Is Packet Dropping a Suitable Congestion Control Mechanism for Vehicular Networks?

Miguel Sepulcre, Jorge Mira, Gokulnath Thandavarayan, Javier Gozalvez
Universidad Miguel Hernandez de Elche (UMH)
UWICORE laboratory, http://www.uwicore.umh.es
Avda. de la Universidad, s/n, 03202, Elche, Spain
msepulcre@umh.es, jorge.mira01@goumh.umh.es, gthandavarayan@umh.es, j.gozalvez@umh.es

***Abstract*— Vehicular networks are based on the wireless exchange of information between vehicles and between vehicles and road side units. To avoid overloading the radio channel, different distributed congestion control mechanisms have been proposed in the literature. A widely adopted mechanism to control congestion is packet rate control. It controls the number of packets that each vehicle transmits to the radio channel through packet dropping. Packet dropping has been shown to effectively improve the radio communications performance thanks to the reduction of the channel load and packet collisions. However, from the application perspective, packets dropped by congestion control mechanisms are not transmitted and are therefore lost. This paper demonstrates for the first time that, while packet dropping can improve the performance at the radio level, it degrades the performance at the application level. This raises the question on whether current congestion control protocols based on packet dropping are actually suitable for vehicular networks.**

***Keywords*—Vehicular networks, V2X networks, congestion control, DCC Access, ETSI, Adaptive, Reactive, packet dropping, ITS-G5, ns-3, open source.**

I. Introduction

Connected vehicles are expected to improve traffic safety and efficiency through V2V (Vehicle-to-Vehicle) and V2I (Vehicle-to-Infrastructure) communications or, in general, through V2X (Vehicle-to-Everything) communications. The widespread deployment of connected vehicles and the increasing number of applications and messages that need to be transmitted could overload the radio channel. To avoid overloading the radio channel, the use of adequate congestion control protocols is of primary importance. This type of protocols are used to control the channel load and interferences generated. To this aim, they dynamically adapt the number of packets transmitted by each vehicle, their transmission power or their data rate.

Packet rate control is one of the most commonly used congestion control mechanisms. It makes use of packet dropping to control the number of packets that each vehicle transmits to the radio channel. Packet rate control is part of the DCC (Decentralized Congestion Control) framework defined by ETSI (European Telecommunications Standards Institute). It is generally assumed that the use of packet dropping improves the performance of vehicular networks. This is the case because packet dropping reduces the number of packets transmitted over the radio channel, and therefore the channel load and packet collisions. This assumption can be considered true if the performance is measured at the radio level, because the probability of correctly receiving a packet that has been effectively transmitted increases. However, from the application perspective, packets dropped by congestion control mechanisms are not transmitted and are therefore lost (i.e. not received by any nearby vehicle). As a result, the impact of congestion control on the performance of vehicular applications could be negative if the packet loses due to packet dropping are not compensated by the reduction of the channel load and packet collisions.

Previous studies have focused on the performance evaluation of congestion control mechanisms at the radio level. To the authors knowledge, there is no research that analyzes the performance at the application level to determine if using congestion control based on packet dropping actually has a benefit or not. In this context, the main goal of this paper is to analyze for the first time the impact of congestion control based on packet dropping on the application-level performance of vehicular networks and compare it with the radio-level performance. This work utilizes the DCC Access protocol defined by ETSI as a benchmark. We have conducted this analysis for multiple scenarios and configurations, including different traffic densities, different types of packets with the same and different priorities, and the two DCC Access approaches defined by ETSI (Reactive and Adaptive). The obtained results demonstrate that congestion control based on packet dropping has a negative impact on the performance at the application level. The obtained results therefore question the suitability of this type of protocols for vehicular networks.

II. ETSI DCC Access

To control the radio channel congestion in vehicular networks, ETSI has defined a DCC framework that includes components at different layers of the protocol stack. All these components have been particularly designed for the ITS-G5 technology. Most of the research conducted to date has focused on the study and optimization of the DCC Access component of the DCC framework. This component controls the data traffic injected to the radio channel through packet dropping at the Access layer. As part as DCC Access, two approaches have been specified by ETSI [1]: Reactive and Adaptive. Different studies have shown that the performance and stability of the Reactive approach can be significantly challenged. For example, [2] evaluated its performance for platooning, showing its limitations due to its low granularity and threshold parametrization. It also identified ways for improvement, including more appropriate control criteria or novel control

The authors acknowledge the support of the Spanish Ministry of Economy and Competitiveness and FEDER funds under the project TEC2017-88612-R, and the support of the Horizon2020 Programme through the TransAID project (Grant Agreement no. 723390).

algorithms. The work in [3] evaluated the performance of the Reactive approach in terms of radio coverage range, delay and DCC stability considering a time-varying traffic density, showing the need for certain stabilization times. The Adaptive approach was initially proposed in [4], which demonstrated its good stability and convergence properties independently of the number of vehicles. The work in [4] also showed that the performance of the Adaptive approach can be significantly higher than the performance achieved by the Reactive one. Both approaches have also been compared in [5], where the superior performance of Adaptive over Reactive was demonstrated in terms of packet error ratio or packet inter reception time. It is important to note that the focus of these studies was on the radio communications performance, and did not analyze the impact of packets losses due to dropping on the application performance.

A. Control mechanisms

DCC Access controls the data traffic injected by each vehicle to the radio channel for ITS-G5. It makes use of Prioritization, Queuing and Flow Control mechanisms, as described below.

All packets received by the DCC Access component from the upper layers are first classified according to their priorities. Four different priorities are differentiated, depending on the four DCC Profiles (DPs): DP0, DP1, DP2 and DP3, where DP0 has the highest priority. At the lower layers, these priorities are mapped to the traffic categories of the ITS-G5 EDCA (Enhanced Distributed Channel Access). DCC Access implements 4 different queues, each of them for one packet priority or DCC Profile. Each queue follows a first-in-first-out (FIFO) scheduling policy so that the packet that has been waiting longer in the queue is transmitted first. The DCC Access queuing mechanism drops those packets that have been waiting in the queue for a time longer than their lifetime. When a queue is full, no more packets are accepted.

Finally, flow control is applied to de-queue packets from the DCC queues and send them to the lower layers for their radio transmission. Packets with higher priorities are de-queued first. A packet is only de-queued if there is no packet with a higher priority waiting in its corresponding queue. As a result, lower priority packets can suffer from starvation and never be transmitted. To control the rate of transmitted packets per vehicle, two approaches have been defined in [1]: Reactive and Adaptive. Both approaches adapt the time between consecutive packet transmissions based on the CBR (Channel Busy Ratio). CBR is defined as the percentage of time that the channel is sensed as busy. These two approaches are described below.

B. Reactive approach

The Reactive approach is based on a state machine where the current state depends on the CBR and each state can only be reached by a neighboring state. The *Restrictive* state is the most stringent one, i.e. the one reached with the highest CBR. The *Relaxed* state is the least stringent one. Intermediate states called *Active 1*, *Active 2*, … *Active n*, can also be defined. The number of states is not fixed and can be configured. For each state, different radio transmission parameters can be defined to control the channel load depending on the CBR. The ETSI specification allows the adaptation of the data rate or the transmission power, but in related studies only the packet rate (i.e. the time between packet transmissions or T_{off}) is configured differently in each state. In [1], a possible parameter setting is provided as Informative Annex. In this setting, 5 states are defined, and only the packet rate is adapted following the T_{off} values shown in Table I for scenarios where the packet duration T_{on} is below 0.5ms. Following this table, if a vehicle requires the transmission of e.g. 8 packets per second and the channel load is 51% (State Active 3), DCC will only allow the transmission of 4 packets per second and will drop the rest. Other configurations are possible, but the one in Table I is one of the most used ones.

TABLE I. MAPPING OF CBR VALUES TO STATES AND T_{OFF} FOR T_{ON}=0.5MS [1]

State	CBR	Packet rate	T_{off}
Relaxed	< 30%	20 Hz	50 ms
Active 1	30% to 39%	10 Hz	100 ms
Active 2	40% to 49%	5 Hz	200 ms
Active 3	50% to 65%	4 Hz	250 ms
Restrictive	> 65%	1 Hz	1000 ms

C. Adaptive approach

The Adaptive approach makes use of a linear control process to adapt the packet rate of each vehicle. This process is designed to make the CBR converge to a target value CBR_{target}=68%. To this aim, it adapts the parameter δ, which is a unitless value that represents the maximum fraction of time that a vehicle is allowed to transmit. The parameter δ is updated every 200 ms based on the difference between the current CBR and the target CBR. Then, the computed δ is used to calculate the time between packet transmissions (T_{off}), taking into account the duration of the current packet (T_{on}). This approach has been shown to converge to a stable solution in steady state [4]. More details about the parameters and equations can be found in [1].

III. EVALUATION

A. Settings

To conduct this study, we have developed a module for the network simulator ns-3 that implements the DCC Access protocol defined in [1] and described in the Section II. This module has been tested with ns-3.26 and ns-3.20, but can be easily integrated in other versions of ns-3. We also plan to integrate it into the iTETRIS (an Integrated Wireless and Traffic Platform for Real-Time Road Traffic Management Solutions) simulation platform [6]. The source code developed by the authors for this work is made available in [7], where instructions and examples have also been published.

The simulations conducted consider a highway scenario of 5km with 6 lanes (3 lanes in each direction). In this scenario, vehicles move at 100 km/h, and 3 different traffic densities have been simulated: 60, 120 and 180 veh/km. The simulation time was set to 100 seconds. Only the vehicles located in the 2km around the center of the simulation scenario are used to collect the results to avoid boundary effects.

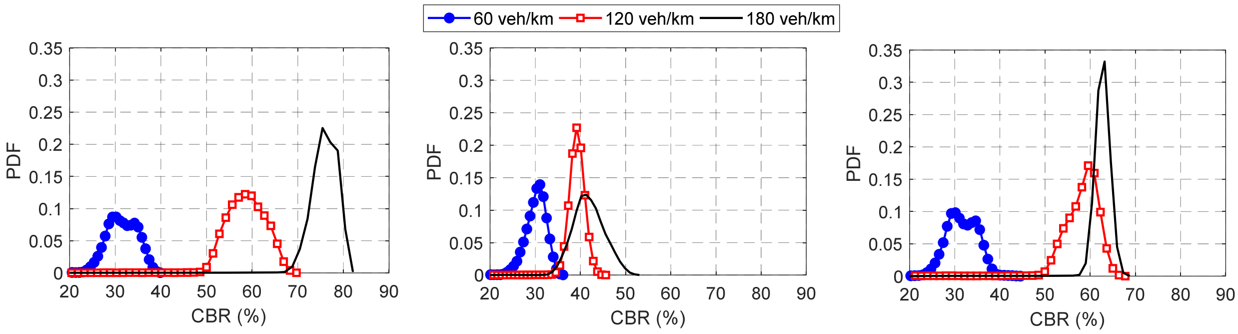

Fig. 1. PDF of the CBR experienced for different traffic densities and DCC configurations, when all packets are configured with the same priority (DP2).

At the Facilities layer, all vehicles periodically generate CAMs (Cooperative Awareness Messages) and CPMs (Collective Perception Messages). CAMs contain the position, speed and basic status information of the transmitting vehicle. CPMs contain the position and speed of the objects that are detected by the vehicles' onboard sensors (e.g. radar, lidar, etc.). Following [8], in this study CAMs are generated at the Facilities layer at 3 Hz with a payload of 350 Bytes. Based on [9], each vehicle generates CPMs at a rate of 8.5 Hz with a payload of 261 Bytes. CAMs are configured with a DP2 priority following the ETSI specifications. Different simulations have been performed considering CPMs as DP2 and DP3, since ETSI has not decided yet its final value [10].

DCC Access has been configured following [1]. The queue length of all DCC queues in this study is 10 packets, and the life-time of each packet is equal to 1 second. The CBR is locally measured by each vehicle every 100 ms. The Reactive approach has been configured with 5 states following Table I. The Adaptive approach has been configured following [1], and in particular with a CBR_{target} of 68%.

All vehicles are equipped with an ITS-G5 radio interface. All packets are transmitted using the 6Mbps data rate (QPSK – Quadrature Phase Shift Keying – with coding rate 1/2) over a 10MHz channel at the 5.9GHz frequency band. The transmission power is 23dBm. The sensing power threshold has been set equal to -85dBm. The propagation effects are modeled using the Winner+ B1 propagation model following the 3GPP guidelines for V2X simulations [11].

B. Performance at application and radio levels

To understand the impact of DCC, it is important to analyze the channel load experienced and the behavior of each DCC approach. Fig. 1a plots the PDF (Probability Density Function) of the CBR experienced by the vehicles when DCC is not used for the three traffic densities considered. As it can be observed, a CBR of around 20%-40% is experienced in the low traffic density scenario. The CBR increases to 50%-70% for the intermediate traffic density, and it goes up to around 70%-80% for the highest traffic density.

Fig. 1b shows the CBR when all vehicles use the Reactive approach and all packets are configured with the same priority (DP2). Following Table I, the Reactive approach limits the packet rate to 10Hz for CBR levels between 30% and 40%. As a result, Reactive drops around 5% of packets in the low traffic density scenario, and reduces the maximum CBR experienced to around 35%. The percentage of packets dropped by DCC can be observed in Fig. 2. In the intermediate traffic density scenario, Reactive limits the packet rate to 4Hz-5Hz and reduces the CBR to around 40%, dropping more than 30% of packets (Fig. 2). In the highest traffic density scenario, the packet rate is limited to 1Hz most of the time and the CBR is reduced to 35%-50%. In this scenario, more than 50% of packets are dropped by DCC. The results in Fig. 1 therefore show how the Reactive approach drastically reduces the CBR through packet dropping, while a significant portion of the bandwidth remains free.

Fig. 1c shows how the Adaptive approach limits the maximum CBR experienced to 68%. In the low and intermediate traffic density scenarios, the CBR is most of the time below the target 68%. As a result, it does not drop nearly any packet in these two scenarios, as shown in Fig. 2. The Adaptive approach starts dropping a significant percentage of packets in the high traffic density scenario, because the CBR approximates to the target value (68%). In this scenario, more than 25% of packets are dropped by the Adaptive approach, as illustrated in Fig. 2. The results shown in Fig. 2 clearly show that a non-negligible percentage of packets can be dropped by DCC and therefore not transmitted.

To evaluate the performance of DCC, we use the PDR (Packet Delivery Ratio) metric at the radio and application levels. At the radio level, the PDR is defined as the probability of correctly receiving a message that has been effectively transmitted over the radio channel. It has been computed as the ratio between received and transmitted packets. The PDR at the

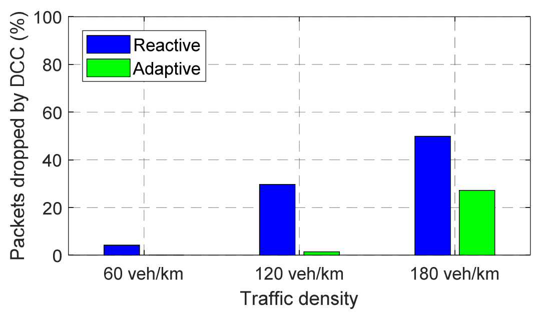

Fig. 2. Packets dropped by DCC for different traffic densities and DCC configurations, when all packets are configured with the same priority.

radio level is therefore influenced by the propagation and interference (i.e. packet collisions) effects. At the application level, the PDR is defined as the probability of correctly receiving a packet generated by the Facilities layer. It is hence calculated as the ratio between received and generated packets. The PDR at the application level is therefore influenced by the packets dropped by DCC, in addition to the propagation and interference effects. When a packet is dropped by DCC, it is not effectively transmitted, and represents a packet loss for the application.

Fig. 3 shows the PDR levels at the radio and application levels for the high traffic density scenario (180 veh/km) when all packets are configured as DP2. The results in Fig. 3a show that both Reactive and Adaptive can improve the PDR at the radio level compared to the scenario without DCC. However, Fig. 3b shows that, despite the improvements at the radio level (Fig. 3a), congestion control based on packet dropping actually degrades the performance at the application level. These results give reasons to doubt about the suitability of current packet dropping mechanisms for vehicular networks to control congestion. This degradation is produced because a very significant percentage of packets generated by the Facilities layer are not correctly received; they are simply dropped by DCC and not even transmitted. The degradation observed at the application level is especially higher for the Reactive approach, due to its higher percentage of packets dropped (Fig. 2) and lower PDR at the radio level. The lower PDR at the radio level is produced because vehicles tend to synchronize with each other when using Reactive, change their T_{off} simultaneously and transmit at the same time. As a result, Reactive generates a significant amount of collisions [5].

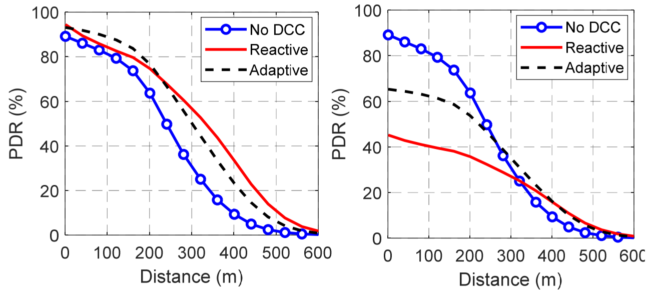

(a) Radio level (b) Application level

Fig. 3. PDR at the radio and application levels for different DCC configurations, when all packets have the same priority (DP2). Traffic density: 180 veh/km.

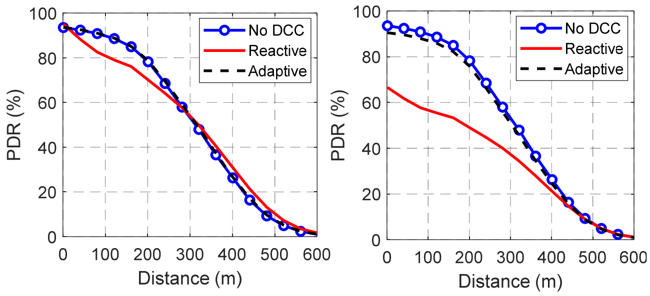

(a) Radio level (b) Application level

Fig. 4. PDR at the radio and application levels for different DCC configurations, when all packets have the same priority (DP2). Traffic density: 120 veh/km.

Fig. 4 shows analogous results, but considering a traffic density of 120 veh/km. In this scenario, almost no packets are dropped by the Adaptive approach (Fig. 2). As a consequence, nearly the same PDR curves (at the radio and application levels) are obtained with the Adaptive approach and without DCC. In this scenario, the Reactive approach still drops an important number of packets and therefore degrades significantly the PDR at the application level. In the low traffic density scenario (60 veh/km - figure not shown due to space limitations), almost no packets are dropped by DCC, and therefore all PDR curves overlap, irrespective of the traffic density and DCC approach. The obtained results for medium and low densities show that using DCC does not provide any benefit when evaluating the PDR at the application level.

C. *Packet prioritization*

We have also evaluated a scenario where CAMs and CPMs have different priorities. CAMs have been configured as DP2 (higher priority) and CPMs as DP3 (lower priority). Following the simple queuing mechanism described in Section II, a CPM is only de-queued by DCC and transmitted when there is no CAM in the queue. As a consequence, in the scenarios considered, DCC does not discard almost any CAM and only CPMs are discarded. This can be observed in Table II, that presents the percentage of packets discarded for the three traffic densities considered. The table shows that in scenarios with different packet priorities, packet dropping nearly exclusively affects the packets with lowest priority. In this study, this results in that up to 70% of the low priority packets are discarded, while almost no high priority packets are dropped.

TABLE II. PACKETS DROPPED BY DCC WHEN CAMS ARE CONFIGURED AS DP2 AND CPMS AS DP3

Traffic density	Reactive		Adaptive	
	CAM	**CPM**	**CAM**	**CPM**
60 veh/km	0.09 %	5.9 %	0.08 %	0.8 %
120 veh/km	0.05 %	40.3 %	0.03 %	2.2 %
180 veh/km	0.10 %	70.4 %	0.07 %	37.2 %

D. *Mixed traffic scenarios*

Vehicles can use either Reactive or Adaptive approaches. We have therefore extended the previous analysis with a scenario with a mix of vehicles using Reactive and Adaptive considering an intermediate traffic density of 120 veh/km and all packets configured as DP2. Fig. 5 shows the PDR at the application level considering a mix of vehicles using Reactive and Adaptive. These results demonstrate that the best performance is achieved when DCC is not used. We can also observe that the vehicles using Reactive experience a significantly lower application performance than the vehicles using Adaptive. This is mainly due to the higher percentage of packets dropped by Reactive. In fact, in mixed scenarios, the Adaptive approach does not drop almost any packet in any of the scenarios. This is the case because the vehicles using Reactive start dropping packets at lower CBR values than the vehicles using Adaptive, and maintain the CBR below the target of Adaptive. More than 90% of packets are dropped by

DCC Reactive in the scenario where only 25% of vehicles use Reactive and 75% use Adaptive. When the percentage of vehicles using Reactive increases, more vehicles drop packets (i.e. the ones using Reactive), and therefore each of them has to drop less packets to reduce the CBR. As a result, the PDR experienced at the application level by vehicles using Reactive improves as the percentage of vehicles using Reactive increases. However, the increase of this percentage degrades the PDR of the vehicles that use Adaptive. This is the case because they also tend to synchronize with the vehicles using Reactive. This synchronization provokes packet collisions, degrading the PDR at the radio and application levels. The results obtained show again that the best performance at the application level is achieved when DCC is not used. When DCC is used, the results demonstrate that the highest performance is achieved when no vehicle implements the Reactive approach and therefore discourage its implementation.

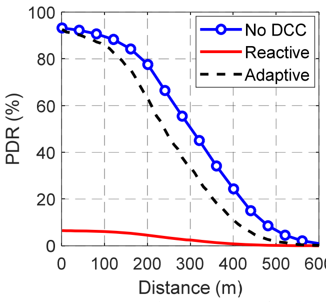

(a) 25% Reactive – 75% Adaptive

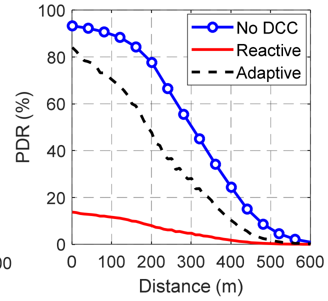

(b) 50% Reactive – 50% Adaptive

Fig. 5. PDR (Packet Delivery Ratio) at the application level for scenarios with a mix of vehicles using Reactive and Adaptive, when all packets are configured with the same priority (DP2). Traffic density: 120 veh/km.

IV. Conclusions and Dicsussion

This paper has demonstrated that congestion control based on current packet dropping can significantly degrade the performance of vehicular networks at the application level. This degradation is produced because packet dropping provokes that a significant portion of the packets generated by the applications are not transmitted and therefore lost. The results obtained considering DCC Access show that using congestion control based on packet dropping is worse than not using congestion control in all the considered scenarios (including congested ones). The obtained results suggest that future alternatives are needed, and this could influence the upcoming standardization processes. To reduce the channel load without affecting the application-level performance, a more intelligent generation of messages should be investigated to e.g. control redundant information and avoid the transmission of unnecessary packets [12]. In addition, the use of message compression was proposed and evaluated for the first time in [13] for vehicular networks. CAMs can be compressed between 4% and 14% approximately and higher compression gains could be achieved for longer packets or using more advanced data compression algorithms. An alternative to avoid reducing the packet rate while maintaining the CBR is the reduction of the transmission power and/or the increase of the data rate (modulation and coding scheme) [14]. Reducing the transmission power decreases the communication and interference ranges. Increasing the data rate decreases the packet duration, and also reduces the communication range. This reduction of the communication range to maintain the packet transmission rate can be a reasonable option from the application's perspective, especially for those vehicles driving slower (that do not require a long communication range). Further studies will be needed to analyze all these mechanisms as a complement or substitute of packet dropping for congestion control.

References

[1] "Intelligent Transport Systems (ITS); Decentralized Congestion Control Mechanisms for Intelligent Transport Systems operating in the 5 GHz range; Access layer part", TS 102 687 v1.2.1, ETSI, April 2018.

[2] N. Lyamin, A. Vinel, D. Smely and B. Bellalta, "ETSI DCC: Decentralized Congestion Control in C-ITS," *IEEE Communications Magazine*, vol. 56, no. 12, pp. 112-118, December 2018.

[3] A. Alonso Gómez et al., "Dependability of Decentralized Congestion Control for Varying VANET Density," *IEEE Transactions on Vehicular Technology*, vol. 65, no. 11, pp. 9153-9167, Nov. 2016.

[4] G. Bansal et al., "LIMERIC: A Linear Message Rate Control Algorithm for DSRC Congestion Control," *IEEE Transactions on Vehicular Technology*, vol. 62, no. 9, pp. 4182-4197, July 2013.

[5] A. Rostami et al., "Stability Challenges and Enhancements for Vehicular Channel Congestion Control Approaches," *IEEE Trans. on Intelligent Transportation Systems*, vol. 17, no. 10, pp. 2935-2948, Oct. 2016.

[6] M. Rondinone et al., "iTETRIS: a modular simulation platform for the large scale evaluation of cooperative ITS applications" *Simulation Modelling Practice and Theory*, vol. 34, pp. 99-125, May 2013.

[7] "ETSI DCC Access implementation for ns-3". Accessed: Feb. 2020. [Online]. Available: https://github.com/msepulcre/DCC-ns3

[8] V. Martinez and F. Berens, "Survey on ITS-G5 CAM statistics", CAR 2 CAR Communication Consortium, TR2052 V1.0.1, Dec. 2018. Accessed: Feb. 2020. [Online]. Available: https://www.car-2-car.org/fileadmin/documents/General_Documents/C2CCC_TR_2052_Survey_on_CAM_statistics.pdf

[9] G. Thandavarayan, M. Sepulcre and J. Gozalvez, "Analysis of Message Generation Rules for Collective Perception in Connected and Automated Driving", in *Proc. 30th IEEE Intelligent Vehicle Symposium*, Paris, France, 9-12 June 2019, pp. 134-139.

[10] "Intelligent Transport System (ITS); Vehicular Communications; Basic Set of Applications; Analysis of the Collective -Perception Service (CPS)", TR 103 562 v2.1.1, ETSI, Dec. 2019.

[11] "Study on LTE-based V2X Services", TR 36.885 V14.0.0, 3GPP, Technical Specification Group Radio Access Network, 2016.

[12] G. Thandavarayan, M. Sepulcre and J. Gozalvez, "Redundancy Mitigation in Cooperative Perception for Connected and Automated Vehicles", in *Proc. 91st IEEE Vehicular Technology Conference (VTC2020-Spring)*, Antwerp, Belgium, 25-28 May 2020.

[13] M. Sepulcre, P. Tercero, J. Gozalvez, "Can Beacons be Compressed to Reduce the Channel Load in Vehicular Networks?", in *Proc. IEEE Vehicular Networking Conference (VNC)*, Taipei, Taiwan, 5-7 Dec. 2018.

[14] M. Sepulcre et al., "Why 6Mbps is not (always) the Optimum Data Rate for Beaconing in Vehicular Networks", *IEEE Transactions on Mobile Computing*, vol. 16, no. 12, pp. 3568-3579, Dec. 2017.