Generation Rules for the Collective Perception Service

Garlichs, Keno and Hendrik-Jörn Günther, and Wolf, Lars C.

Authors post-print published on 2019-12-13
Originally published in *IEEE Vehicular Networking Conference* (VNC)
Publisher version available at TBD
DOI: TBD

Abstract:

The Collective Perception service is currently being standardized in Europe. Besides the message format, very im- portant design aspects are the generation rules governing when to generate a message and which data to include. These rules can heavily affect both the provided service quality for receiving Intelligent Transport System stations and the generated network load on the employed radio channel. This paper contributes a set of generation rules for the Collective Perception Service that provides a trade-off between these two factors. Their effectiveness is detailed in an extensive simulation study. Depending on specific radio channel and Decentralized Congestion Control settings, the simulations also reveal co-existence challenges with other services and messages disseminated on the same channel, e.g., the legacy Cooperative Awareness Message. The latter findings are thereby generalized as they apply to any additional message that needs to be disseminated along with the legacy Vehicle-to-X messages.

Generation Rules for the Collective Perception Service

Keno Garlichs*, Hendrik-Jörn Günther†, Lars C. Wolf*
*Technische Universität Braunschweig, Germany
Institute of Operating Systems and Computer Networks
{garlichs, wolf}@ibr.cs.tu-bs.de
†*Volkswagen Group of America, USA*
hendrik-joern.guenther@vw.com

***Abstract*—The Collective Perception service is currently being standardized in Europe. Besides the message format, very important design aspects are the generation rules governing when to generate a message and which data to include. These rules can heavily affect both the provided service quality for receiving Intelligent Transport System stations and the generated network load on the employed radio channel. This paper contributes a set of generation rules for the Collective Perception Service that provides a trade-off between these two factors. Their effectiveness is detailed in an extensive simulation study. Depending on specific radio channel and Decentralized Congestion Control settings, the simulations also reveal co-existence challenges with other services and messages disseminated on the same channel, e.g., the legacy Cooperative Awareness Message. The latter findings are thereby generalized as they apply to any additional message that needs to be disseminated along with the legacy Vehicle-to-X messages.**

***Index Terms*—Connected Vehicles, Collective Perception, MCO**

I. INTRODUCTION

The concept of Collective Perception in the context of Vehicle-to-X (V2X) communications refers to the idea of broadcasting information about objects detected by local perception sensors mounted to Intelligent Transport System Stations (ITS-Ss) [1]. Broadcasting of the already standardized Cooperative Awareness Message (CAM) [2] enables ITS-Ss to share detailed state information about themselves only, but not about their observations. Collective Perception aims at also sharing information about (non-connected) objects detected by a disseminating ITS-S. The concept enables novel safety applications based on received information about objects located outside of the range of a station's perception sensors. In combination with roadside mounted sensors, the concept also allows to realize applications warning oncoming traffic of the presence of, e.g., Vulnerable Road Users (VRUs) [3] in hidden line-of-sight conditions. With ongoing standardization in both the European Telecommunications Standards Institute (ETSI) [4], [5] and SAE [6], Collective Perception has gained increased attention in research and V2X industry.

Collective Perception Messages (CPMs) will be generated at the facility layer of the communications stack in a CP service entity. As such, the CPMs are competing for channel access with messages from other services operating on the same channel (e.g. the Cooperative Awareness (CA) service). Depending on the employed communications technology on an ITS-S, lower layer operations may thereby defer message transmission or even drop messages, if the observed channel utilization does not allow an ITS-S to transmit the queued messages. In the context of the Intelligent Transport System (ITS)-G5 stack, e.g., Decentralized Congestion Control (DCC) [7] mechanisms implement a gatekeeper functionality ensuring that the generated channel utilization of a particular ITS-S stays within defined limits [8]. The gatekeeper thereby takes message prioritization and defined inter-message transmission deferrals [9] into account in order to control the generated channel utilization regardless of the specific message content. It is therefore the task of the service generating the message, to ensure that relevant data is included in the message to be transmitted, while aiming at reducing the message size to help reduce the resulting channel utilization.

This paper proposes a CPM generation process that aims at addressing this trade-off. Based on the concept employed by the CA service, detected objects are included in a message to be transmitted based on their dynamics and data age with respect to the last transmitted CPM. This procedure reduces both the message generation frequency and the resulting encoded message size, as not all objects need to be included in every message. In case of complex perception sensor setups (e.g. 360°-sensors or sensor clusters at intersections), the likelihood of perceiving a very large number of objects increases. The proposed generation rules therefore introduce a message segmentation mechanism that allows for generating consecutive independent CPMs. This contribution assesses the proposed generation rules in an extensive simulation study gauging their behavior in different traffic scenarios and radio configurations. Since Multi-Channel Operation (MCO) is currently considered for future ITS operations [10], this study includes an MCO variant, disseminating the CPM on a separate channel.

The remainder of this paper is structured as follows: Section II introduces related approaches and provides further information about the concept of CP. The proposed generation rules addressing the aforementioned trade-off are detailed in Section III. The assessment and analysis of these rules as part of an extensive simulation study are presented in Section IV. Section V concludes the results and discusses follow-up work and further refinements.

II. RELATED WORK

To our knowledge, the first concepts for Collective Perception have been introduced as part of the Ko-FAS project which aimed at realizing applications targeting non-line-of-sight scenarios. The messages and concepts developed as part of this projects focused on sensor data fusion mechanisms [11], and the analysis of communication related parameters [12] when sharing sensor data using a broadcast mechanism. Both aspects combined create challenges in the context of V2X communications, as localization accuracy largely depends on Global Navigation Satellite System (GNSS) characteristics at a certain location and thereby complicate the data fusion process, when combining received remote sensor data. Furthermore, updates about detected objects of a particular transmitter are received much less frequently compared to locally mounted sensors [13], as is assumed by traditional sensor fusion approaches. The proposed message format of this project, however, was not adaptive to the current perception situation and created a large overhead unsuitable for standardization purposes.

The effect of an additional message on an ITS-G5 based V2X communications system has been further studied by the authors of [1] and [14]. Their conducted simulation studies analyze the challenges associated with fitting either an additional message into the Control Channel (CCH) or into one of the available Service Channels (SCHs). The results indicate that with increasing market penetration of V2X communications, the required channel capacity increases significantly when transmitting both their proposed Environment Perception (EP) message [15] and the already standardized CA message on the same channel. In addition, it is shown that DCC operations lead to further message drops, thereby affecting not only the additional EP message but also the legacy CAM transmission. The selection of message priority levels and message generation frequency mechanisms have a significant impact on the effectiveness and the feasibility of sending an additional message besides the legacy CAM. Their work, however, tied the generation of the proposed EP message to the generation rules of the CAM instead of adapting message generation to the characteristics of perception scenarios. This results in simultaneous injection of both the CA and EP message for transmission in the DCC queues, and therefore, message drops at the gatekeeper.

The generation rules which we proposed within the ETSI standardization and discuss in this paper therefore aim at exploiting the different dynamic natures of detected objects in traffic scenarios to reduce both the number of objects included in a CPM and the generation frequency, while maintaining a sufficient update rate suitable for sensor fusion algorithms. The authors of [16] employ the generation rules initially developed in our contribution and assess complementing metrics in a simplified highway scenario. Our contribution assesses the proposed rules both in more complex and realistic traffic scenarios along with the impact of DCC and MCO operations.

In parallel to the development of suitable message generation rules for the CPM, related research also focuses on data redundancy mitigation techniques further reducing message generation frequency: In case multiple ITS-Ss are able to perceive the same object with their local perception sensors, the authors of [17] propose a mechanism in which each transmitting ITS-S maintains a representation of the assumed tracking state of a particular detected object in surrounding ITS-Ss's environment models, i.e. a vehicle's internal representation of its current driving environment. A transmitting station would then only include the detected object in a CPM in case a predetermined assumed quality threshold of the representation of this object in another station's environment model is exceeded. Stations also detecting and tracking the same object are thereby matching objects from received CPMs and omit inclusion of this object in consecutive CPMs as long as their quality estimation of this object's representation in surrounding ITS-Ss' environment models does not exceed the aforementioned threshold. However, resulting challenges regarding false object association and its consequences such as falsely omitted objects from generated CPMs are not addressed and subject to further research.

III. GENERATION RULES FOR THE CP SERVICE

Ongoing standardization work at ETSI is developing the dedicated CP message which provides a modular format employing several containers for detailing information about the state of the disseminating ITS-S, its perception capabilities and detected objects [4], [5]. The proposed CPM format consists of four separate containers to transport information about detected objects: The mandatory *Management Container* includes information about the position and type of the transmitting ITS-S. Depending on the former, more detailed information about the dynamic state (e.g., for vehicles) or references to mounting positions on MAP-data (e.g., for infrastructure-mounted sensors) may be provided as part of the *Station Data Container*. An optional *Sensor Information Container (SIC)* details the sensory capabilities of the disseminating ITS-S by outlining the available perception ranges and sensor mounting positions. All detected objects and obstacles are encoded in a *Perceived Object Container (POC)*, whereas each object can be detailed by a variable set of parameters.

As for the related work presented above, rather than sharing raw-sensor data, the standard development focuses on transmitting analyzed perception data which has been processed by a high-level data fusion framework for object detection and tracking. The modular message format requires a distinct set of rules governing the inclusion of the available containers, when generating a CP message. On the one hand, CP messages would only need to be generated and transmitted if an ITS-S is perceiving at least one object. On the other hand, transmitting a CP message that indicates to receiving ITS-S that no object or obstacle has been perceived in the perception area of the transmitting ITS-S also represents valuable information for safety applications. For this purpose, a container providing information about the sensory capabilities of the transmitting ITS-S should still be transmitted at a low rate, thereby constituting a beaconing mechanism. Under the assumption that a receiving ITS-S injects received objects into a sensor data fusion component predicting the movement of tracked objects, the CP message has to continuously provide updates about detected objects to reduce prediction errors on the receiver side. These errors, however, will largely depend on the dynamics of an observed object. In case of a stationary object, updates will have to be provided less frequently, whereas more dynamic objects (e.g., vehicles), need to be reported more often. However, more frequent updates increase the channel load, as more messages need to be transmitted.

A. Perceived Object Selection

Our proposed message generation process, as depicted in Figure 1, aims at addressing the aforementioned trade-off, i.e., balancing the amount of transmitted data and the resulting data quality. The process consists of three sub-processes that are checked cyclically, but only if a minimum time duration between two consecutive message generation attempts $T_{\text{Gen\_CPM}}$ has passed. This parameter is set to a value between $0.1\,\text{s}$ and $1\,\text{s}$ and may be set by other entities of the communications stack, such as the DCC entity to adapt message generation to the available channel resources.

Whenever a CPM is generated, the CP service queries the environment model of the ITS-S it runs on and extracts a list of all objects fulfilling minimum object confidence criteria, each associated with a set of object states such as position, speed, etc. This ensures that 'ghost'-objects and other temporary objects of low plausibility or data quality are not considered for transmission in a CPM. A harmonized concept for applicable confidence criteria is required to purposefully include received object information into a vehicle's environment model. Its definition is currently undergoing further investigation in other research projects and out of the scope of this paper.

However, the aforementioned list therefore contains all candidates that are selected for transmission, if one of the following conditions is fulfilled:

- **Novelty:** The object is included in the newly generated message if it has never been selected for transmission before. This mechanism is based on a unique object identifier that is maintained as long as consecutive sensor measurements are assigned to this object.
- **Distance:** The object is included in the newly generated message if it has moved by more than $4\,\text{m}$ with respect to the last time this particular object has been selected for transmission in a CPM. In case of a moving transmitting ITS-S, this distance differential shall refer to the absolute distance travelled (i.e. the measured relative distance corrected by the movement of the transmitter). This requirement is derived from the CA service [2].
- **Speed:** The object is included in the newly generated message if its speed has changed by more than $0.5\,\text{m/s}$ with respect to the last time this particular object has been selected for transmission in a CPM. In case of a moving transmitting ITS-S, this speed differential shall refer to the absolute speed change of the object (i.e. the measured relative speed corrected by the speed of the transmitter). This requirement is also derived from the specification of the CA service [2].
- **Age:** The object is included in the newly generated message if it has not been included in a CPM for more than $1\,\text{s}$. This requirement enables the indication of a standstill object and is relevant in dynamic scenarios in which the generally unknown number of receivers varies.

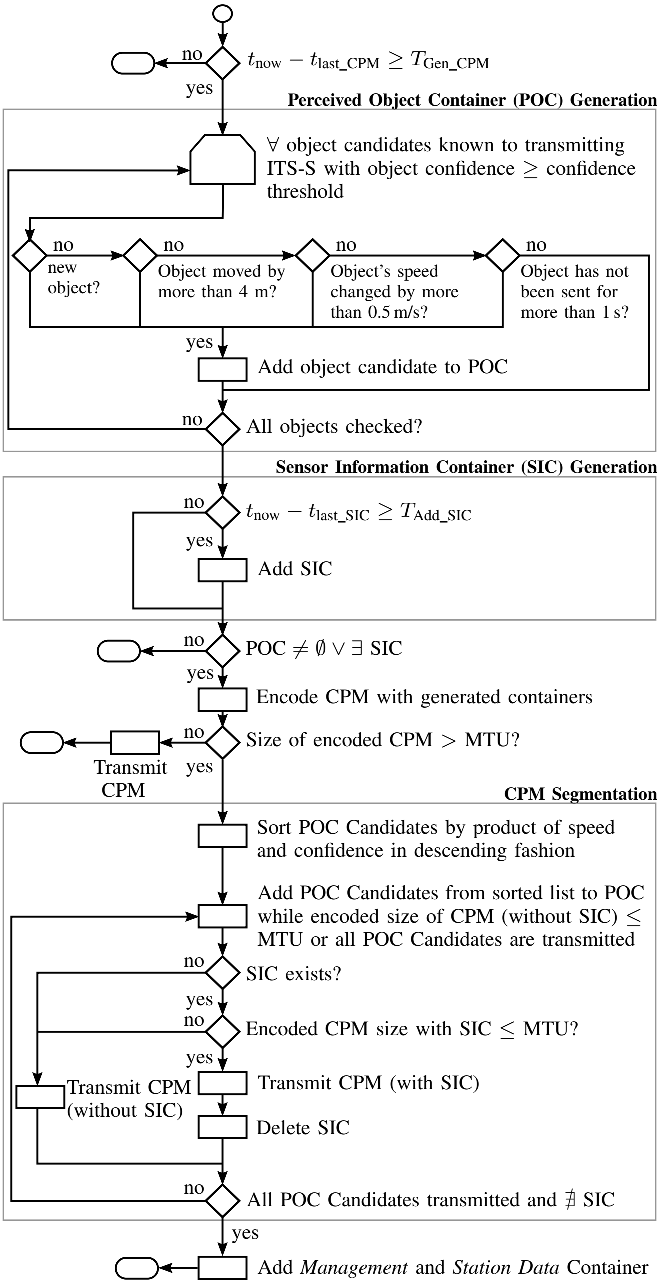

Fig. 1. Simplified Flow-Chart for CPM Generation

It has to be noted that due to the selection of these thresholds based on the specification of the CA service, different threshold-levels may be applied in case a disseminating ITS-S detects the presence of VRUs which are generally exhibiting different dynamic properties compared to vehicles. Each object candidate fulfilling any of the above conditions is included in the POC of the newly generated message. As a result of these conditions, the process has to be able to keep track of the spatial-temporal state of each object with respect to its last inclusion in a message. Furthermore, the CPM provides a data field to indicate the total number of objects that the disseminating ITS-S is aware of. This allows receiving ITS-S to expect further object reports in consecutive messages.

B. Sensor Information Container Inclusion

After evaluating these conditions for all object candidates, the process proceeds to checking if a *Sensor Information Container* also needs to be included in the currently generated CPM. This container should be included in a CPM less often, as the information contained in the SIC is generally static, i.e. the sensory properties of the disseminating ITS-S are persistent. The container is appended every $T_{\text{Add\_SIC}}$.

In case neither a POC nor a SIC has been generated as part of the previous process steps, no CPM is generated at all.

Otherwise, the CPM is generated by encoding all generated containers (including the required *Management* and *Station Data Container*). Message encoding makes use of the Abstract Syntax Notation One (ASN.1) Unaligned Packed Encoding Rules (UPER) [18]. The resulting encoded payload size is evaluated against the Access Layer (AL)-specific Maximum Transmission Unit (MTU): In case the encoded message size does not exceed the allowed MTU, the message is passed on to the lower layers for transmission.

C. Message Segmentation

If the resulting message size exceeds the MTU, message segmentation occurs. This mechanism is especially relevant for dense traffic situations in which a large number of dynamic objects has been detected and selected for transmission, despite the inclusion rules detailed in section III-A aiming at reducing this number. Each generated message segment can thereby be interpreted as a standalone message on the receiver-side. To increase spectral efficiency, every segment contains the relevant information required to perform coordinate transformation and alignment processes and can be decoded despite not receiving other segments. Nevertheless, the existence of further message segments can be indicated as part of the generated message. Depending on the implementation and message priority selection procedures of the lower layer, secondary messages may be subjected to a lower likelihood of being transmitted due to DCC operations and regulatory duty cycle requirements.

Consequently, when generating the message segments, the proposed process first sorts all the perceived object candidates that have been selected for transmission as detailed in section III-A by the product of an object's speed and the associated confidence in a descending fashion. This ensures that more dynamic objects, and objects of a high measurement quality (i.e. low inter-object state variable covariances) are selected for the first message segment to be generated. For assembling consecutive CPM segments, POC candidates are sequentially flagged for transmission as part of the current segment, as long as the resulting UPER-encoded payload size of the resulting segment including the required *Management* and *Station Data Container* (but not the SIC) is less than or equal to the AL-specific MTU. In case sufficient space exists in the segment to be generated to also include the SIC, it is also added to and encoded in the segment. Once the SIC has been added to a segment, it is no longer considered for consecutively generated segments for the current message generation iteration. Each generated segment is passed on to the lower layers for transmission. This process is repeated for all identified POC candidates until all candidates are included in one segment.

IV. SIMULATION AND ANALYSIS

The effectiveness of the proposed generation rules for the CP service as well as the potential benefit of incorporating MCO is assessed in an extensive simulation study. The following subsections detail the employed simulation environment, the simulation scenarios and parameter variations and analyze the simulation results.

A. Artery

All simulations have been performed with the simulation framework Artery [19]. This framework couples the microscopic traffic simulator Simulation of Urban Mobility (SUMO) and the discrete event simulator Open Modular Network Testbed in C++ (OMNeT++)[1] with its framework INET[2] which provides realistic propagation and shadowing models. Within Artery, the upper layers of the ETSI ITS-G5 communication stack are represented by the open-source implementation Vanetza[3]. Artery also features an environment model and an abstract implementation of local perception sensors that can be enabled for vehicles. The CP service is implemented as a service in Artery and leverages the local perception sensors to acquire information about surrounding objects to be reported in CPMs according to the generation rules detailed in Section III. The service also evaluates received CPMs to generate metrics required for the analysis of the simulation results. The implementation of the CP service in Artery is explained in more detail in [20].

B. Simulation Scenarios

Two different scenarios were simulated for the study. The Luxembourg SUMO Traffic (LuST) scenario [21] is well established in V2X related research and has been used in many studies before. It reflects 24 h of realistic traffic demand in the state of Luxembourg. This scenario is used to gather information on resulting realistic message sizes to be expected and whether at all there is a need for message segmentation. The scenario consists of eleven 5 s time slots every 2 h throughout the day to capture all daytime-specific traffic effects.

---

[1] omnet.com [2] inet.omnetpp.org/ [3] github.com/riebl/vanetza

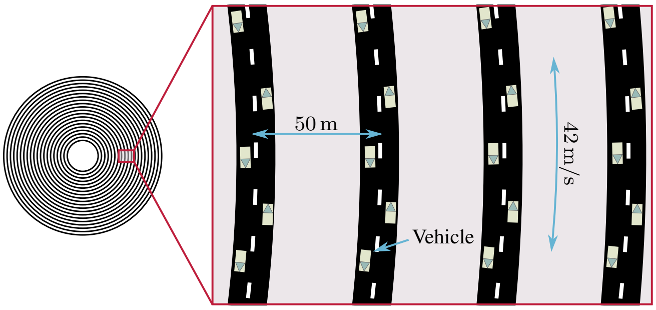

Fig. 2. Layout and Detail of the Spider scenario

Due to the large number of vehicles in the scenario (more than 4000 vehicles), simulating the entire day would be computationally unfeasible.

The second scenario, an artificial setup referred to as the *Spider* scenario due to its road layout, has been created for this study to examine the performance of the CP service in most challenging conditions with respect to the radio channel. The scenario consists of 20 concentric circular two-lane roads evenly spaced at 50 m with each lane being occupied with opposing traffic, as depicted in Figure 2. All vehicles in the scenario are travelling at a constant speed of 42 m/s (151 km/h) and all 40 lanes are populated as densely as the underlying driving model of SUMO allows, resulting in a total of 1800 vehicles within the radius of one kilometer around the center. To assess the CP service in challenging channel load conditions, this setup has been selected to result in the highest possible message generation rate of 10 Hz for the proposed generation rules. After a warm-up period of 2 s, the simulation is run for 10 s which is sufficient due to the homogeneous traffic pattern.

C. Parameterization

Each scenario is assessed in different parameter variations as listed in Table I. In the more realistic LuST scenario, a combined front-facing mid- and long-range sensor (± 40°, 65 m and ± 5°, 150 m, respectively) has been used. In the Spider scenario, vehicles are equipped with a 360° radar sensor with a range of 150 m. This has been chosen to detect even more surrounding objects and to increase the resulting message sizes. The market penetration rate, i.e., the percentage of vehicles equipped with V2X hardware and therefore the capability to transmit both CAMs and CPMs, is varied between 5 % and 100 %, influencing the overall effectiveness of the CP service as well as the channel load. Furthermore, the *dynamic* generation rules introduced in Section III are compared to a *static* baseline generating a complete message every 100 ms, with the SIC being added every 1000 ms.

As for the radio configurations, the CA message is always transmitted in the CCH with the required DCC priority level of DCC Profile (DP)2 [22]. Transmission of the CPM is varied for Single-Channel Operation (SCO) and MCO setup. For the *SCO DP2* configuration, the CPM shares the CCH and DCC profile DP2 with the CA. For the *SCO DP3* configuration, the CPM is sent on the CCH with priority DP3. For the *MCO* configuration, the CPM is disseminated on the separate SCH 1. As the CAM and CPM are the only messages in their respective channels, different DCC priorities do not have any influence. Fixed transmission and CP service parameters are summarized in Table II. A full factorial experiment design has been selected, resulting in 432 simulation runs.

TABLE I
VARIABLE PARAMETERS USED FOR THE SIMULATION STUDY

Scenario	Market Penetration	Generation Rules	Transmission Configuration
LuST (11 slots of 5 s), Spider (10 s)	5, 10, 25, 50, 75, 100	static, dynamic	SCO DP2, SCO DP3, MCO

TABLE II
FIXED PARAMETERS USED FOR THE SIMULATION STUDY

Parameter	Configuration
Data Bitrate	6 Mbit/s
Transmission Power	200 mW (23 dBm)
Radio Propagation Model	Two Ray Interference Model [23]
Maximum Interference Range	1500 m
DCC Finite State Machine	TRC, 3x Active, 500 µs T_{on} [7]
DCC Queue Length	2
Maximum Segment Size	1100 B
$T_{\mathrm{Add\_SIC}}$	1 s
$T_{\mathrm{Gen\_CPM}}$	100 ms

D. Analysis

This section presents the findings and observations that can be drawn from the simulation study.

1) Realistic Traffic Profile (LuST scenario): The LuST scenario captures traffic in urban, suburban and rural environments throughout an entire day. It is found that the resulting message size increases with the traffic volume in the scenario, as shown for the SCO DP3 variant in Figure 3 at a market penetration rate of 100 %. For the rush hour traffic peaks at 10 AM and 4 PM to 6 PM (cf. [21]), message sizes increase up to 1100 B, with median message sizes ranging between 50 B and 110 B. For all time slots, the proposed dynamic rules outperform the static rules in terms of message size

By taking object dynamics into account, CPMs are also generated less frequently. Those two effects together lead to a reduced channel load as depicted in the boxplot of Figure 4 for the SCO DP3 variant. The slot of 8 AM thereby represents the time of highest traffic volume to showcase these effects under most challenging conditions. With growing market penetration rate (i.e., number of vehicles sending CPMs) the Channel Busy Ratio (CBR) increases. However, the level caused by the dynamic rules is on average less than 50 % of the level caused by the simple, static generation rules. This shows that the dynamic generation rules meet their design goal of heavily reducing channel load.

Reducing the channel load however is only one half of the trade-off. The aim of the CP service is to generate awareness about other, potentially non-communicating objects in the surroundings. Therefore, reducing the channel load should

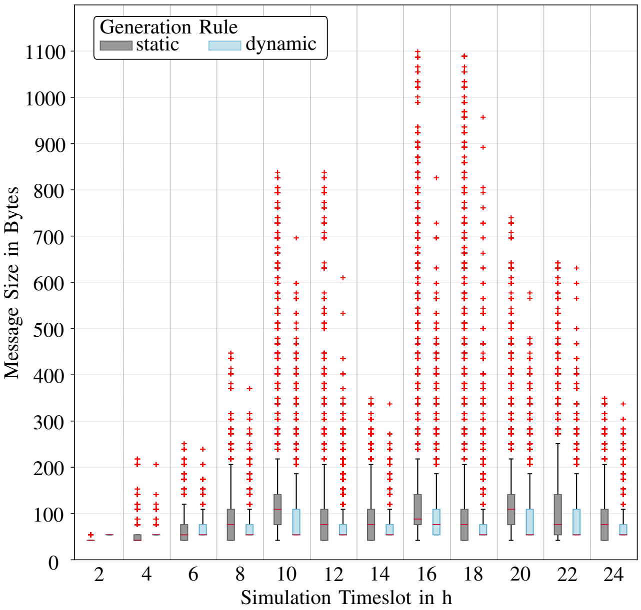

Fig. 3. Message sizes for both generation rules throughout a day (LuST, 100 % market penetration, SCO DP3)

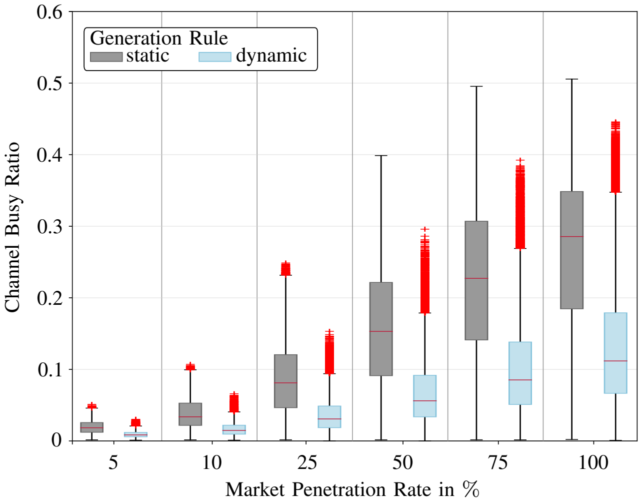

Fig. 4. Comparison between the static and dynamic generation rules at varying market penetration rates (LuST, 8 AM)

not lead to an extensive degradation of the service quality provided to applications. Among other metrics, this can be measured by the amount of distinct objects a vehicle is aware of only due to receiving a CPM, counting both the sender of the CPM, as well as the Perceived Objects (POs) contained in the messages. A boxplot of those counts, aggregated over all vehicles at 8 AM at 100 % market penetration rate is shown in Figure 5. On average, the number of objects known by CPM reception does not differ significantly between the two generation rules and transmission configurations, with differences resulting from timing effects at the border of the transmission range. However, the resulting channel load is heavily reduced for dynamic generation rules as shown in Figure 4 as mentioned before.

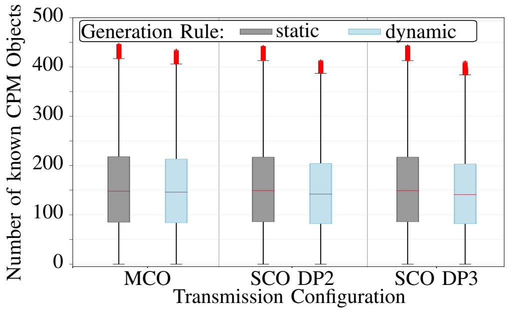

Fig. 5. Number of objects known to a host vehicle as a result of CPM reception. The displayed metric accumulates unique known objects over a sliding window of 100 ms. Objects not reported via CPM for more than 1 s are no longer considered. (LuST, 8 AM)

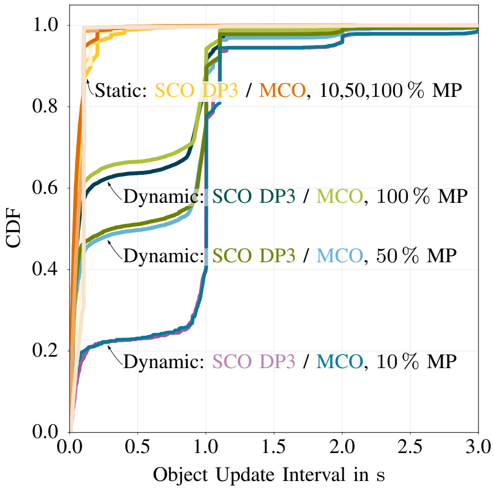

Fig. 6. Time between updates for unique objects for different transmission variants, generation rules and market penetration rates (MP) (LuST, 8 AM)

The rate at which an ITS-S receives updates about each unique object might also be negatively influenced by the dynamic generation rules. Figure 6 shows a Cumulative Distribution Function (CDF) of the intervals between two updates (regardless of the sender of the containing CPM) per unique object at 8 AM for the LuST scenario. For the dynamic rules, the intervals decrease with increasing market penetration as more ITS-Ss transmit information about the same object. The static generation rules achieve much shorter inter update intervals at the cost of producing at least twice the channel load on average. However, this is expected as the dynamic generation rules aim at reducing the updates in less dynamic situations e.g., vehicles waiting in front of a traffic light.

2) High Load Traffic Profile (Spider scenario): For the Spider scenario, the effect of the different radio configurations and DCC parameters is studied in more challenging channel load conditions. Due to the scenario setup, the effect of the different generation rules can be neglected, as the dynamic rules will trigger message generation at the same frequency as the static rules due to all vehicles driving at a constant

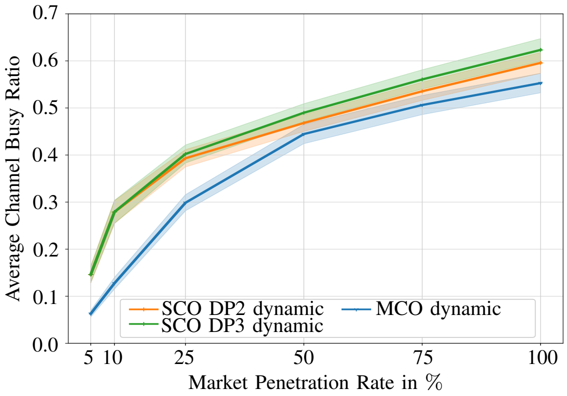

Fig. 7. Average Channel Busy Ratio ($\mu \pm \sigma$) for different market penetration rates and transmission configurations on the channel utilized by the CPM (Spider scenario, Dynamic Message Generation Rules)

TABLE III
MAPPING OF CBR TO PACKET RATE ALLOWED BY DCC [7]

State	**CBR**	**Max. Packet Rate**
Relaxed	$<$ 30 %	20 Hz
Active 1	30 % to 39 %	10 Hz
Active 2	40 % to 49 %	5 Hz
Active 3	50 % to 65 %	4 Hz
Restricted	$>$ 65 %	1 Hz

speed of 42 m/s (this will always trigger speed condition in Section III-A). For the same reason, vehicles also send CAMs at the defined maximum rate of 10 Hz with an average size of about 42 B.

Due to the scenario layout, vehicles always perceive between 27 to 30 surrounding objects, resulting in message sizes of 900 B ± 120 B ($\mu \pm \sigma$).

Large messages sent at 10 Hz are expected to put a high load on the radio channel utilized by the CPM, as shown in Figure 7. At 100 % market penetration rate, the CBR is well above 50 % for all three configurations. In addition to that, Figure 7 yields another observation. For the SCO variants, both CAM and CPM are sent in the same radio channel (i.e., the CCH). Since both are sent at 10 Hz, they are expected to generate a comparable channel load. This is evident at 5 % and 10 % market penetration as the CBR for the MCO variant is approximately 50 % of the SCO variants. This expected behaviour vanishes with increasing penetration rates while the difference between the SCO DP2 and SCO DP3 variants increases. This is a result of DCC reacting to the high channel load by reducing each node's message budget according to Table III. The maximum packet rate is applicable per channel and therefore for the sum of CAM and CPMs in the SCO variants. For penetration rates larger then 50 %, the MCO variant resulting in CBRs nearly as high as for the SCO cases indicates a large amount of packets dropped by the DCC gatekeeper for both SCO variants. In general, the SCO DP3 variant exhibits the highest overall channel load as both messages do not compete for the same DCC queue. As a result, more messages are disseminated than in the SCO DP2 case.

Figure 8 shows boxplots of the number of distinct objects the receiving ITS-Ss are aware of due to the CAM (Figure 8a) and the CPM (Figure 8b) at 100 % market penetration. The baseline of the generated awareness is the leftmost variant (MCO) in which both messages use a separate channel for dissemination and therefore do not compete for transmission opportunities. Here, the CAM and CPM account for a median of about 480 and 1015 objects, respectively. In the SCO DP2 variant, both of those numbers are reduced as the messages are sent in the same channel. As both messages have the same priority (DP2), it is up to the specific implementation of the communications stack to determine which message gets dropped upon queue insertion. Nevertheless, both services will suffer from such a configuration. In our implementation, the CPM replaces a CAM in the queue and the resulting awareness generated decreases by less than 10 % whereas for the CA message, the number of known objects drops by almost 50 %. For the SCO DP3 configuration, however CPM transmission is reduced significantly. This becomes apparent in Figure 8b, as the median number of objects known due to the CPM drops to around 100, while CAM transmission is not influenced.

This shows that it is not preferable to co-locate both services in the same channel in high density scenarios, which confirms the findings of [14].

In addition, for both scenarios, the need for message segmentation for a maximum size of a CPM of 1100 B has been evaluated. In the LuST scenario, segmentation did not occur at all and even in the Spider scenario, only 0.6 % of all messages had to be segmented. However, the need for segmentation might emerge, with future sensing mechanisms, increasing the number of perceived objects.

V. CONCLUSION

This paper suggests a set of generation rules for the CP service which is currently standardized at the ETSI. The generation rules aim at reducing the channel load caused by the CP service while maintaining a high service quality in terms of the number of known objects as well as the update frequency per unique object known. An extensive simulation study has shown that the trade-off between channel load and service quality can be addressed by the proposed generation rules. These achieve a similar service quality while heavily reducing the channel load when compared to a simple, static baseline generation method.

Furthermore, the simulation results show that in high-load situations, the CPM and the CAM compete for transmission opportunities if sent in the same channel. This results in a severe degradation of service quality for one or even both services, depending on the specific implementation of the communications stack. Assigning the CPM to a separate channel improves the situation.

Future work focuses on more elaborated MCO approaches by combining higher data rates and reduced transmit power. Alternative DCC algorithms [24], [25] potentially reducing CP service degradation have to be evaluated as well. Additionally, incorporating lower-layer feedback about message drops into the generation rules may further increase CP service quality.

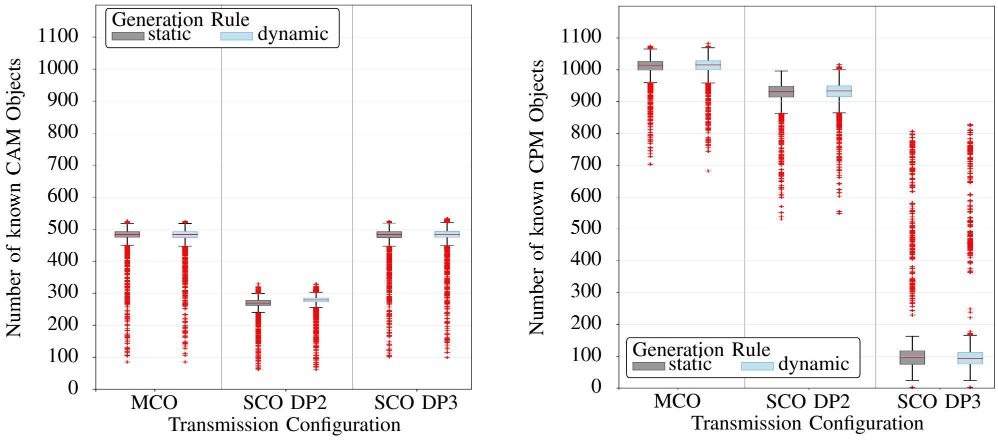

(a) Awareness generated by receiving CAs message only (b) Awareness generated by receiving CPMs only

Fig. 8. Boxplots for total number of known objects by each communicating vehicle (no duplicates) over a sliding window of 1000 ms for indicated message type and different transmission configurations (Spider scenario, 100 % market penetration)

Although the proposed rules already reduce the resulting channel load and message sizes, the current simulation results also hint at further potential to reduce the number of generated messages, thereby decreasing additional header overhead. By implementing a *look-ahead* mechanism predicting the object states to the next likely CPM generation time, premature object inclusion could further reduce channel utilization.

REFERENCES

[1] H.-J. Günther, O. Trauer, and L. Wolf, "The Potential of Collective Perception in Vehicular Ad-hoc Networks," in *IEEE International Conference on ITS Telecommunications (ITST)*, 2015.

[2] *ETSI EN 302 637-2 V1.4.1 - Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Part 2: Specification of Cooperative Awareness Basic Service*, ETSI Std., 2019.

[3] "Honda's V2X Technology at a Crossroads," https://www.autonews.com/article/20181022/MOBILITY/181029939/honda-s-v2x-technology-at-a-crossroads, accessed: 2019-07-04.

[4] *ETSI TR 103 562 V0.0.20 - Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Analysis of the Collective Perception Service*, ETSI Std., 2019.

[5] *ETSI TS 103 324 V0.0.14 - Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Specification of the Collective Perception Service*, ETSI Std., 2019.

[6] *J 2945/8 Cooperative Perception System (standardization in progress)*, SAE International Std., 2012.

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

[8] *ETSI EN 302 571 V2.1.1 - 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*, ETSI Std., 2017.

[9] *ETSI TS 102 724 V1.1.1 - Intelligent Transport Systems (ITS); Harmonized Channel Specifications for Intelligent Transport Systems operating in the 5 GHz frequency band*, ETSI Std., 2012.

[10] T. Leinmüller, P. Spaanderman, and A. Festag, "Next Steps for Multi-Channel Operation in EU V2X Systems," *ITS World Congress and Exhibition*, 2016.

[11] A. Rauch, F. Klanner, R. Rasshofer, and K. Dietmayer, "Car2X-Based Perception in a High-Level Fusion Architecture for Cooperative Perception Systems," in *IEEE Intelligent Vehicles Symposium (IV)*, 2011.

[12] A. Rauch, F. Klanner, and K. Dietmayer, "Analysis of V2X Communication Parameters for the Development of a Fusion Architecture for Cooperative Perception Systems," in *IEEE Intelligent Vehicles Symposium (IV)*, 2012.

[13] A. Rauch, S. Maier, F. Klanner, and K. Dietmayer, "Inter-Vehicle Object Association for Cooperative Perception Systems," in *IEEE Intelligent Transportation Systems Conference - (ITSC)*, 2013.

[14] H.-J. Günther, R. Riebl, L. Wolf, and C. Facchi, "The Effect of Decentralized Congestion Control on Collective Perception in Dense Traffic Scenarios," *Computer Communications*, vol. 122, pp. 76–83, 2018, Elsevier.

[15] H.-J. Günther, B. Mennenga, O. Trauer, R. Riebl, and L. Wolf, "Realizing Collective Perception in a Vehicle," in *IEEE Vehicular Networking Conference (VNC)*, 2016.

[16] G. Thandavarayan, M. Sepulcre, and J. Gozalvez, "Analysis of Message Generation Rules for Collective Perception in Connected and Automated Driving," in *IEEE Intelligent Vehicles Symposium (IV)*, 2019.

[17] T. Higuchi, M. Giordani, A. Zanella, M. Zorzi, and O. Altintas, "Value-Anticipating V2V Communications for Cooperative Perception," in *IEEE Intelligent Vehicles Symposium (IV)*, 2019.

[18] *X.691: Data Networks and Open System Communication - OSI networking and system aspects - Abstract Syntax Notation One (ASN.1)*, ITU-T Std., 2002.

[19] R. Riebl, H.-J. Günther, C. Facchi, and L. Wolf, "Artery: Extending Veins for Vanet Applications," in *IEEE International Conference on Models and Technologies for Intelligent Transportation Systems (MT-ITS)*, 2015, pp. 450–456.

[20] K. Garlichs, M. Wegner, and L. C. Wolf, "Realizing Collective Perception in the Artery Simulation Framework," in *IEEE Vehicular Networking Conference (VNC)*, 2018.

[21] L. Codeca, R. Frank, and T. Engel, "Luxembourg SUMO Traffic (LuST) Scenario: 24 Hours of Mobility for Vehicular Networking Research," in *IEEE Vehicular Networking Conference (VNC)*, 2015.

[22] CAR 2 CAR Communication Consortium, "Basic System Profile 1.3.0," 2018.

[23] C. Sommer, S. Joerer, and F. Dressler, "On the Applicability of Two-Ray Path Loss Models for Vehicular Network Simulation," in *IEEE Vehicular Networking Conference (VNC)*, 2012.

[24] G. Bansal, B. Cheng, A. Rostami, K. Sjoberg, J. B. Kenney, and M. Gruteser, "Comparing LIMERIC and DCC approaches for VANET channel congestion control," in *IEEE International Symposium on Wireless Vehicular Communications (WiVeC)*, 2014.

[25] G. Bansal, J. B. Kenney, and C. E. Rohrs, "LIMERIC: A Linear Adaptive Message Rate Algorithm for DSRC Congestion Control," *IEEE Transactions on Vehicular Technology*, vol. 62, no. 9, pp. 4182–4197, 2013.