C^2RC: Channel Congestion-based Re-transmission Control for 3GPP-based V2X Technologies

Gaurang Naik*, Jung-Min (Jerry) Park*, Jonathan Ashdown†
*Bradley Department of Electrical and Computer Engineering, Virginia Tech, Arlington, VA, USA
†Air Force Research Laboratory, Rome, NY, USA
{gaurang, jungmin}@vt.edu, jonathan.ashdown@us.af.mil

***Abstract*—The 3rd Generation Partnership Project (3GPP) is actively designing New Radio Vehicle-to-Everything (NR V2X)—a 5G NR-based technology for V2X communications. NR V2X, along with its predecessor Cellular V2X (C-V2X), is set to enable low-latency and high-reliability communications in high-speed and dense vehicular environments. A key reliability-enhancing mechanism that is available in C-V2X and is likely to be re-used in NR V2X is packet re-transmissions. In this paper, using a systematic and extensive simulation study, we investigate the impact of this feature on the system performance of C-V2X. We show that statically configuring vehicles to always disable or enable packet re-transmissions either fails to extract the full potential of this feature or leads to performance degradation due to increased channel congestion. Motivated by this, we propose and evaluate Channel Congestion-based Re-transmission Control (C^2RC), which, based on the observed channel congestion, allows vehicles to autonomously decide whether or not to use packet re-transmissions without any role of the cellular infrastructure. Using our proposed mechanism, C-V2X-capable vehicles can boost their performance in lightly-loaded environments, while not compromising on performance in denser conditions.**

I. Introduction

The benefits of Vehicle-to-everything (V2X) communications in terms of reducing vehicle crashes and fatalities and improving traffic efficiency are well known. During the conception of Dedicated Short Range Communications (DSRC)—the first radio access technology to enable direct V2X communications—the focus was on enabling safety applications that alerted drivers in potentially hazardous situations. Today, however, connected vehicles are considered far more important not only in terms of enabling basic safety applications but also as a precursor to fully autonomous driving. A quick look at the standardization activities around V2X technologies bears testimony to the growing importance of V2X communications. The underlying standard for DSRC—IEEE 802.11p is evolving to IEEE 802.11bd, while the recently developed Cellular V2X (C-V2X) is already undergoing enhancements in New Radio V2X (NR V2X) [1].

C-V2X was developed by the 3rd Generation Partnership Project (3GPP) in its Rel. 14, while the standardization of it's successor, NR V2X, will be completed in 3GPP Rel. 16. We collectively refer to C-V2X and NR V2X as 3GPP-based V2X technologies. Direct vehicular communications in these technologies are supported over the sidelink[1] interface. 3GPP-based V2X technologies have an added advantage over IEEE 802.11-based technologies in that vehicles equipped with C-V2X or NR V2X can rely on the widespread cellular infrastructure for centralized resource allocation wherever possible. However, since the presence of cellular coverage cannot always be relied upon, both C-V2X and NR V2X support direct vehicular communications in the absence of the cellular infrastructure as well. Direct V2X communications in the absence of cellular coverage is enabled through sidelink mode 4 in C-V2X and sidelink mode 2 in NR V2X [1].

One of the fundamental challenges in the design of robust V2X systems is the design of a medium access control (MAC) protocol that is efficient, distributed and scalable. In distributed systems, the MAC protocol remains a major bottleneck in translating the high performance of the physical (PHY) layer into reliable application layer performance. Therefore, efforts are underway to improve the MAC efficiency in NR V2X. As such, NR V2X is set to introduce several new features derived from 5G NR. However, the fundamental aspects of the MAC layer design are likely to remain similar across the two 3GPP-based V2X technologies [1]. One key reliability-enhancing mechanism that is considered in both, C-V2X and NR V2X, is the use of packet re-transmissions. In the broadcast mode of communications (the default mode in C-V2X and one of the three modes in NR V2X), re-transmissions are also referred to as *blind* re-transmissions because if configured to do so, transmitters re-transmit every packet regardless of whether the earlier transmission(s) are received successfully or not. This differentiates re-transmissions in C-V2X and NR V2X from the conventional hybrid automatic repeat request (HARQ) used in cellular 3GPP technologies.

Re-transmissions are beneficial from the PHY layer perspective because multiple copies of a packet can provide frequency and time-domain diversity. At the receiver, a packet is successfully received as long as one of its copies is correctly decoded by the receiver. Thus, the probability of packet errors is reduced, thereby enhancing the link-level

This work was partially sponsored by NSF through grants 1563832, 1642928, and 1822173, and by the industry affiliates of the Broadband Wireless Access & Applications Center (BWAC).
Approved for Public Release; Distribution Unlimited: 88ABW-2016-2948.

[1]This interface, defined in 3GPP Rel. 12, enables direct communications between user equipment without the involvement of the cellular infrastructure.

performance. From the MAC layer perspective, however, it has been shown in the context of C-V2X that the benefit of achieving higher transmission reliability diminishes as the traffic density[2] increases [2]. The 3GPP Rel. 14 allows for static configuration of a User Equipment (UE), i.e., vehicle, to transmit each packet once (i.e., no re-transmissions) or twice (i.e., one re-transmission) as well as a dynamic configuration whereby UEs can autonomously decide whether or not to re-transmit a packet. Similar provisions are likely to be provided in 3GPP Rel. 16 as well, except the number of permissible re-transmissions is likely to be more than one [1]. However, to the best of our knowledge, the problem of how to enable UEs to autonomously (i.e., without assistance from the cellular infrastructure) decide when to activate re-transmissions remains unaddressed, and it is the subject of this paper.

In this paper, through an extensive simulation study carried out using network simulator-3 (ns-3), we show that static re-transmission configurations provide superior performance only in specific scenarios. Although in theory, re-transmissions are beneficial (especially from the PHY layer perspective) and is often claimed to be a key reliability-enhancing feature in C-V2X and NR V2X [3], in practice, this feature can have the opposite effect and seriously degrade system performance if it is not implemented carefully. Thus, we argue the need for a re-transmission control mechanism, i.e., a mechanism that allows the vehicles to autonomously decide whether or not to re-transmit packets based on locally available information. We then propose Channel Congestion-based Re-transmission Control (C^2RC)—a simple, yet effective, probabilistic re-transmission control mechanism using which the UEs can use re-transmissions to maximize system performance in lightly-loaded scenarios while avoiding re-transmitting in heavily loaded conditions, thereby avoiding performance degradation. Although the discussions in this paper on the impact of re-transmissions and the proposed re-transmission control mechanism apply to both C-V2X and NR V2X, our simulation study is restricted to C-V2X sidelink mode 4 because the specifications for NR V2X are not yet available. Throughout this paper, we refer to the terms vehicles and UEs interchangeably. The main contributions of this paper are as follows.

- We investigate the impact of re-transmissions on C-V2X sidelink mode 4 performance. Based on our observations, we motivate the need for re-transmission control mechanisms in C-V2X sidelink mode 4.
- We perform a systematic study on the impact of re-transmissions on the observed channel congestion. We show that beyond a certain congestion level, switching on re-transmissions leads to the re-use of resources, which can potentially degrade the overall system performance.
- We propose C^2RC—a mechanism that allows C-V2X UEs to enjoy the benefits of using packet re-transmissions when they are truly beneficial while minimizing the impact of re-transmissions in congested environments.

[2]We say that the transmission density increases when the number of vehicles increases and/or the packet transmission rate (in Hz) increases.

II. Related Work

The authors in [4] and [5] study the performance of C-V2X sidelink mode 3, where the UEs are within cellular coverage. On the other hand, the authors in reference [2] and [6] evaluate the performance of C-V2X mode 4 in the highway and urban scenarios, respectively. In [2], the authors show that packet re-transmissions are not always helpful. Furthermore, the authors show the impact of several transmission parameters such as the packet transmission rate, modulation and coding scheme, etc. on the system performance of sidelink mode 4 in highway scenarios. The authors in [7] perform a comprehensive study on the impact of different PHY and MAC layer parameters on the system performance of C-V2X sidelink mode 4.

An analytical model to capture the performance of C-V2X mode 4 has been developed in [8]. Through their analytical model, the authors accurately characterize the performance of C-V2X mode 4 in different scenarios. However, due to the complexity of the model, the performance gain (or lack thereof) due to re-transmissions is ignored in [8]. The aforementioned works provide several insights on the performance of C-V2X. However, the study of re-transmissions in C-V2X as a reliability-enhancing mechanism has not received much attention in the literature, with only brief discussions in [2].

III. Background & Simulation Setup

The C-V2X sidelink mode 4 algorithm is defined in 3GPP Rel. 14. In the interest of space, we do not describe the C-V2X mode 4 algorithm in detail in this paper. We refer the interested reader to 3GPP Rel. 14 [9], [10] for a complete description of the algorithm, while reference [7] investigates its performance in detail. In what follows, we provide a brief overview of the algorithm required to understand the re-transmission control mechanism, C^2RC, proposed in Sec. V.

When the C-V2X sidelink mode 4 algorithm is triggered by a UE, it first senses the operating channel for a fixed 1 second interval (this interval is referred to as the *sensing window*) and uses these sensing results to select a *resource* within the *selection window*[3]. A resource in C-V2X constitutes a group of contiguous Resource Blocks (RBs) in the frequency domain and lasts for 1 sub-frame (i.e., 1 msec) in time. Each time a resource is selected, the UE reserves the selected resource for a number of subsequent transmissions. This number is referred to as the re-selection counter (denoted as C_{resel} in [10]). C_{resel} is randomly chosen in [5, 15] if the packet transmission rate is 10 Hz, in [10, 30] for 20 Hz and in [25, 75] for 50 Hz. After each packet transmission, C_{resel} is decremented by 1, and if C_{resel} becomes zero, the mode 4 algorithm is triggered to select a new resource with a probability known as the *reslection probability*, p_{resel}. Finally, when the UE transmits a packet, in addition to the actual payload, it broadcasts the resource selection information to aid other UEs in selecting the best resource for their respective transmissions.

[3]The selection window is the time duration within which the packet must be transmitted in order to satisfy its latency requirement.

The results presented in this paper have been obtained using our ns-3 simulator. Our ns-3 simulator implements all functionalities of C-V2X sidelink mode 4, including the mode 4 algorithm as per the 3GPP specifications [9], [10]. This algorithm can be used in all deployment scenarios. However, in urban environments, vehicles are more likely to be within cellular coverage. Therefore, resource allocation of vehicle-generated packets can be done by the eNodeB using sidelink mode 3. In the rest of this paper, we restrict our focus to the performance of C-V2X mode 4 in highway environments.

The performance metric chosen in our study is the packet delivery ratio (PDR)—the ratio of the number of packets received from a certain transmitter to the number of packets sent by that transmitter. In cases where a transmitter sends each packet twice, each packet is successfully decoded at the receiver as long as one of its copies is received successfully. Losses in a wireless communication system can either occur due to packet collisions or wireless channel characteristics. Our ns-3 simulator handles both types of losses. The wireless channel for a transmitter-receiver pair is modeled as per [11], which provides a look-up table for the probability of packet error for a given signal-to-interference-plus-noise ratio (SINR). For each received packet, the SINR is computed as the ratio of the signal power to the noise power (in case of no collisions) or noise plus interference power (in case of packet collisions).

The 3GPP provides guidelines on simulation topologies for the performance evaluation of C-V2X [12]. For the highway scenario, 3GPP recommends that at least a 2 kilometer (km) road-stretch must be simulated, with three 4 m wide lanes in each direction. Scenarios are identified by the average velocity of vehicles. Given the average velocity, the inter-vehicle distance in each lane is computed as the distance traveled by a vehicle in 2.5 seconds. The two standard scenarios specified in [12] are for 70 km/hr and 140 km/hr, which represent medium (120 vehicles/km) and sparse (60 vehicles/km) density of vehicles, respectively. The simulation parameters used in this paper are as outlined in Table I.

Parameter	Value	Parameter	Value
Noise Figure	9 dB	Bandwidth	10 MHz
Sub-channel size	24 RBs	Packet size	190/300 Bytes
Loss Model	WINNER+ B1	Tx power	23 dBm
Modulation scheme	QPSK	Code rate	0.7 (190 B) 0.5 (300 B)
RSRP Thresh. see [7]	−110 dBm	p_{resel} see [2], [7]	1

TABLE I: Simulation Parameters

IV. Is Re-transmission Control Needed?

A. *Impact of Re-transmissions on System Performance*

We first study the performance of C-V2X sidelink mode 4 in the two standard highway scenarios given in [12]. The purpose of this study is to understand if re-transmission control mechanisms are needed and if so, under what scenarios are such mechanisms useful. Consistent with [6], we refer to the 70 km/hr and 140 km/hr scenarios as the *Highway Slow* and *Highway Fast* scenarios, respectively. For these two highway scenarios, Fig. 1 shows the average PDR observed across all C-V2X receivers as a function of the receiver's distance from a given transmitter. The packet transmission rate, specified in Hz (10, 20 or 50) indicates the number of packets transmitted by each C-V2X UE every second. 1 Tx/packet indicates that each packet is transmitted only once, whereas 2 Tx/packet indicates that each packet is always re-transmitted. Hereafter, we refer to the static configuration where re-transmissions are always disabled or enabled at all UEs as the always OFF and always ON modes, respectively.

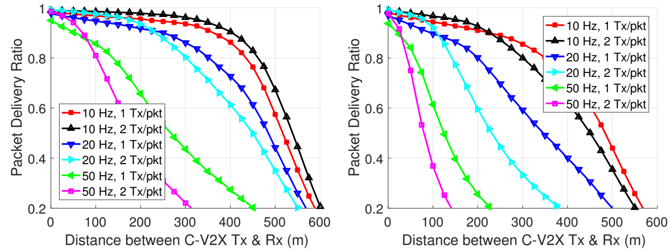

(a) Highway Fast (b) Highway Slow

Fig. 1: C-V2X mode 4 performance in 3GPP-defined scenarios

Fig. 1 shows that except for the 10 Hz, Highway Fast scenario, the always ON mode performs inferior to the always OFF mode beyond a certain distance between the transmitter and receiver. For example, in the Highway Fast scenario with 20 Hz transmission rate, if a transmitter has to communicate with a receiver greater than 220 m away, the transmitter is better off in sending each packet only once (as opposed to sending each packet twice) in terms of the average PDR observed at the receiver. The primary reason for this behavior is channel congestion. The always ON mode can require up to twice[4] the number of resources used in the always OFF mode. Beyond a certain traffic density, when re-transmissions are switched ON, C-V2X UEs end up re-using a fraction of resources (i.e., a C-V2X UE re-uses a resource already reserved by another C-V2X UE), which leads to packet collisions and consequently, performance degradation.

Further, we see that in all scenarios, the always ON mode results in a higher PDR at small transmitter-receiver distances. For small distances, even if a collision occurs on a particular resource, the SINR can be sufficient to decode the strongest signal, which is typically sent by the transmitter located very close to the receiver. Therefore, packets transmitted by UEs that are located close to the receiver have a high probability of being successfully decoded, especially if each packet is transmitted twice. However, as the traffic density increases, the always ON mode results in the re-use of a significant number of resources, i.e., a large fraction of resources suffer from collisions. In such cases, packets sent by far-away transmitters are less likely to be decoded, resulting in poor performance relative to the always OFF mode at large distances.

[4]Ideally, as long as UEs in the always OFF mode cumulatively use up less than 50% of the available resources, switching from the always OFF to the always ON mode should double the number of used resources.

B. Need for Re-transmission Control

It must be noted from Fig. 1 that at relatively low densities (e.g., Highway Fast, 10 Hz), re-transmissions improve the system performance for all transmitter-receiver distances. It is only when the traffic density exceeds a certain limit that the always ON mode begins to perform inferior to the always OFF mode for large transmitter-receiver distances. We leave the determination of the exact traffic density and the distance at which this behavior occurs to future work. Our focus in this paper, rather, is to develop a mechanism using which UEs can autonomously decide whether or not to enable re-transmissions, thereby leading to superior performance in sparse scenarios, while mitigating the negative impact of re-transmissions on performance in dense conditions.

The communication requirements of vehicular safety applications differ in terms of latency, reliability, and range. Some applications require high reliability with a short range. For such applications, the always ON mode is suitable. However, this performance gain at short distances comes at the cost of significant loss at larger distances as shown in Fig. 1. For certain safety applications, a large communication range is desired. Thus, high reliability at short distances alone is not sufficient to support all safety applications. In practical scenarios, vehicles are likely to participate in multiple safety applications simultaneously. For example, each vehicle will process packets received from near-by vehicles to detect another vehicle in its blind spot, while using a packet from a far-away vehicle to detect an approaching emergency vehicle. Considering this diverse set of safety applications, to guarantee performance across a wide range of transmitter-receiver distances, a re-transmission control mechanism is necessary.

V. Design of C^2RC

A. Channel Congestion due to Re-Transmissions

A detailed explanation of the impact of re-transmissions on system performance, particularly the observed performance loss at certain densities, requires a closer look at the congestion level observed in the channel when the C-V2X mode 4 algorithm is used. Toward this objective, we conduct simulations for a hypothetical scenario where all vehicles can sense each other. In this scenario, each vehicle must account for every other vehicle's transmissions while selecting a resource for its transmissions. The objective of studying this hypothetical scenario is to study the impact of switching ON re-transmissions on resource utilization when the C-V2X mode 4 algorithm is used. Before we do so, however, we must define a metric to quantify the congestion level in the network. We use the 3GPP-defined metric—Channel Busy Ratio (CBR), which is defined at any time instant as the ratio of the number of resources *utilized* in the previous 100 sub-frames (i.e., 100 milliseconds) to the total number of resources available in previous 100 sub-frames. Note that this definition of CBR requires a threshold energy value beyond which a resource is declared as utilized. We elaborate on what threshold to use in the next sub-section. For the sake of the following discussions, it can be assumed that all UEs sense each other at above the chosen energy threshold.

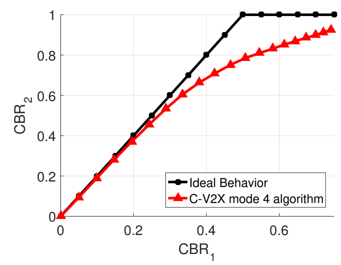

Fig. 2: Effect of re-transmissions on the observed CBR. All UEs can sense each other at above the energy threshold.

For a given network configuration, let CBR_1 be the average CBR observed at a UE in the always OFF mode, while CBR_2 denotes the average CBR observed in the always ON mode. Fig. 2 shows the variation in CBR_2 as a function of CBR_1. Observe that as long as CBR_1 is under $\sim$0.3, transmitting a packet twice nearly doubles the observed congestion level, which is intuitive behavior. However, beyond this CBR_1 level, switching on re-transmissions leads to re-utilization of some of the resources, leading to packet collisions on some of the utilized resources even as other resources in the channel remain idle. When CBR_1 increases further, the curve significantly deviates from linear behavior, implying that a non-negligible fraction of resources is, in fact, re-used. Thus, we can infer from Fig. 2 that after the observed CBR level at a UE exceeds a certain level (not necessarily 0.3 as we show in Sec. V-D), using the always ON mode is likely to negatively impact the overall system performance. Consequently, if the observed CBR is large enough, re-transmissions should be disabled, at least for some UEs, to prevent performance degradation.

B. Choice of Energy Threshold

In the previous sub-section, we noted that the definition of CBR requires the selection of an energy threshold (say γ) such that if the resource energy is greater than γ, the resource is declared as occupied. In the following discussions, we elaborate on the choice of an appropriate energy threshold.

For the simulations reported in this sub-section, we vary the number of vehicles in the 2 km stretch, with the inter-vehicle distance maintained such that vehicles are distributed uniformly along the 2 km road length. Fig. 3 shows the ratio of CBR_2 to CBR_1 as a function of the number of vehicles for two different packet transmission rates (i.e., 10 Hz and 20 Hz). Note from our discussions in the previous sub-section that re-transmissions begin to negatively affect the system performance when doubling the number of transmissions does not result in a corresponding increase in the observed congestion level. The number of UEs after which re-transmissions hurt the overall system performance is 125 and 70 for the 10 Hz and 20 Hz packet transmission rates (not shown in Fig. 3), respectively. For values of γ above -105 dBm, the ratio of CBR_2 to CBR_1 does not vary significantly with the number of vehicles. The non-linear behavior of ratio becomes

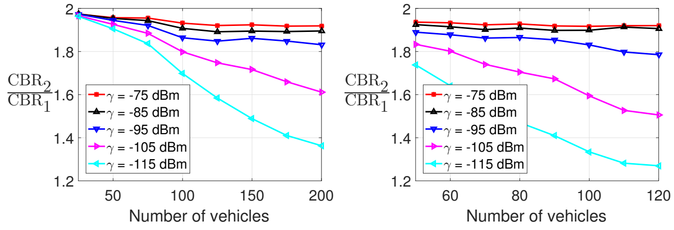

(a) 10 Hz (b) 20 Hz

Fig. 3: Impact of energy threshold, i.e., γ.

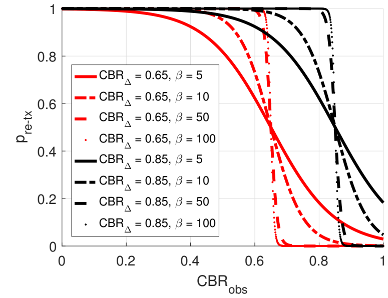

Fig. 4: Impact of CBR_Δ and β on $p_{\text{re-tx}}$.

prominent only at $\gamma = -105$ dBm and below. Although results are presented in this paper only for two specific cases, a similar trend was observed when parameters such as packet transmission rate and sub-channel size were varied. Thus, *an energy threshold, i.e., γ, of -105 dBm is suitable for estimating the congestion level on the operating channel.*

C. Channel Congestion-based Re-transmission Control

In C^2RC, when the UE uses the C-V2X mode 4 algorithm for resource selection, it decides on whether to enable or disable re-transmissions for the upcoming C_{resel} transmissions (see Sec. III) based on the CBR level observed in the cycle of previous C_{resel} transmissions. Let the average CBR observed over the previous C_{resel} transmissions be CBR_{obs}. This CBR_{obs} can readily be computed at C-V2X UEs since its computation only depends on the measurement of Sidelink Received Signal Strength Indicator (S-RSSI) [13], which is measured during each cycle of C_{resel} transmissions. The C-V2X UE then computes a probability, $p_{\text{re-tx}}$, with which the it will enable re-transmissions for the upcoming C_{resel} transmissions. Based on our discussions in Sec. V-A, it is clear that UEs must re-transmit packets if CBR_{obs} is low and refrain from re-transmitting packets if CBR_{obs} is too high. Thus, the probability $p_{\text{re-tx}}$ must meet the following criteria.

- The probability $p_{\text{re-tx}}$ must be a non-increasing function of the observed CBR level, CBR_{obs}.
- The rate at which $p_{\text{re-tx}}$ drops from 1 to 0 with CBR_{obs} must be configurable.

In this paper, we choose the shifted and inverted logistic function to define $p_{\text{re-tx}}$ as follows:

$$p_{\text{re-tx}} = 1 - \frac{1}{1 + e^{-2\beta(\text{CBR}_{\text{obs}} - \text{CBR}_\Delta)}},$$

where $\text{CBR}_\Delta \in [0, 1]$ and $\beta \in (0, \infty)$ are tunable parameters. The significance of CBR_Δ and β can be understood from Fig. 4. CBR_Δ signifies the CBR level at which $p_{\text{re-tx}}$ drops to 0.5, while β controls the rate at which $p_{\text{re-tx}}$ drops from 1 to 0. If $\beta = \infty$, $p_{\text{re-tx}}$ resembles a shifted and inverted unit step function at CBR_Δ, while for small values of β, $p_{\text{re-tx}}$ smoothly transitions from 1 to 0.

D. Selection of C^2RC Parameters

Fig. 5 shows the performance of C^2RC in comparison to the static re-transmission modes for three different traffic densities. Fig. 5a represents a lightly loaded scenario where 100 UEs (10 Hz periodicity) are spread out over the 2 km stretch. As expected, the always ON mode outperforms the always OFF mode for all distances between the transmitter and the receiver. This is because the channel congestion is low enough that switching on re-transmissions does not lead to significant re-use of resources. The choice of parameters CBR_Δ and β in this scenario must be such that C^2RC must enable re-transmissions with a high probability in this scenario. We see that as CBR_Δ progressively increases, the resulting C-V2X system performance keeps improving and approaches the performance of the always ON mode as $\text{CBR}_\Delta = 0.65$. A CBR_Δ value beyond 0.65 yields no performance improvement as the performance is bounded by that of the always ON mode.

Next, Fig. 5b represents a moderately loaded scenario with 200 UEs, each transmitting at 20 Hz. We see that as the channel load increases (from what was shown in Fig. 5a), an increase in CBR_Δ leads to a deterioration in system performance, albeit marginally (when $\beta = 100$). This is because the channel congestion is now large enough that doubling the number of transmissions (as a result of always re-transmitting) leads to significant re-use of resources and contributes negatively to the system performance. This trend is the reverse of what was shown in Fig. 5a. Further, although not shown in Fig. 5b, the parameter β has negligible impact on the system performance if CBR_Δ is small (0.1 through 0.55). However, at large values of CBR_Δ (such as 0.65), choosing a small β results in a non-negligible probability of re-transmissions even if the observed CBR is very high. For example, Fig. 4 shows that even if the CBR at a UE is 0.8, $\text{CBR}_\Delta = 0.65, \beta = 5$ results in $p_{\text{re-tx}} = 0.2$. This results in a slightly inferior system performance when $\beta = 5$ is selected.

Finally, Fig. 5c shows the performance in a heavily loaded scenario where 300 UEs transmit at 50 Hz. In this case, the observed CBR at all UEs in the always OFF mode is very high ($\sim$0.9). Consequently, any reasonable choice[5] of CBR_Δ has a negligible impact on the resulting system performance. However, similar to the moderately loaded scenario, if a small value of β is chosen at large values of CBR_Δ (e.g., 0.65), the system performance will deteriorate as shown in Fig. 5c.

Note from Fig. 5b and 5c that when C^2RC is used, the average PDR at C-V2X transmitter-receiver distances in the [0-160] m (moderate scenario) and [0-30] m (dense scenario)

[5]Naturally, this is true when $\text{CBR}_\Delta < 0.9$. However, at such congestion levels, it is obvious that re-transmissions will hurt the system performance.

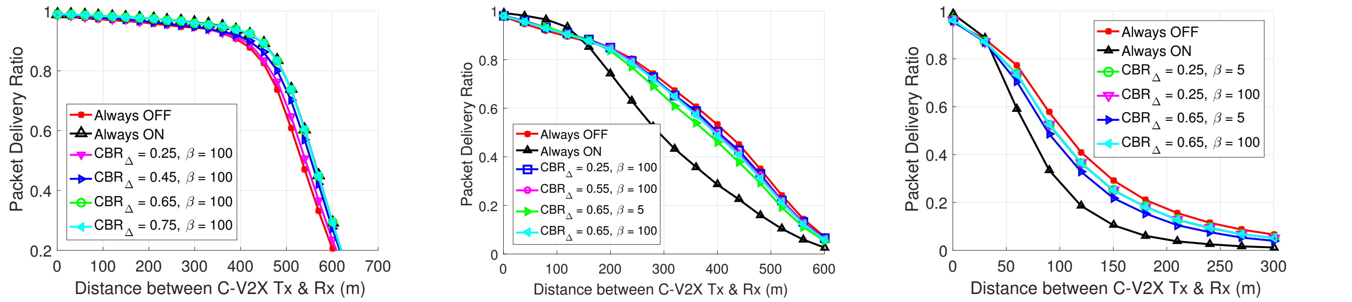

(a) Light, 10 Hz, 100 vehicles (b) Moderate, 20 Hz, 200 vehicles (c) Heavy, 50 Hz, 300 vehicles

Fig. 5: Performance of C^2RC in different traffic densities.

range[6] is in-between the average PDR of the two static modes. Thus, a small improvement over the always OFF mode is observed for smaller transmitter-receiver distances. However, a substantial improvement is observed over the always ON mode for large transmitter-receiver distances. Because the UEs are randomly dropped along the road, in some instances, certain sections of the road can be less congested than others, which is likely in real-life scenarios as well. A marginal performance loss relative to the always OFF mode at larger distances occurs due to re-transmissions from a few C-V2X UEs that re-transmit because the locally observed CBR may be low enough to trigger the second transmission even though the channel is congested at far-away receivers.

Based on the above discussions, we infer that a small value of CBR_Δ is more suitable for moderate channel loads, while a large value of CBR_Δ is optimal for low channel loads. Yet, increasing $CBR_\Delta > 0.65$ yields very little performance improvement in light conditions. Furthermore, at moderate densities, even though an increase in CBR_Δ degrades the performance, this degradation is negligible as long as β is large. Based on our simulation results, $\beta \geq 100$ prevents performance degradation even when CBR_Δ is as high as 0.65.

In summary, *our simulations reveal that* $CBR_\Delta = 0.65$ *and* $\beta = 100$ *strikes a balance between avoiding performance loss at high densities and improving performance at low densities*. This is, however, true as long as the chosen energy threshold for declaring a resource as occupied is $\gamma = -105$ dBm. Although simulation results for a limited set of scenarios have been shown in Fig. 5, the implications of selecting different values for CBR_Δ and β as discussed here were tested to be true in a wide range of traffic densities.

VI. Conclusions

In this paper, we take an in-depth look at the performance implications of packet re-transmissions, which is used in C-V2X and is likely to be re-used in NR V2X. We show that this feature, although useful, must be used with caution. Since both C-V2X and NR V2X must operate in distributed environments without any assistance from the cellular infrastructure, it is critical for the UEs to decide on whether or not to re-transmit packets based only on the locally observed channel congestion. To address this, we propose and evaluate C^2RC—a probabilistic re-transmission control mechanism that allows the UEs to maximize performance when re-transmissions are truly beneficial while preventing performance degradation in scenarios where UEs are better off transmitting only once.

[6] 160 m and 30 m are the distances beyond which the always ON mode performs worse in the moderate and heavy load scenarios, respectively.

VII. Disclaimer

Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the U.S. government or AFRL.

References

[1] G. Naik, B. Choudhury, and J.-M. Park, "IEEE 802.11 bd & 5G NR V2X: Evolution of radio access technologies for V2X communications," *IEEE Access*, vol. 7, pp. 70169–70184, 2019.
[2] R. Molina-Masegosa and J. Gozalvez, "LTE-V for Sidelink 5G V2X Vehicular Communications: A New 5G Technology for Short-Range Vehicle-to-Everything Communications," *IEEE Vehicular Technology Magazine*, vol. 12, no. 4, pp. 30–39, 2017.
[3] Qualcomm Inc., "Accelerating C-V2X commercialization." https://www.qualcomm.com/documents/path-5g-cellular-vehicle-everything-c-v2x, September 2017.
[4] A. Bazzi, B. M. Masini, A. Zanella, and I. Thibault, "On the Performance of IEEE 802.11 p and LTE-V2V for the Cooperative Awareness of Connected Vehicles," *IEEE Transactions on Vehicular Technology*, vol. 66, no. 11, pp. 10419–10432, 2017.
[5] A. Bazzi, B. M. Masini, and A. Zanella, "How many vehicles in the LTE-V2V awareness range with half or full duplex radios?," in *2017 15th International Conference on ITS Telecommunications (ITST)*, pp. 1–6, IEEE, 2017.
[6] R. Molina-Masegosa and J. Gozalvez, "System level evaluation of LTE-V2V mode 4 communications and its distributed scheduling," in *2017 IEEE 85th VTC Spring*, pp. 1–5, IEEE, 2017.
[7] A. Bazzi, G. Cecchini, A. Zanella, and B. M. Masini, "Study of the Impact of PHY and MAC Parameters in 3GPP C-V2V Mode 4," *IEEE Access*, vol. 6, pp. 71685–71698, 2018.
[8] M. Gonzalez-Martín, M. Sepulcre, R. Molina-Masegosa, and J. Gozalvez, "Analytical Models of the Performance of C-V2X Mode 4 Vehicular Communications," *IEEE Transactions on Vehicular Technology*, vol. 68, no. 2, pp. 1155–1166, 2019.
[9] 3GPP, "3GPP TS 36.213: Evolved Universal Terrestrial Radio Access (E-UTRA); Physical layer procedures (v14.3.0, Release 14)," June 2017.
[10] 3GPP, "3GPP TS 36.321: Evolved Universal Terrestrial Radio Access (E-UTRA); Medium Access Control (MAC) protocol specification (v14.3.0, Release 14)," June 2017.
[11] Huawei, HSilicon, "R1-160284: DMRS enhancement of V2V." 3GPP TSG RAN WG1 Meeting # 84, St. Julian's Malta, 2016.
[12] 3GPP, "3GPP TR 36.885: Study on LTE-based V2X Services (v14.0.0, Release 14)," June 2016.
[13] Qualcomm Incorporated, "R1-1611594: Congestion control for V2V." 3GPP TSG RAN WG1 Meeting # 87, Reno, USA, 2016.