). As a result, the all-slots strategy has the potential to anticipate the transmission of a TB compared to the one-slot strategy. However, this comes at the expense of a higher computational cost since re-evaluation checks must be done for all slots within $[n+1, m-T_3]$. A vehicle that has selected a new resource at slot m' with the all-slots strategy continues performing re-evaluation checks in every slot until the slot $m'-T_3$. This process is repeated iteratively until it is not possible to define a new selection window (e.g., because $T_{2min}>T_2$), or when the slot at which the vehicle should execute the re-evaluation has passed.

C. First-slot strategy

The one-slot and all-slots re-evaluation check strategies are defined in 3GPP standards, and it is up to UE implementation to decide to use one of them [2]. The all-slots strategy has the potential to reduce the transmission latency compared to the one-slot strategy, but this is achieved at the expense of a significantly higher computational cost for the re-evaluation mechanism. In this paper, we propose a new re-evaluation check strategy designed with the objective to reduce the transmission latency while controlling the computational cost of the re-evaluation mechanism. We refer to this new proposal as the first-slot strategy.

A vehicle implementing the first-slot strategy performs re-evaluation checks sequentially in all slots in the range $[n+1, m-T_3]$ until it detects the first re-evaluation. Fig. 3 shows an example in which a vehicle implementing the first-slot strategy performs re-evaluation checks from the slot $n+1$ until the slot $n'=s$ where it detects a re-evaluation. It then executes the resource replacement process and selects new resources at slot m' . After this resource replacement, the vehicle does not perform any additional re-evaluation checks when implementing the first-slot strategy. The objective is to increase the likelihood of anticipating the transmission of a TB compared to the one-slot strategy, while reducing the computational cost compared to the all-slots strategy. The likelihood to anticipate the transmission of a TB (i.e., $m'<m$) with the first-slot strategy depends on the slot at which the re-evaluation is detected. The earlier it is detected, the higher the likelihood that $m'<m$. This is the case because the size of the

range $([s+T_l, m])$ that anticipates the TB transmission increases. Compared to the one-slot and all-slots strategies, the first-slot strategy does not execute a re-evaluation check T_3 slots before the selected resources if the vehicle detects a re-evaluation in a previous slot.

IV. SIMULATION ENVIRONMENT

We have evaluated the performance of the three different re-evaluation check strategies using an NR V2X mode 2 simulation platform implemented by the authors in ns3 following the 3GPP's standards [5][2] and guidelines [4]. The evaluation scenario is a 5 Km-long highway with 3 lanes in each driving direction where vehicles drive at 70 km/h [4]. We evaluate the performance under vehicle densities of 50 veh/km and 100 veh/km. Following the traffic models reported by 3GPP in [4], we consider that vehicles transmit packets of variable size aperiodically. The packet size follows a uniform distribution in the range [200,1200] bytes with a 200-byte step. The packet generation follows the aperiodic models that 3GPP defines for medium and high data traffic intensity scenarios. The minimum and average inter-packet arrival time for the medium traffic intensity scenario are equal to 50 ms and 100 ms, respectively. These values are equal to 10 ms and 20 ms, respectively, for the high traffic intensity scenario. For both scenarios, the latency requirement or PDB of each generated packet is set equal to the minimum inter-packet arrival time. For the sake of simplicity, we refer to these scenarios as high and medium intensity when the PDB is 10 ms and 50 ms, respectively.

NR V2X mode 2 is configured to operate at 5.9 GHz with a channel bandwidth of 20 MHz. We have set the subcarrier spacing to 30 kHz. With this setting, vehicles can utilize up to 4 sub-channels (of 12 RBs each) in each slot to transmit the generated packets. The sensitivity is configured to -103.5 dBm following [6], and the transmission power is set to 23 dBm. The pathloss and shadowing models follow reference 3GPP models reported in [4]. The shadowing model includes spatial shadowing correlation and follows a log-normal distribution with mean and standard deviation equal to 0 dB and 3 dB, respectively [4]. Vehicles transmit the TBs using the MCS index 13 (with 16QAM modulation and coding rate equal to 0.5). With this MCS, vehicles utilize 1, 2, 3, 3, 4 and 4 sub-channels to transmit packets of 200, 400, 600, 800, 1000 and 1200 bytes, respectively, and their associated SCI. The transmissions of the TB [7] and SCI [8] are modeled with Block Error Rate (BLER) – Signal to Interference plus Noise Ratio (SINR) lookup tables from 3GPP working documents.

At the MAC layer, the RRI is set to the minimum inter-packet arrival time following the results obtained in [3]. The probability P to select the same resources when *Reselection Counter* is depleted is configured equal to 0. The percentage threshold of resources that must be available after the vehicle executes the step 1 of mode 2 is set to its minimum value (i.e., 20%) since this value minimizes the executions of step 1. The RSRP threshold that determines whether a vehicle must exclude a resource or not with step 1 is configured to its minimum value (i.e., -128 dBm) following the results reported in [9]. The parameters T_1 , T_2 and T_{2min} configure the limits of the selection window. They are set equal to 2 slots, PDB (20 and 100 slots for 10 ms and 50 ms) and 2 slots, respectively. $T_{proc,0}$ and T_0 configure the limits of the sensing window, and they are set equal to 1 slot and 2200 slots, respectively. T_3 defines the last slot where a re-evaluation check can be executed in the range $[n+1, m-T_3]$, and it is set equal to 5 slots following [5].

V. RESULTS

We utilize the following metrics to evaluate the three re-evaluation check strategies:

- End-to-end latency: time elapsed between the moment at which vehicles generate a packet and the moment at which other vehicles receive it. Since the objective is to evaluate and compare the different re-evaluation check strategies, this metric is computed for the packets that experience at least one re-evaluation detection (i.e. those for which the re-evaluation check has triggered a resource replacement).
- Re-evaluation Detection Ratio (ReDR): ratio of transmitted packets that experience at least one re-evaluation detection to the total number of transmitted packets.
- Multiple Re-evaluation Detection Ratio (MReDR): ratio of transmitted packets that experience more than one re-evaluation detection to the total number of transmitted packets.
- Packet Delivery Ratio (PDR): ratio of successfully received packets to the total number of transmitted packets.
- Number of Re-evaluation Checks (NReC): number of re-evaluation checks that vehicles execute per packet. This metric only considers packets that experience at least one re-evaluation detection.
- Channel Busy Ratio (CBR): ratio of sub-channels with a measured RSSI higher than a threshold computed in a window of $100 \cdot 2^{\mu}$ slots.

A. Latency

This subsection analyzes the impact of the re-evaluation check strategies on the end-to-end latency experienced by the packets that detect at least one re-evaluation. Fig. 4 shows the cumulative distribution function (cdf) of this latency for the medium and high intensity scenarios when the vehicle density is 50 veh/km. We should note that the upper-bound of the latency is 50 ms (Fig. 4.a) and 10 ms (Fig. 4.b), respectively, regardless of the re-evaluation check strategy. This is the case because these limits correspond to the length of the selection window or PDB in each scenario, and all strategies are bounded by the same length of selection window. Fig. 4 shows that the all-slots and first-slot strategies reduce considerably the latency with respect to the one-slot strategy. The one-slot strategy delays the start of the new selection window created after the re-evaluation detection on average 56 slots (or 28 ms) and 11 slots (or 5.5 ms) with respect to the slot where the packet is generated for the medium and high traffic scenarios, respectively. All-slots and first-slot only delay the start of the new selection window on average 3 slots (or 1.5 ms) and 4 slots (or 2 ms), respectively. Then, first-slot and all-slots can select resources that are located earlier than the ones selected by one-slot. This justifies the lower latencies reported in Fig. 4 for all-slots and first-slot strategies with respect to one-slot.

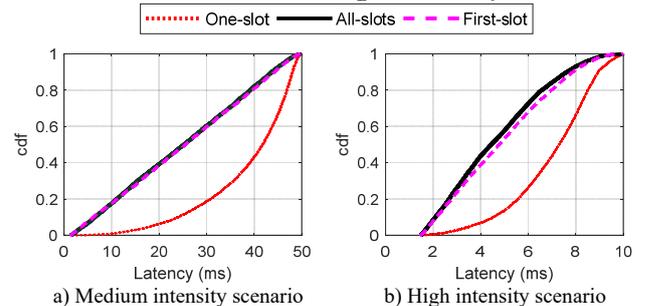


Fig. 4. CDF of the end-to-end latency for a vehicle density of 50 veh/km (similar trends observed for 100 veh/km).

Table I reports the operation and performance of the three re-evaluation check strategies under different evaluated conditions. Table I shows that the one-slot strategy increases by more than 50% and 40% the average latency compared to the all-slots and first-slot strategies in the medium and high intensity scenarios, respectively. These differences affect the packets that experience at least a re-evaluation detection, as per the definition of the latency metric. As a result, the impact of the re-evaluation check strategies on the NR V2X mode 2 latency increases with the ReDR ratio. Table I shows that the ratio of packets that experience at least a re-evaluation detection (ReDR) increases with the experienced CBR, and the CBR depends on the vehicular density and the traffic intensity.

Fig. 4 and Table I show that the all-slots and first-slot strategies achieve similar latency values. However, the all-slots strategy continues performing re-evaluation checks after the first re-evaluation detection. These additional checks result in that all-slots might detect multiple re-evaluations (MReDR in Table I), and perform multiple resource replacements. This possibility increases again with the CBR since the probability that two vehicles select the same resources augment with the CBR [3]. In any case, the small MReDR values reported for all-slots even for the high CBR scenarios justify its similar latency with the first-slot strategy.

TABLE I. OPERATION AND PERFORMANCE OF THE RE-EVALUATION CHECK STRATEGIES

PDB	Veh/km	CBR	Strategy	Avg. latency	ReDR	MReDR	NReC
50 ms	50	24 %	One-slot	38.6 ms	19 %	0 %	1
			All-slots	25.1 ms	19 %	0 %	46
			First-slot	25.4 ms	19 %	0 %	1
	100	39 %	One-slot	38.5 ms	31 %	0.3 %	1
			All-slots	24.8 ms	35 %	0.4 %	46.1
			First-slot	24.7 ms	35 %	0 %	1.1
10 ms	50	67 %	One-slot	7 ms	31 %	1.6 %	1.04
			All-slots	4.67 ms	38 %	4.2 %	6.4
			First-slot	4.9 ms	38 %	0 %	1.5
	100	88 %	One-slot	7.02 ms	30 %	1.2 %	1.04
			All-slots	4.74 ms	37 %	4.1 %	6.5
			First-slot	4.96 ms	37 %	0 %	1.6

B. Reliability and computational cost

The previous results show that all-slots and first-slot strategies significantly reduce the latency compared to the one-slot strategy. However, the one-slot strategy significantly reduces the computational cost compared to the all-slots strategy. This is visible in Table I that reports the number of re-evaluation checks that vehicles execute per packet (NReC). One-slot minimizes the NReC since it only performs a re-evaluation check T_3 slots before the selected resources are allocated. On the other hand, all-slots checks for re-evaluation in all slots, and this augments the NReC. We should note that a significantly higher NReC is observed under the scenario with high PDB. This is the case because the selection windows are larger (in number of slots) under high PDBs, and so is the number of re-evaluation checks conducted by all-slots.

All-slots and first-slot achieve similar latencies, and significantly lower than those experienced with the one-slot strategy (Fig. 4). However, Table I shows that the proposed first-slot strategy is capable to achieve such low latency values while significantly reducing the computational cost compared to the all-slots strategy. Actually, Table I shows that the first-slot strategy results in NReC values very similar to those observed with the one-slot strategy, while significantly reducing the latency. The first-slot strategy reduces by ~98% and ~75% the average number NReC of re-evaluation checks

compared to the all-slots strategy in the medium and high intensity scenarios, respectively. The reduction in computational cost is experienced also under both densities (Table I).

The first-slot strategy can achieve the same low latency values as the all-slots strategy, while reducing the computational cost to the same level as the one-slot strategy. These gains are obtained without sacrificing the reliability of NR V2X mode 2. Actually, Fig. 5.a shows that the three re-evaluation check strategies result in nearly identical PDRs for the packets transmitted after a re-evaluation detection. The resources selected after a re-evaluation detection and resource replacement might impact the transmissions of other vehicles. However, Fig. 5.b shows that three strategies result again in nearly the same PDR for all transmitted packets.

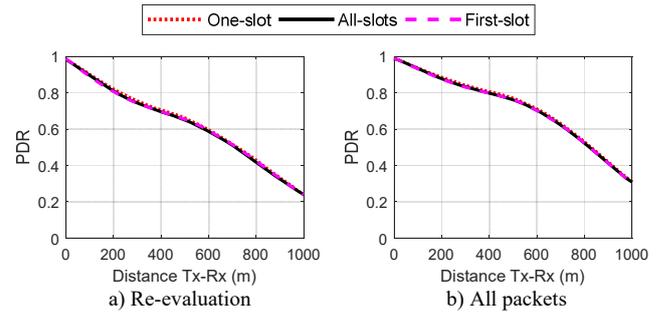


Fig. 5. PDR of packets that experienced at least a re-evaluation detection (a) and PDR experienced by all packets (b). Medium intensity with 50 veh/km (same trends observed in the other evaluated scenarios).

VI. CONCLUSIONS

This paper has presented and evaluated a set of re-evaluation check strategies for NR V2X mode 2. The re-evaluation mechanism has been introduced in NR V2X mode 2 to detect and avoid possible packet collisions. The 3GPP standard defines two strategies (one-slot and all-slots) for vehicles to decide when to perform a re-evaluation check and how many times to execute a re-evaluation check. The conducted evaluation has shown that the all-slots strategy can significantly reduce the transmission latency at the expense of a high computational cost compared to the one-slot strategy. This study presents an alternative re-evaluation check strategy (first-slot) that reduces the latency to values similar to the all-slots strategy with a low computational cost similar to the one of one-slot strategy. These benefits are achieved without degrading the reliability of the NR V2X mode 2 transmissions compared to the standardized re-evaluation check strategies.

REFERENCES

- [1] M. H. C. Garcia, et al., "A Tutorial on 5G NR V2X Communications," *IEEE Communications Surveys & Tutorials*, vol. 23, no. 3, pp. 1972-2026, Feb. 2021.
- [2] 3GPP, "TS 38.321 NR; Medium Access Control (MAC) protocol specification (v16.7.0, Release 16)", Tech. Spec., Jan. 2022.
- [3] A. Molina-Galan, et al., "How Does 5G NR V2X Mode 2 Handle Aperiodic Packets and Variable Packet Sizes?", *Proc. IEEE HPSR 2022*, pp. 183-188, Taicang, Jiangsu, China, 6-8 June 2022.
- [4] 3GPP, "TR 37.885 Study on evaluation methodology of new Vehicle-to-Everything use cases for LTE and NR", Tech. Rep., June 2019.
- [5] 3GPP, "TS 38.214 NR; Physical layer procedure for data (v16.8.0, Release 16)", Tech. Spec., Jan. 2022.
- [6] 5GAA, "V2X Functional and Performance Test Report; Test Procedures and Results", Tech. Report, Apr. 2019.
- [7] Huawei, HiSilicon, "R1-1900852. Link level evaluations on sidelink for NR V2X," 3GPP WG1 Ad-Hoc Meeting 1901, Taipei, Jan. 2019.
- [8] Ericsson, "R1-1903180. Link level evaluations of NR PSCCH," 3GPP WG1 Meeting #96, Athens, Greece, March 2019.
- [9] V. Todisco, et al., "Performance Analysis of Sidelink 5G-V2X Mode 2 Through an Open-Source Simulator", *IEEE Access*, vol. 9, pp. 145648-145661, Oct. 2021.