Revisiting Message Generation Strategies for Collective Perception in Connected and Automated Driving

Quentin Delooz and Andreas Festag
Technische Hochschule Ingolstadt / CARISSMA
Ingolstadt, Germany
Email: Quentin.Delooz@carissma.eu
Email: Andreas.Festag@thi.de

Alexey Vinel
Halmstad University
Halmstad, Sweden
Email: alexey.vinel@hh.se

***Abstract*—Collective perception enables vehicles to exchange pre-processed sensor data and is being standardized as a 2nd generation V2X communication service. The European standardization in ETSI foresees the exchange of detected objects and defined a dedicated message type (Collective Perception Message, CPM) with rules to decide when and with which objects the message should be generated, referred to as generation rules. The choice of these rules is not straightforward and influences both channel load and perception quality. For the object inclusion, ETSI currently follows a similar policy as for the generation of Cooperative Awareness Messages (CAM): The objects are filtered based on their dynamics. We regard this approach as conservative. The present paper revisits the generation rules for the CPM and applies two approaches for object inclusion to the CPM – the conservative strategy of ETSI and a more 'greedy' strategy. We assess the performance by discrete-event simulations in a scenario representing a city with realistic vehicle densities and mobility patterns. The simulations take into account the effects imposed by decentralized congestion control. Considering that ETSI currently follows the conservative strategy, we conclude that the application of a greedy strategy improves the perception quality in low-density scenarios.**

***Keywords*–V2X; vehicular communications; collective perception; message generation.**

I. INTRODUCTION

Sensor data sharing using vehicle-to-everything (V2X) communications is an effective and low-cost solution to enhance the perception range of a vehicle's sensors. It is the basis for various advanced use cases for connected and automated driving. Recently, the European Telecommunications Standards Institute (ETSI) has completed a study item for sensor data sharing [1], named 'Collective Perception' (CP). CP is based on the periodic exchange of messages with the direct neighbours in communication range. The study item implies important design decisions including the definition of the Collective Perception Message (CPM) and features of the communication protocol towards the future standard.

The collective perception complements other communication services. Specifically, in the European system for V2X communications, the Cooperative Awareness (CA) service enables vehicles to report their position and driving dynamics to others through Cooperative Awareness Messages (CAMs) [2]. Similarly, the CPM carries objects lists, the vehicle's sensor configuration, and other data fields. All message types are transmitted in the bandwidth-limited wireless channels in the 5.9 GHz band allocated for road safety and traffic efficiency applications. Depending on the message frequency and the number of objects included, CPMs can considerably increase the channel load [3]. Decentralized Congestion Control (DCC) limits the overall data rate a vehicle is allowed to transmit over the wireless channel, but introduces additional delays or even drop messages under high channel load [4]. Following the ETSI study item for the CP service [1], the CP protocol defines several mechanisms to reduce the load generated by CPMs. Although DCC achieves the stabilization of the network, it can severely affect the performance of the CP service.

For the selection of objects to include in a CPM, i.e., *inclusion rules*, ETSI has adopted the strategy used for the CAM: By default, a CAM is broadcasted at a rate of 1 Hz. Then, depending on the vehicle dynamics (position, speed, and heading variation over time), the rate increases to up to 10 Hz [2]. In the case of the CPM, the reasoning is the same but applied to each object. Correspondingly to the CAM specification, the Society of Automotive Engineers (SAE) established the Basic Safety Message (BSM) for the DSRC system [5]. Though specified for the same purpose, the default BSM rate is 10 Hz, but independent of the dynamics of the vehicle. In this paper, we apply the BSM approach for the CPM object inclusion, i.e., a detected object will be transmitted at the rate of 10 Hz.

Both CAM and BSM address the trade-off between channel usage and message rate. The CAM generation can be seen as a *conservative* strategy as it uses the channel only when needed even though more transmission resources would be available. In contrast, the BSM generation rules will always send at the maximum rate if the DCC allows it. This strategy can be regarded as a *greedy* approach, which saturates the channel faster. However, in comparison to a conservative approach, it reduces the time between updates for an object.

In this paper, we present the design of the ETSI collective perception service [1] as a decomposition into components for message sending rules, object inclusion, and redundancy mitigation. Inspired by the BSM generation rules, we compare the 'conservative' strategy currently defined by ETSI for object inclusion rules with a 'greedy' approach. The evaluation relies on simulations using the OMNeT++- based ARTERY framework [6] and LIMERIC for DCC [7]. We consider a realistic scenario with urban, suburban, and highway traffic (*LuST* [8]) to evaluate the performance of both approaches.

The remainder of this paper is organized as follows: After reviewing existing work in Section II, we give an overview of ETSI collective perception in Section III and provide technical background on DCC and LIMERIC in Section IV. Section V describes our simulation environment and parameters used to assess the CP performance. Sections VI and VII provide an analysis of the obtained results and conclude the paper.

II. Related Work

Thandavarayan et al. [9] analyse two different policies, which define the object inclusion and message sending rules for the CP service. The *fixed* policy includes all the detected objects and CPMs are generated at a fixed rate. The *dynamic* policy filters the objects based on their dynamics, similarly to the sending rules of CAMs [2]. Additionally, if no object has to be transmitted, the generation of CPMs is omitted. The authors compare both policies in a highway scenario with different vehicle densities, all generating CPMs, but without considering DCC. Garlich et al. [10] analyse the same policies as in [9], but take into account DCC (a reactive approach, see Section IV), message sending rules and different channel configurations. The paper applies two different scenarios: a realistic (*LuST*) and an artificial one ('spider'). The authors of the two papers greatly contributed to the standardization process of the CP service.

Compared to our paper, [9] and [10] do not make a clear distinction between inclusion and sending rules. Specifically, the *fixed* and *dynamic* policies combine different inclusion rules with different sending rules, making the comparison hard to interpret. Additionally, the authors focus on the corner case where all vehicles send V2X messages; even considering a fast-growing rate of V2X-equipped vehicles, this should not happen before years. In the present paper, we make a clearly separate inclusion and sending rules. We analyse in deep the *LuST* scenario and focus on cases where the number of vehicles able to send CPM is low. Additionally, we consider LIMERIC [7], an adaptive DCC approach, which is more permissive compared to the reactive approach in [10].

III. Overview of collective perception standardized in ETSI

Based on the study item of ETSI about Collective Perception in [1], we decompose the CP service into components for triggering, inclusion, redundancy, and sending rules, which are periodically checked and subsequently executed (see Figure 1). We note that the component names do not correspond directly to the terms in [1], but our proposal eases the understanding of the mechanisms and their relationship. In addition, we do not consider the segmentation of CPMs as in [1]. Instead, if the size of the CPM is larger than the maximum message size of 1,100 B, we randomly remove objects from the message until the maximum message size is reached. Object removal in our scenario occurs rarely and can therefore be neglected. In the following, we explain each component.

A. Checking time

The checking time determines the frequency with which the rules are periodically inspected. It can be regarded as a sleeping time of the algorithm, i.e., the time duration in which a CPM cannot be generated. The value should be less than or equal to the minimum interval between two consecutive CPMs, i.e., 100 ms. Though [1] does not define a checking time, we can assume the same value as specified for the CA service [2] (see Section III-B). In addition, most of the existing research publications use a value of 100 ms.

B. Triggering rules

These rules define the time to wait between the generation of two consecutive CPMs. In [1], the lower and upper bound of the CPM transmission interval time is set to 100 ms and 1 s, respectively. DCC regulates the transmission rate of the CP service between these bounds. If DCC allows, CP triggers the generation of a CPM and set its content with the rules defined by the next components in Figure 1. Both *checking* and *triggering* rules are independent of the *conservative* and *greedy* policy.

C. Locally perceived environment

This component subsumes the pre-processed sensor data as a set of detected, tracked, and classified objects in a vehicle. In general, depending on its technical characteristics, each sensor type represents an object differently. However, using the CP service, the objects are represented in a standardized format, i.e., by their descriptions including position and speed relative to a reference position of the sending vehicle.

D. Inclusion rules

This component filters less relevant objects and these with a confidence level below a pre-defined threshold. The ETSI study item on collective perception [1] defines relevance criteria based on the objects' dynamics, type, and last transmission time. The criteria for object dynamics rely on the CAM generation rules [2], i.e., on the object's difference in position, speed, and heading since the last object inclusion. For the confidence level and threshold, we note that [1] does not define the parameter values.

Figure 2 depicts the decision tree that is executed to decide if an object should be included in the generated CPM, or not. We stress that these inclusion rules rely only on the perception of the sending vehicles, i.e., the objects received via V2X communication are not taken into account. Furthermore, the inclusion rules are static and are applied irrespective of the channel load or the vehicle's driving situation.

In the present paper, we compare the inclusion rules defined in [1], i.e., the *ETSI rules*, with an approach where all objects detected are included in the generated CPM, i.e., *no-filtering*. In analogy to the channel usage vs. message generation trade-off discussed in Section I, the ETSI inclusion rules are seen as *conservative* and non-filtering as *greedy*.

E. Redundancy mitigation rules

A redundant transmission occurs when the same object is received multiple times from different senders. The redundancy mitigation rules omit the transmission of objects which were already received. In [1], these rules are only applied if the channel load is larger than a (still undefined) threshold. Additionally, [1] proposes several strategies. For example, the frequency-based approach omits locally perceived objects from the new CPM if a certain number of previously sent CPMs in a given time window already included information about the same objects. As proposed by [3], some of these rules could also be considered as *inclusion rules*.

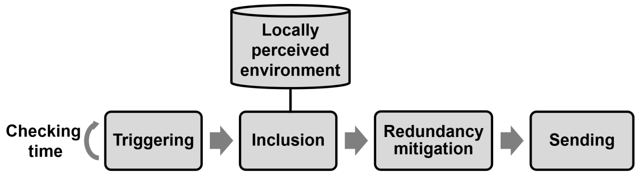

Figure 1. Components for the message generation in ETSI collective perception derived from [1] (without segmentation).

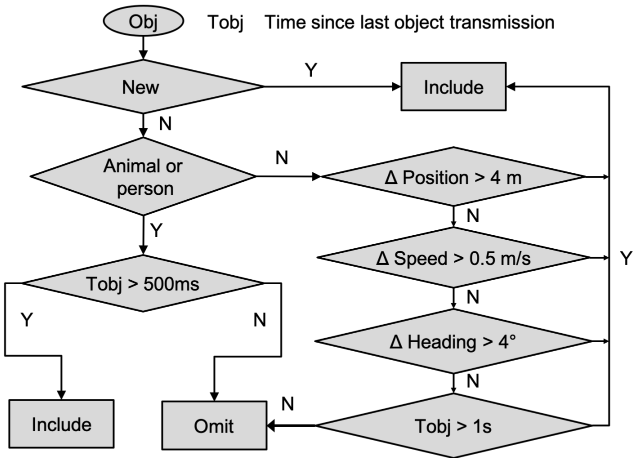

Figure 2. Rules for object inclusion as defined in [1], which corresponds to the *conservative* strategy in this paper.

A typical scenario where the redundancy mitigation rules would be effective is a road intersection. In that case, potentially many vehicles detect the same object, e.g., a pedestrian, and start transmitting information about it. The added information brought by each vehicle would be small in comparison to the 'cost' of its transmission, especially in areas with a dense number of vehicles.

F. Sending rules

The sending rules represent the last decision point to decide if the generated CPM should be sent to the lower levels of the protocol stack. Following [1], at least one of these three conditions should be respected for the CPM to be generated: *(i)* there is at least one object to send, *(ii)* the last CPM with sensor information and *(iii)* the duration since the generation of the last CPM is at least 1 second.

In [10], the authors compare the non-filtering and the ETSI inclusion rules. However, the authors apply different sending rules for the inclusion strategy: With the ETSI inclusion rules, a CPM is generated if objects are present to be transmitted. For the non-filtering approach, a fixed rate of 10 Hz is applied, independently whether objects have to be transmitted. Therefore, it is hard to analyse the effect of the inclusion rules and the sending rules separately, especially if the sensors' parameters do not allow vehicles to detect always at least one object. In the present study, CPMs will only be generated if there is information, such as objects or sensors data, to send, and independently of the used inclusion rules.

IV. Decentralized Congestion Control (DCC)

In the following, we provide background information on DCC in general in Section IV-A and specifically on LIMERIC in Section IV-B.

A. General

DCC is a set of mechanisms in the V2X protocol stack that ensure the stability of the network and fairness in resource usage among network nodes. Its principal function is to measure the channel load (channel busy ratio, CBR) and to control the data that a station generates. DCC is standardized by ETSI in several standards.

DCC is a cross-layer functionality with interacting entities at different layers. The access layer functionality [11] provides traffic shaping for the injected packets. Practically, it implements a 'gatekeeper' that realizes a First-In-First-Out (FIFO) queuing system for each channel. A gatekeeper has multiple queues for the packets to be sent and a single server, which dispatches always the non-empty queue with the highest priority (simple priority queue). When a packet enters the gatekeeper, and the queue is not full, DCC allows the transmission of the packet and sets its transmission parameters. If the queue is full or the lifetime of the packet expires during the waiting time in the queue, the packet is discarded.

To determine when a packet can be transmitted to the MAC layer, ETSI standardized two types of strategies for the gatekeeper [12]: reactive and adaptive. Both strategies respect the DCC requirements specified in [13]:

- $0 < T_{on} < 4ms$: T_{on} is the maximum duration of a packet transmission.
- duty cycle $<= 3\,\%$: it means that a station can occupy at most 3%, i.e., 30 ms, of channel time.
- $T_{off} >= 25ms$: T_{off} is the duration before the gatekeeper re-opens after the transmission of packet and allows a new packet to be transmitted. In other words, the maximum packet transmission frequency is 40 Hz.
- if $CBR >= 0.62$, $T_{off} >= 1{,}000ms$.

The reactive approach defines a set of states for which values of the T_{off} time are assigned to specific CBR thresholds. The higher the measured CBR, the longer a station needs to wait between two consecutive transmissions. [12] proposes two sets of states, each one depending on the maximum allowed transmission time. Effectively, the reactive approach sets a predefined rate based on the measured CBR. In contrast, an adaptive method shares the channel resources between the stations in communication range such that the CBR converges to a predefined maximum value. The LInear MEssage rate Integrated Control (LIMERIC) algorithm [7] meets the ETSI requirements for the adaptive DCC approach and is used in the simulation of the present paper.

B. LIMERIC

Instead of directly adapting the transmission rate, LIMERIC adjusts the duty cycle δ every 200 ms. The duty cycle is the allowed ratio of the transmitter total "on" time relative to 1 s. [12] defines the algorithm to adapt δ depending on the observed CBR. [14] provides some insights about the reason behind the chosen LIMERIC parameters and proposes a dual-α approach to improve LIMERIC's convergence time

 ISBN: 978-1-61208-795-5

TABLE I. SUMMARY OF LIMERIC'S PARAMETERS

Parameter	Description	Values
α_{low}	Convergence parameter	0.016
α_{high}	Convergence parameter	0.1
th	To choose between the α	0.00001
α	Convergence parameter	as in [14]
β	Convergence parameter	0.0012
CBR_{target}	Convergence point	0.68
δ_{max}	Max allowed duty cycle	0.03
δ_{min}	Min allowed duty cycle	0.0006
δ_{init}	Initial δ	0.0153
G^{+}_{max}	Upper born used to update δ	0.0005
G^{-}_{max}	Lower born used to update δ	–0.00025
T_{CBR}	Interval for CBR value update	100 ms

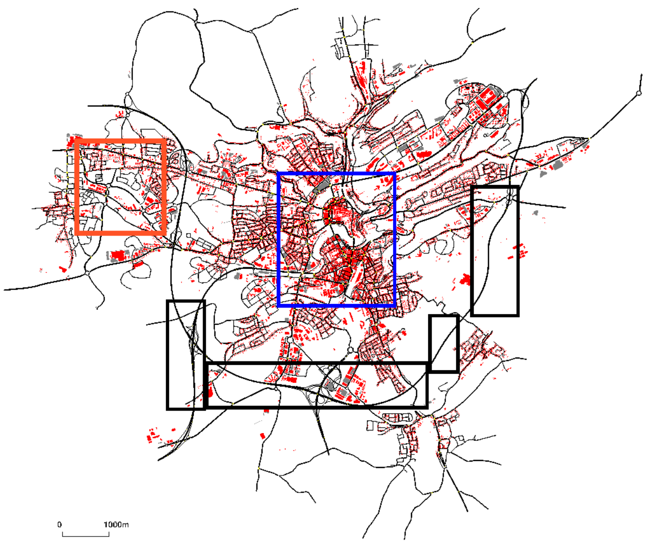

Figure 3. Topology of Luxembourg in *LuST*.

and fairness during transition phases. We decided to use this modification. Table I summarizes the LIMERIC parameters used in our simulations.

From the allowed duty cycle determined by LIMERIC, [12] derives T_{off} to enforce the rate by

$$T_{off} = \min(\max(\frac{Ton_{pp}}{\delta}, 25\,\text{ms}), 1\,\text{s}) \qquad (1)$$

with Ton_{pp} being the transmission time of the last transmitted packet. We note that the reactive strategy considers only the CBR and makes some simple assumptions for the packet size. In contrast, the adaptive strategy takes into account the size of the transmitted packet to enforce the allowed duty cycle.

V. SIMULATION ENVIRONMENT

This section presents the used simulation framework, the V2X services deployed and their respective message formats, and how the vehicles are equipped in the simulations.

A. Simulation framework

For the evaluation of the CPM generation strategies, we used the discrete-event simulator ARTERY [6] to model the V2X communications following ETSI standards. *ARTERY* relies on VANETZA, INET and OMNeT++ (v5.4.1), and implements the V2X protocol stack based on ITS-G5 (see [6] for details). To model the traffic and mobility of the vehicles, we used the microscopic road traffic simulator *SUMO* (v1.0.1) [15] with the popular Luxembourg scenario, a.k.a. *LuST* [8]. Figure 3 shows the topology of the SUMO map for the Luxembourg scenario (see [3] for the distribution of vehicles). We note that the LuST scenario was validated with real mobility data for SUMO version 0.26. Since we have used a newer version of SUMO, the traffic mobility model cannot be regarded as formally validated but still represents a realistic scenario.

Each simulation run is executed for a duration of 13 s with 10 s of warmup. The warmup phase gives time to LIMERIC to converge to the desired δ for each vehicle.

For the *LuST* scenario, we have chosen a snapshot at 8 a.m. This corresponds to a rush-hour with around 5,000 vehicles in the simulated environment. Within the scenario, we selected three distinct areas: urban, suburban, and highways, respectively represented in Figure 3 by the blue, orange, and black squares. In terms of vehicle density, the *urban* area will face the highest density and the *suburban* the lowest. Table II shows the vehicle dynamics depending on the area; the dominant dynamic parameter are marked in bold. For example, for the *highway* area, the dominant parameter is Δ position, i.e., the vehicle speed, because it triggers the generation of a CAM first. Following the vehicle dynamics parameters for CAMs [2], the theoretically resulting CAM transmission rate is presented in the last row of Table II. For the urban area, the two parameters Δ position and Δ speed result in an almost equal CAM transmission rate. Therefore, we indicated both as dominant in Table II.

B. V2X services and CPM format

Both CA and CP services are enabled. The CA service operates on the Control Channel (CCH) and the CP service on the Service Channel 1 (SCH1) of the 5.9 GHz frequency band. We consider that the vehicles can receive and send at the same time on both channels and that there is no interference between them. The fading model used is the one integrated into ARTERY called *VanetNakagamiFading* [16].

The CAM and the CPM formats rely on [1] and [2], respectively. Specifically, the CPM consists of an ITS PDU header and several containers, including containers for management and station data containing information about the sender such as position, heading and velocity, 0 to 127 Sensor Information Containers (SICs), and 0 to 127 Perceived Object Containers (POCs). Using default values, the size of a SIC varies from 11 to 88 bytes and a POC from 20 to 46 bytes. The CPM format is specified in ASN.1 and encoded by the Unaligned Packed Encoding Rules in ASN.1 as specified by ETSI (see Section 6.8.3 and Annex A of [1]). We have used the ASN.1 open-source compiler *asn1c*.

TABLE II. AVERAGE CHANGES OF THE VEHICLE DYNAMICS DURING A 100 ms TIME INTERVAL IN THE *LuST* SCENARIO

	All	Urban	Suburban	Highway
Δ position $[m]$	**1.415**	**0.777**	**0.996**	**2.47**
Δ speed $[m/s]$	0.071	**0.097**	0.102	0.042
Δ heading [°]	0.006	0.006	0.01	0.007
CAM frequency (Hz)	3.54	1.94	2.49	6.175

 ISBN: 978-1-61208-795-5

TABLE III. SUMMARY OF THE SIMULATION PARAMETERS

Parameter	Values
Protocol stack	ITS-G5
Frequency band	5.9 GHz
Channel number	SCH1 (176) for CP service
(IEEE numbering scheme)	CCH (180) for CA service
Channel model	*VanetNakagamiFading*
DCC	LIMERIC
Inclusion rules	{Etsi, No-filtering}
Scenarios	*LuST*
PVE	{10, .., 90, 100}
Time of simulation	8 a.m.
Number of vehicles	≈ 5,000
Simulation time	13 s (incl. 10 s of warmup)
Number of repetitions	2
Vehicle sensor equipment	{60 & 174 m, +/-10 & +/-45°}
	{150 m, 360°}

C. Vehicle equipment and object detection

It is assumed that with the increasing deployment of C-ITS, the ratio of vehicles equipped with V2X technologies will grow over time. The larger the ratio, the higher gets the generated data load on the channel. To analyse the impact of the V2X equipment rate on the performance of the filtering approaches, we varied the V2X equipment rate (PVE = Percentage of Vehicles Equipped) and used the values PVE = {10 20, ..., 100} %.

For object detection, the vehicles have local sensors mounted on them and we used two different configurations. In the first one, each vehicle has two radars with respectively a range of 60 and 174 m, and a field of view of +/– 45 ° and +/– 10 °, respectively. Both radars are located in front of the vehicle and are facing ahead. In the second configuration, each vehicle is equipped with a radar with a range of 150 m, and a FOV of 360 °. The first configuration simulates the early development of sensor perception. The second one grossly reproduces the future perception capabilities of vehicles.

The method to detect objects is the same as explained in [17]. In brief, each ITS-S mounted with sensors detects an object if one of the four corners of the object is in the line of sight of one of the sensors. The information retrieved from the perception is idealistic, i.e., all object attributes are always available and no errors in object detection occur.

VI. Performance Evaluation

The evaluation compares the performance of the conservative and the greedy strategy for object inclusion in the CP service for different values of the PVE. In the evaluation, we also vary the area type (urban U and highway H) and the sensor vehicle equipment (Field of view of {+/-10 & +/-45°} and 360°). For example, the simulation *U-CPM-conservative (360)* corresponds to the *conservative* strategy for object inclusion in the *urban* scenario and with a vehicle sensor configuration for an FOV of *360°*. We collected different metrics to assess distinct aspects of the *CPM-conservative* and *CPM-greedy* inclusion rules. For readability reasons, we only include results for the urban and highways areas, which cover most of the interesting points to discuss.

A. Network-related metrics

The *Channel Busy Ratio* (CBR) gives a measure of the channel occupancy. The *Packet Error Rate* (PER) is the rate of unsuccessfully decoded messages on the number of received ones. The higher the CBR, the larger the PER is likely to be. Figure 4a shows the CBR on the SCH1, i.e., the CBR obtained with the CP Service, for the urban (U) and highways (H) areas in the *LuST* scenario. In general, the CBR for the urban area is higher than for the highway one. The same applies to Figure 4b showing the PER obtained for the same scenarios. The highest CBR is observed with the *greedy* approach and a sensor with a FoV of 360°. With this configuration, the PER is around 25 % at PVE=100 %. Still, the CBR does not reach the targeted CBR defined by LIMERIC, and the channel is not saturated. The ETSI configuration does not generate sufficient data for a CBR higher than 0.3. The maximum average observed PER obtained for ETSI is around 25 % with the *CPM-greedy (360)* configuration in the urban scenario.

B. Application-related metrics

Figure 5 shows the results obtained for the *Number of Objects Detected* (NOD) and the *Time Between Update* (TBU) metrics. The NOD metric represents how many objects a vehicle was aware of during the last second. The TBU metric expresses the average time between two consecutive updates of the same object. Both CAMs and CPMs contribute to these metrics. In the urban scenario and with the 360° sensor, the *greedy* approach provides the highest number of NOD for any PVE. Interestingly, with this configuration, when the PVE reaches 50%, the NOD does not vary significantly anymore. We could question the necessity to have more vehicles sending CPMs if we can assert that CPM transmitter distribution is uniform. Additionally, the *U-CPM-conservative (360)* and the *U-CPM-greedy* have comparable results. The urban area with ETSI inclusion rules has the highest TBU. However, the updating rate is always less than if only the CA service would be used (see Table II). The smallest TBU, around 50 ms, is obtained in the highway area, a 360° sensor, and the *greedy* inclusion rules. The difference between *greedy* and *conservative* is more important in the urban than for the highway. This is expected since in highway scenarios, vehicles exhibit higher dynamics than in the urban area (see Table II).

C. Ratio of Resource Used (RRU)

Figure 6 shows the average *Ratio of Resources Used* (RRU) by a CPM. Following (1), the RRU is defined by Ton_{pp}/δ. For example, if $\delta = 0.001$ and the transmission time of a CPM is 200µs, then the RRU taken by this CPM will be 0.2. This means that the transmitted CPM takes 20% in a 1 s time duration of the channel access time determined by LIMERIC. If the RRU is lower than 0.1, DCC always allows the CP service to generate CPMs at the maximum rate of 10 CPM/s. Except for the *U-CPM-greedy (360)* configuration, the RRU is in average constant, independently of the PVE. This can be explained by two reasons: the average size of CPM for each configuration is the same, independently of the PVE. The second reason is the convergence time of LIMERIC when the parameter δ increases. Indeed, even with 10 s of warmup, corresponding to 50 updates of δ, its highest possible value would be around 0.023 from δ_{init}. This is not a problem as the resulting T_{off} is always shorter than 100 ms. With the dual-alpha approach [14], δ converges faster when decreasing.

For the *U-CPM-greedy (360)*, the δ decreases enough to observe the RRU increasing. It means that with more

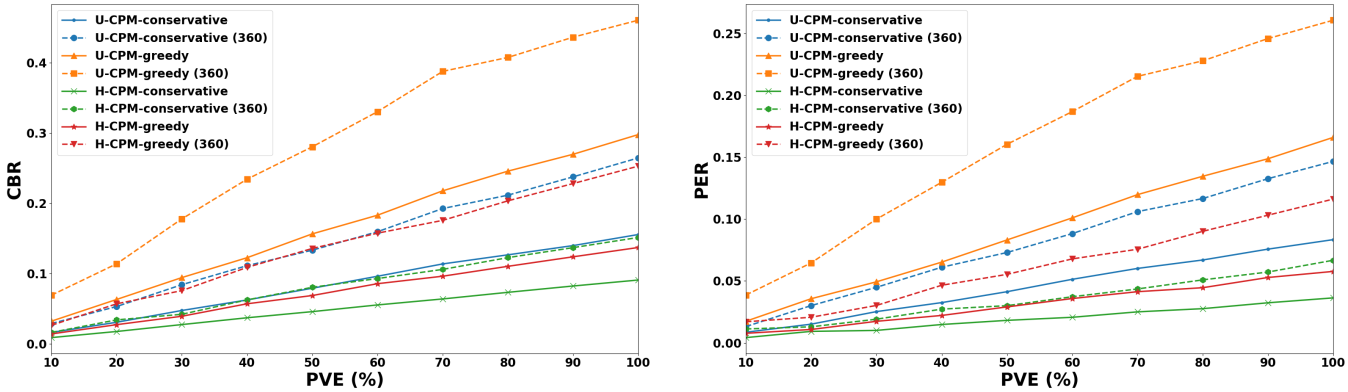

(a) Average Channel Busy Ratio (CBR)

(b) Average Packet Error Rate (PER)

Figure 4. Network-related metrics for different values of percentage of vehicles equipped (PVE) with V2X capabilities in the *LuST* scenario.

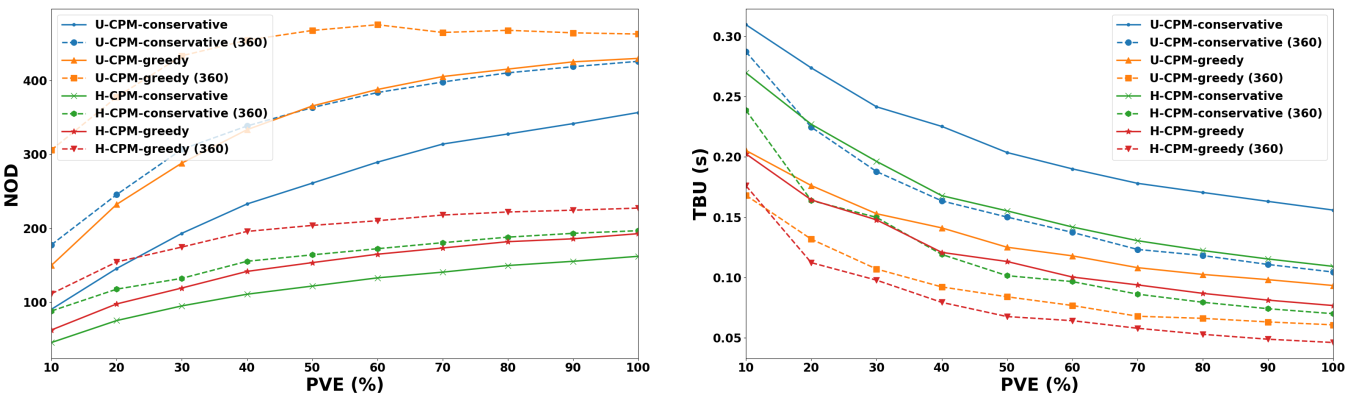

(a) Average Number of Objects Detected (NOD)

(b) Average Time Between Updates (TBU)

Figure 5. Applications-related metrics for different PVE values in the *LuST* scenario.

transmitting vehicles and objects to detect, LIMERIC would only start reducing the transmission rate of the vehicles. Only in the configuration *U-CPM-greedy (360)*, the RRU increases with the PVE. Still, in average the RRU remains under 0.1. Therefore, the CP service can generate 10 CPM/s.

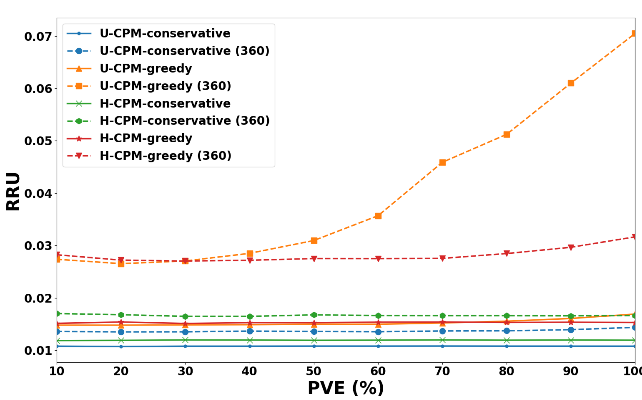

Figure 6. Average Ratio Resource Used (RRU).

D. Discussion

Even if not all kinds of objects, such as pedestrians or obstacles, are present in the simulations, both scenarios provide insights about the trade-off between channel load and perception quality. In the *LuST* scenario, the CP service is not able to saturate the channel even without filtering and with a PVE of 100 %. Notably, the *conservative* inclusion rules underutilize the channel resources, while largely available, at low PVE and in areas with a small density of vehicles. Compared to *conservative*, the *greedy* approach provides always a better perception quality while not saturating the channel.

VII. Conclusion and Future Work

In this paper, we addressed the CP service as currently defined by ETSI and presented it as a decomposed system with message sending rules, object inclusion, and redundancy mitigation. For the object inclusion rules, we analysed two strategies, i.e., *conservative* and *greedy*. In comparison, the *greedy* strategy, which does not filter objects, provides a reduced time between updates and a higher number of perceived objects at the cost of higher channel usage. Still, in the considered scenario the target CBR of LIMERIC is never reached. Also, DCC does not have to reduce actively the message rate. While both *greedy* and *conservative* approach differently consider the trade-off between the perception quality and channel usage, at a low percentage of equipped vehicles and in areas with a small density of vehicles, the *greedy* approach allows for a higher channel utilization and for better performance.

In our future work, we will study the improvement of the CP service when the inclusion rules dynamically adapt to the channel load. This approach smoothly combines the *greedy* and *conservative* approach: When the channel load is low, fewer objects are filtered and, the quality of perception is increased. When the channel is close to saturating, we switch to the conservative approach and filter more objects. We will also consider other filtering approaches, such as redundancy mitigation rules, to reach the same goal.

ACKNOWLEDGMENT

This work was supported by the German Research Foundation DFG within the priority program Cooperative Interacting Cars (CoInCar) (SPP 1835) and in part by the Swedish Knowledge Foundation within SafeSmart ”Safety of Connected Intelligent Vehicles in Smart Cities” Synergy Project (2019–2023). The authors would like to thank H. Güenther (Volkswagen), R. Riebl (Technical University Ingolstadt), K. Garlich and L. Wolf (Technical University Braunschweig) for valuable discussions.

REFERENCES

[1] ETSI, “Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Analysis of the Collective Perception Service (CPS); Release 2,” Dec. 2019, ETSI TR 103 562 V2.1.1.

[2] ETSI, “Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Part 2: Specification of Cooperative Awareness Basic Service,” Apr. 2019, ETSI EN 302 637-2 v.1.4.1.

[3] Q. Delooz and A. Festag, “Network load adaptation for collective perception in V2X communications,” in 2019 IEEE International Conference on Connected Vehicles and Expo (ICCVE), 2019, pp. 1–6.

[4] S. Kühlmorgen et al., “Evaluation of congestion-enabled forwarding with mixed data traffic in vehicular communications,” IEEE Transactions on Intelligent Transportation Systems, vol. 21, no. 1, 2020, pp. 233–247.

[5] J. B. Kenney, “Dedicated Short-Range Communications (DSRC) standards in the United States,” Proc. IEEE, vol. 99, no. 7, Jul. 2011, pp. 1162–1182.

[6] R. Riebl, H. Günther, C. Facchi, and L. Wolf, “Artery: Extending Veins for VANET applications,” in 2015 International Conference on Models and Technologies for Intelligent Transportation Systems (MT-ITS), pp. 450–456.

[7] G. Bansal, J. B. Kenney, and C. E. Rohrs, “LIMERIC: A linear adaptive message rate algorithm for DSRC congestion control,” IEEE Trans. Veh. Technol., vol. 62, no. 9, Nov. 2013, pp. 4182–4197.

[8] L. Codeca, R. Frank, S. Faye, and T. Engel, “Luxembourg SUMO traffic (LuST) scenario: Traffic demand evaluation,” IEEE Intell. Transp. Syst. Mag., vol. 9, no. 2, 2017, pp. 52–63.

[9] G. Thandavarayan, M. Sepulcre, and J. Gozalvez, “Analysis of message generation rules for collective perception in connected and automated driving,” in 2019 IEEE Intelligent Vehicles Symposium (IV), 2019, pp. 134–139.

[10] K. Garlichs, H. Günther, and L. C. Wolf, “Generation rules for the collective perception service,” in 2019 IEEE Vehicular Networking Conference (VNC), 2019, pp. 1–8.

[11] ETSI, “Decentralized Congestion Control Mechanisms for Intelligent Transport Systems operating in the 5 GHz range; Access layer part,” Apr. 2018, ETSI TS 102 687 V1.2.1.

[12] ETSI, “Intelligent Transport Systems (ITS); Decentralized Congestion Control Mechanisms for Intelligent Transport Systems operating in the 5 GHz range; Access layer part,” Apr. 2018, ETSI TS 102 687 V1.2.1.

[13] ETSI, “Intelligent Transport Systems (ITS); Radiocommunications equipment operating in the 5 855 MHz to 5 925 MHz frequency band; Harmonised Standard covering the essential requirements of article 3.2 of Directive 2014/53/EU,” Feb. 2017, ETSI EN 302 571 V2.1.1.

[14] I. Soto, O. Amador, M. Urueña, and M. Calderon, “Strengths and weaknesses of the ETSI adaptive DCC algorithm: A proposal for improvement,” IEEE Commun. Lett., vol. 23, no. 5, May 2019, pp. 802–805.

[15] P. A. Lopez et al., “Microscopic traffic simulation using SUMO,” in 2018 21st International Conference on Intelligent Transportation Systems (ITSC), 2018, pp. 2575–2582.

[16] L. Cheng, B. E. Henty, D. D. Stancil, F. Bai, and P. Mudalige, “Mobile vehicle-to-vehicle narrow-band channel measurement and characterization of the 5.9 GHz dedicated short range communication (DSRC) frequency band,” IEEE J. Sel. Areas Commun., vol. 25, no. 8, Oct. 2007, pp. 1501–1516.

[17] H. Günther, O. Trauer, and L. Wolf, “The potential of collective perception in vehicular ad-hoc networks,” in 2015 14th International Conference on ITS Telecommunications (ITST), 2015, pp. 1–5.

 ISBN: 978-1-61208-795-5