This is an author-created post print version
The final publication is available at: https://www.sciencedirect.com/science/article/pii/S1084804523000747?dgcid=author

Scalable Cooperative Perception for Connected and Automated Driving

Gokulnath Thandavarayan, Miguel Sepulcre, Javier Gozalvez and Baldomero Coll-Perales
gthandavarayan@umh.es, msepulcre@umh.es, j.gozalvez@umh.es, bcoll@umh.es
Universidad Miguel Hernandez de Elche (UMH), Avda. Universidad s/n, 03202, Elche (Alicante), Spain.

Corresponding author: Gokulnath Thandavarayan (gthandavarayan@umh.es).

Abstract Cooperative perception (a.k.a. collective perception or cooperative sensing) will allow Connected and Automated Vehicles (CAVs) to share information about detected objects. Cooperative perception improves the sensing accuracy, confidence and range of CAVs, and extends their perception of the driving environment. First message generation rules based on the mobility and dynamics of detected objects have been proposed to decide when a cooperative perception message should be generated and what information it should include. Studies have shown that this type of generation rules can compromise the scalability of cooperative perception and vehicular networks, as they tend to transmit significant amounts of redundant information and generate small and frequent cooperative perception messages that increase the communications overhead. To combat these inefficiencies, this paper proposes and evaluates three techniques that combine, for the first time, baseline mobility-based generation rules for cooperative perception messages with mechanisms to control the redundancy and organize the information about detected objects in order to avoid the frequent transmission of small messages. This study demonstrates that the proposed techniques improve the perception of CAVs and reduce the information age. In addition, the techniques reduce the channel load and improve the scalability of cooperative perception services and vehicular networks. The study demonstrates that the most effective technique is based on: (1) first applying the generation rules to decide whether a cooperative perception message should be generated, (2) then applying redundancy control, and finally (3) organizing the information about all detected objects to avoid small and frequent messages.

Keywords Cooperative perception, collective perception, cooperative sensing, CPS, CPM, connected automated vehicles, autonomous vehicles, CAV, RSU, V2X, vehicular networks, congestion control, redundancy mitigation, C-ITS, DCC, ETSI.

I. Introduction

Connected and Automated Vehicles (CAVs) will use V2X (Vehicle-to-Everything) communications to complement the capabilities of their onboard sensors and improve their safety and driving [1-3]. The sensing technology has significantly improved over the last years [4], but the capabilities of onboard sensors like cameras, radars or lidars are still limited under the presence of obstacles or adverse weather conditions, among other factors [5][6]. V2X communications can help overcome these limitations through the exchange of sensor information among CAVs. By sharing this information, CAVs can extend their field of view beyond that of their own onboard sensors' as well as improve the sensors' detection accuracy. The process of exchanging sensor information is generally referred to as cooperative perception, collective perception or cooperative sensing [7][8].

ETSI (European Telecommunications Standards Institute) and SAE (Society of Automotive Engineers) are currently working to define new V2X standards for cooperative perception. SAE has not yet published its standard for cooperative perception [7]. On the other hand, ETSI published a Technical Report on collective perception in December 2019 [8], and is currently working on the corresponding Technical Specification [9]. In [8], ETSI defined the so-called Collective Perception Service (CPS). The CPS includes the Collective Perception Message (CPM) format and the generation rules to decide when a new CPM should be generated and what information it should include. The core of the CPM generation rules included in ETSI's Technical Report was initially proposed in [10], and is based on the exchange of information about detected objects. These rules when a CAV should include an object it has detected with its onboard sensors in a CPM. The decision is based on the mobility of the detected object (e.g., its position and speed). These ETSI CPM generation rules are referred to as baseline generation rules in this paper. Several studies have revealed that current CPM generation rules present certain inefficiencies that can overload the communications channel and limit their scalability [11]. The first inefficiency is related to the fact that the current CPM generation rules can generate a high number of CPMs with a small payload. This is the case because the CPM headers and the protocol headers can be significantly large. Therefore, when the payload is small, most of the bandwidth is used to transmit overhead data. The second inefficiency is the transmission of redundant information. This occurs when multiple CAVs detect the same object simultaneously and include its information on their CPMs. The transmission and reception of redundant information about the same object can unnecessarily overload the communications channel and increase the computing power needed to process the exchanged sensor information. The authors have addressed these two inefficiencies separately in [12] and [13]. In [12], the authors proposed the so-called Look-Ahead technique that groups objects into larger CPMs and reduces the number of CPMs generated with a small payload. In [13], the authors proposed a redundancy mitigation technique that controls the amount of redundant information by reducing the number of vehicles that simultaneously report about the same detected object. Both solutions are compatible and extend the CPM generation rules initially proposed in [10]. They can provide significant benefits in terms of perception and channel load. In fact, Look-Ahead and redundancy mitigation are now part of ETSI specifications [8][9]. These techniques have been so far designed and evaluated independently. Combining them could potentially lead to additional gains because both techniques address different inefficiencies in the baseline generation rules. However, their combination has not been studied yet, and must be carefully designed since both techniques may have opposite effects on the generation of cooperative perception messages. For example, Look-Ahead generates larger CPMs by grouping objects into a smaller number of CPMs, but it can increase the amount of redundancy because objects can be transmitted more frequently than with the baseline CPM generation rules [12]. On the other hand, redundancy mitigation techniques reduce the amount of redundancy transmitted, but they may also increase the generation of small CPMs [13].

In this context, this study progresses the state-of-the art by proposing how Look-Ahead and redundancy mitigation should be combined to improve cooperative perception and the system's scalability. We focus on these two specific techniques because they are part of ETSI specifications on collective perception, and their combination has not been studied yet [8][9]. This study also analyzes the impact of congestion control and the coexistence of CPMs with other messages on cooperative perception, and in particular on the combination of Look-Ahead and redundancy mitigation or control. The analysis considers the DCC (Decentralized Congestion Control) framework defined by ETSI for congestion control, and analyzes the impact of DCC Access and DCC Facilities on the operation and performance of the combined Look Ahead and redundancy mitigation. Congestion control can significantly alter the generation and transmission of CPMs, and this analysis is key for the integration of the proposed techniques in the protocol stack and their consideration in standards. The conducted evaluation demonstrates that the proposed techniques improve the perception of CAVs and reduce the channel load and the information age compared to the baseline CPM generation rules.

2. State of the Art

Cooperative perception can help to overcome the limitations of onboard sensors and thus improve the safety and driving conditions of CAVs by exchanging sensed information. Some studies have proposed and evaluated the exchange of raw sensor information for cooperative perception [14]. However, exchanging raw sensor data would require a high bandwidth that can compromise the system's scalability. Other formats such as layered cost maps [15] are also proposed to specify the sensor data in a grid-based representation. To accurately track objects using a grid-based representation, vehicles need a 3D representation for each layer. This could increase the system's complexity and the computational load. As a consequence, the majority of studies conducted to date consider instead the exchange of processed information about detected objects (e.g., their position, speed, size and type) for cooperative perception. In this paper, we have reviewed and analyzed the most relevant studies on this topic. These studies are classified in Table 1 according to: 1) whether they consider the ETSI baseline generation rules as a benchmark or not; 2) whether they address the problem associated to the overhead; 3) whether they study redundancy mitigation; and 4) whether they analyze the impact of DCC on cooperative perception or not.

Part of the existing studies focused on the definition of the message format and defined first techniques for message rate and content control as an alternative to the baseline generation rules defined by ETSI ([16]-[24] in Table 1). Initial studies [16]-[17] addressed the definition of the message format and the type of information that should be included in cooperative perception messages. Günther et al. considered in [16] the exchange of periodic messages, each message including all the detected objects. They also proposed to include additional information about the transmitter (e.g., its position) and its sensor capabilities (e.g., range and field of view). This additional information helps the receiving vehicles estimate the free-space areas and better locate the objects with respect to the transmitter, which improves the detection time and reduces potential safety risks. Alternatively, the authors in [17] study the potential benefits of including additional information, such as the correlation and higher order derivatives (e.g., the acceleration or yaw rate) of the detected objects on the fusion accuracy. The work in [17] shows that this additional information can improve the detection accuracy and reduce the channel load.

Other studies focused on controlling the rate at which cooperative perception messages are generated. These techniques are crucial to efficiently use the available bandwidth and, at the same time, obtain high perception levels. Initial studies like [24] evaluated if the information about detected objects should be attached to existing awareness messages (in particular, CAM or Cooperative Awareness Messages [38]) or should be transmitted in separate messages. Attaching the information about detected objects to existing awareness messages could reduce the communications overhead. However, it would condition when information about detected objects is transmitted. As a consequence, most of the conducted studies transmit cooperative perception messages separately, i.e.,

Table 1. Comparison of studies about cooperative perception

Reference	Baseline	Overhead	Redundancy	DCC
[16]-[23]	No	No	No	No
[24]	No	No	No	Access Reactive
[11], [25]-[30]	Yes	No	No	No
[12]	Yes	Yes	No	No
[13], [31]- [35]	Yes	No	Yes	No
[10]	Yes	No	No	Access Reactive
[37]	Yes	No	No	Access and Facilities
Our proposal	Yes	Yes	Yes	Access and Facilities

independently to the transmission of existing basic awareness messages.

The transmission of periodic CPMs was shown to be highly inefficient and generated an unnecessarily high channel load. As a consequence, most of the conducted studies have proposed and analyzed different techniques to control the message rate and content [18]-[23] to mitigate this problem. Studies [18] and [19] focused on the adaptation of the message generation rate. The authors in [18] propose to increase or decrease the rate at which CPMs are generated based on each vehicle's priority. The priority depends on whether the vehicle is detecting objects that are not detected by nearby vehicles. The goal is to cover the blind sensor area of nearby vehicles by transmitting the detected objects in this region with higher priority. Similar to [18], the authors in [19] proposed a probabilistic dynamic scheme to control the transmission of CPMs. The study demonstrated that high awareness can be achieved in the network when transmission priority of vehicles is adapted based on its sensing capabilities.

Other studies such as [20] and [21] proposed content control techniques to dynamically select the objects that must be included in each CPM, and thus optimize the message size. The study in [20] proposes a content control technique that includes an object in a CPM depending on its value or utility for its neighboring vehicles. The idea is to omit those objects that are not important for other vehicles. The accurate estimation of the value of each object in a distributed and dynamic environment is challenging, and the same authors partially address this challenge in [21] using deep reinforcement learning. A recent publication [23] decides based on the relative entropy between the sensors tracking accuracy and the V2X tracking accuracy if the object information is valuable for the nearby vehicles and must therefore be included in a CPM. One of the first studies to propose solutions to dynamically control both the rate and content of the CPM was [22]. The study analyzes the impact of such dynamic adaptation on tracking errors and mapping accuracy, and demonstrates that the joint rate and length control mechanism performs better than the adaptation of either the rate or the length. The study also concludes that the objects that are located farther away from the sender but near the edge of the sensors' range should be prioritized.

The most popular adaptation of the message rate and content proposed to date was defined in [10] and is based on the selection of the detected objects depending on their mobility or dynamics (e.g., their speed, acceleration and heading). In this case, the transmitter includes an object in a cooperative perception message if the object's speed, acceleration or heading has significantly changed compared to the last time it was included in a message. This approach has been adopted so far on, for example, ETSI specifications on collective perception [8][9]. The definition of the baseline message generation rules by ETSI triggered the analysis of their performance in multiple studies ([11], [25]-[30] in Table 1).

The performance achieved with these baseline generation rules has been studied using analytical models for different radio access technologies in [25][26], has been evaluated by means of simulation in [8][27][11], and has been tested in prototypes [28]-[30]. The analytical studies in [25] and [26] provide important insights about the perception that CAVs can achieve with these baseline generation rules. They evaluate the impact of the market penetration rate and traffic density on the environmental awareness and information age. They also identify that, in some scenarios, reducing the rate at which objects are included in CPMs would be beneficial to reduce the channel load and interference. The simulation studies conducted in [8], [27] and [11] provide more detailed evaluations of the baseline generation rules. The study in [27] shows that cooperative perception improves traffic safety; for example, the study finds that a "very good" safety, with low risk of collisions without fatalities becomes possible at V2X equipment rates higher than 25%. The study in [11] shows that cooperative perception can improve the perception capabilities of CAVs, but identifies two main inefficiencies of the baseline generation rules that impose important limitations. First, the baseline generation rules can generate a high number of cooperative perception messages with a small payload. The transmission of cooperative perception messages with a small payload increases the channel load with communications overhead rather than useful information about detected objects; e.g. current ETSI CPM includes around 200 bytes of overhead [12]. This overhead includes CPM headers and also protocol headers from the lower layers of the stack (e.g., PHY and MAC). Second, a CAV can receive simultaneously redundant information about the same object from multiple vehicles. Receiving redundant information could, in principle, be considered positive, since it can help improve the detection accuracy and mitigate potential packet losses. However, a high level of redundancy may overload the communications channel and reduce the possibility to transmit cooperative perception messages with critical (and not redundant) information. In addition, redundant information increases the computing power necessary at each vehicle to process the information received. To date, these two inefficiencies have been

Gokulnath Thandavarayan, Miguel Sepulcre, Javier Gozalvez, and Baldomero Coll-Perales, "Scalable Cooperative Perception for Connected and Automated Driving." Journal of Network and Computer Applications, vol. 216, pp. 103655, May 2023. https://doi.org/10.1016/j.jnca.2023.103655

addressed independently in different studies, as it can be observed in Table 1 and is described below.

The problem associated to the overhead generated with messages of small size was addressed in [12] with the so-called Look-Ahead technique [12]. The Look-Ahead mechanism was designed to reduce the overhead generated with CPMs. It complements and extends the baseline generation rules by adding an additional process that groups the detected objects into larger CPMs. As a result, the message payload about detected objects is larger than the headers or overhead. [12] demonstrates that Look-Ahead can effectively reduce the rate at which CPMs are generated and increase their size. At the same time, it is able to reduce the channel load (and overhead), improve the reliability of V2X communications and enhance the perception of CAVs. To the authors' knowledge, no other study has addressed this first inefficiency to date.

The second inefficiency (redundancy) has been addressed in different studies through the design of redundancy mitigation or control techniques ([13], [31]-[35] in Table 1). For example, [31] proposes a probabilistic selection scheme to decide which objects are included in a perception message and suppress redundant transmissions. The scheme allows CAVs to adjust the transmission probability of each detected object based on the position, vehicular density and road geometry information. The study in [32] adapts the redundant information included in each CPM as a function of the channel load and the infrastructure availability. The results show that the proposed solution increases the awareness and reduces the channel load in the network compared to a periodic solution. The study in [34] employs a deep Q-network to determine the usefulness of detected objects based on their dynamic properties, using a reward function designed to reduce redundancy. The results of the study indicate that this approach improves perception and reduces the load on the communication channel. The authors propose in [13] the so-called dynamics-based redundancy mitigation technique. This technique extends the baseline generation rules to filter out the detected objects that have been recently transmitted by a nearby vehicle. To this aim, the transmitter omits from the perception message the detected objects that have not significantly changed their position, speed and heading since the last time it received a perception message with information about these same detected objects. The dynamics-based redundancy mitigation or control technique has been shown to reduce the channel load by up to 70% while achieving similar perception levels for short and medium distances when compared with the baseline generation rules. Other redundancy mitigation or control techniques listed by ETSI in [8] include: frequency-based, dynamics-based, confidence-based, entropy-based, distance-based and object self-announcement redundancy mitigation. The study in [33] evaluates some of these techniques and demonstrates that reducing or controlling redundancy can significantly improve network-related performance metrics (e.g., channel load and packet error rate) without significantly reducing the number of detected objects through perception messages and the time between updates about detected objects. The work in [35] evaluates the self-announcement redundancy mitigation technique defined in [8]. With this technique, the transmitter vehicle omits an object from the CPM if it detects that the object is transmitting its own messages (e.g. CAMs or CPMs). The results of the study show that this technique effectively balances the network load and the object perception. The study in [36] compares the performance of multiple redundancy mitigation techniques. The results showed that the distance-based [8] and dynamics-based [13] redundancy mitigation techniques achieve the best performance.

Look-Ahead and redundancy mitigation or control techniques extend the baseline generation rules and provide significant benefits in terms of perception and channel load. However, the two techniques have been designed and evaluated so far independently and their combination needs to be carefully analyzed since they may have opposite effects on the generation of cooperative perception messages. Both techniques are included in by ETSI in [8] and [9]. The combination of the two types of techniques has not been studied yet, and in fact, there are different ways to combine them as we show in this study. The combination must carefully look at how and when each object is included in a cooperative perception message since each decision affects not only to the messages transmitted by a CAV, but also the messages transmitted by other nearby CAVs because of the use of redundancy mitigation. In addition, it is important that the combination and evaluation considers the impact of congestion control mechanisms since these mechanisms decide and impact which messages are ultimately transmitted based on the channel load. It is important to highlight that only a few studies have analyzed so far the impact of congestion control schemes on the effectiveness of cooperative perception, as shown in Table 1. Günther et al. studied in [24] the impact of DCC Access on the periodic generation of cooperative perception messages, and in [10] when the baseline message generation rules are considered. The study conducted in [37] analyzed the impact of both DCC Access and DCC Facilities on cooperative perception considering the baseline generation rules, and demonstrated the importance of the DCC configuration on both perception and latency. To the authors' knowledge, none of the existing studies has evaluated such impact when considering Look-Ahead or redundancy mitigation techniques, and this evaluation is key for their integration in the ETSI ITS Architecture and thus for their future adoption in the standards.

In this context, this study evaluates the proposed solutions on top of the DCC framework defined by ETSI for congestion control, including DCC Access and DCC Facilities. DCC Access [39] is a mandatory component that operates as a gatekeeper in the Access layer to control the number of messages transmitted by each vehicle. To this aim, the upper bound of the fraction of time that each vehicle can transmit is calculated as a function of the channel load. If the time consumed by the messages generated by the upper layers is higher than this upper bound, the messages can be internally dropped (i.e., not transmitted) at the Access layer. Two approaches, namely Reactive and Adaptive, are defined for DCC Access [39] and both are evaluated in this paper. The DCC Access Reactive approach makes use of a state machine, with each state mapped to a range of channel load levels and a message rate. The DCC Access Adaptive approach uses a linear control process so that each vehicle adapts its message transmission rate to converge to a target channel load considering the duration of the messages. DCC Access can also be combined with DCC Facilities [40], which has been recently approved by ETSI. DCC Facilities is an optional component in the Facilities layer that controls the number of messages that each application/service can generate for each vehicle. To this aim, DCC Facilities makes use of the upper bound of the fraction of time that the vehicle can transmit used by DCC Access and information about the previously generated messages to adapt the message generation rate of each application/service. Previous studies have shown the limitations of DCC Access and the potential of DCC Facilities to improve both perception and information age when using the baseline generation rules [37].

The proposed solutions are compared in this study with the baseline generation rules defined by ETSI, as well as with the Look-Ahead technique originally proposed in [12] and the redundancy mitigation technique originally proposed in [13]. These three techniques used for the comparison have been selected because they are part of ETSI specifications [8][9], and, to the authors knowledge, there is no other combination technique proposed in the literature.

3. Collective Perception Service

The CPS is a Facilities-layer service defined by ETSI [8][9] that includes the necessary functions and interfaces to implement collective perception, and defines the CPM format and the CPM generation rules.

3.1. CPM format

The CPM format includes an ITS (Intelligent Transport System) PDU (Protocol Data Unit) header and five different containers: a Management Container, an Originating Vehicle or RSU (Road Side Unit) Container, a Sensor Information Container, a Perceived Object Container and a Perceived Region Container. The message ID, the station ID and the protocol version are included in the ITS PDU header. The Management Container is mandatory and includes information about the originating station (vehicle or RSU). The Originating Vehicle or RSU Container is optional and contains data elements with information about the originating station. The Sensor Information Container is also optional and includes information about the sensing capabilities of the transmitter. One Sensor Information Container is included for each sensor. The Perceived Object Container is also optional and is used to exchange information about the mobility and characteristics of the detected objects. One Perceived Object Container is included for each detected object. Lastly, the Perceived Region Container is optional and defines the actual perception capabilities available to the originating station, offering additional (often dynamic) details to the information provided in the sensor information container.

3.2. CPM generation rules

The CPM generation rules define when a vehicle or RSU should generate a CPM and the information that the CPM should include (e.g., which detected object must be included in the Perceived Object Container). The CPM generation rules are defined by ETSI in [9], and will be referred to as the baseline generation rules in this study. The baseline generation rules establish that a vehicle has to check every *T_GenCpm* if a new CPM should be generated. *T_GenCpm* should be set between 100 ms and 1000 ms and can be adapted by DCC based on the channel load. For every *T_GenCpm*, a vehicle should generate a new CPM if it has detected a new object (i.e., an object that the vehicle or RSU has not transmitted before), or if at least one of the following conditions is satisfied for any of the previously detected objects:

1. The absolute difference (ΔP) between the current position of the object and its position the last time it was included in a CPM is higher than 4 m.
2. The absolute difference (ΔS) between the current speed of the object and its speed the last time it was included in a CPM is higher than 0.5 m/s.
3. The time difference (ΔT) between the current time and the last time the object was included in a CPM is higher than 1 s.

A vehicle includes in a new CPM all new detected objects and those objects that satisfy at least one of the previous defined conditions (i.e., ΔP>4m or ΔS>0.5m/s or ΔT>1s). The vehicle still generates a CPM every second even if none of the detected objects satisfy any of the previous conditions. The information about the onboard sensors is included in the CPM only once per second.

3.3. Redundancy mitigation or control

Redundancy mitigation or control techniques extend the baseline generation rules to control the number of detected objects included in the CPM and reduce redundancy. Different redundancy mitigation techniques are defined by ETSI in [8][9]. Among them, this study focuses on the dynamics-based redundancy mitigation technique because it is based on the mobility of the objects and thus can be considered as a natural extension of the baseline generation rules, and because it provides one of the best performances [36]. This dynamics-based redundancy mitigation technique was first proposed and evaluated in [13], and then included by ETSI in [8][9]. It will be referred to as RM in this paper.

The dynamics-based redundancy mitigation technique does not include in a CPM an object that satisfies the baseline generation rules if the object's position or speed have not significantly changed since the last time the transmitting vehicle received information about this object in a CPM from any other vehicle. The objective is to avoid transmitting information about detected objects that have been recently included in CPMs transmitted by other vehicles, since this information has probably been received by most of the nearby vehicles.

The redundancy mitigation technique is executed after the baseline generation rules, and only when the baseline generation rules indicate that a new CPM must be generated. When executed, the technique computes for every detected

Gokulnath Thandavarayan, Miguel Sepulcre, Javier Gozalvez, and Baldomero Coll-Perales, "Scalable Cooperative Perception for Connected and Automated Driving." Journal of Network and Computer Applications, vol. 216, pp. 103655, May 2023. https://doi.org/10.1016/j.jnca.2023.103655

object that should be included in the CPM according to the baseline generation rules, the change in its absolute position (ΔP_R) and speed (ΔS_R) since the last time the object was received in a CPM transmitted by other vehicles. If $\Delta P\_R \leq P\_Threshold$ and $\Delta S\_R \leq S\_Threshold$, the object is not included the CPM. When the *P_Threshold* and/or the *S_Threshold* parameters decrease, the detected objects are omitted from the CPM less frequently, increasing the redundancy.

The evaluation of the dynamics-based redundancy mitigation technique in [13] demonstrated its potential to reduce the amount of information transmitted about detected objects to control the level of redundancy. Compared to the baseline generation rules, RM reduces the redundancy by 25% with *P_Threshold*=1m and 64% with *P_Threshold*=4m in a scenario with low traffic density. The redundancy increases with the traffic density, because more vehicles detect and transmit information about the same objects, but the gains obtained with RM are maintained for higher densities. This reduction of the object redundancy results in a reduction of the channel load. More specifically, [13] showed that the channel load is reduced by 17% and 58% compared to the baseline generation rules in the low traffic density scenario with *P_Threshold* equal to 1m and 4m, respectively. This reduction is even higher for the high traffic density scenario (26% and 68%).

The reduction of redundancy has a negative effect on the perception, especially at long distances. The perception is estimated in [13], and in this paper, with the object perception ratio. This metric is defined as the probability to detect an object within the observation time window. A vehicle successfully detects an object if it receives at least one CPM with information about that object during the observation time window. In the low traffic density scenario, [13] showed that the distance at which an object perception ratio of 0.95 is obtained decreases from 338m (baseline) to 322m (RM with *P_Threshold*=1m) or to 244m (RM with *P_Threshold*=4m). Similar trends were obtained for higher traffic densities.

The application of a redundancy mitigation technique has also an important impact on the number of CPMs generated and their size [13]. The number of CPMs generated per second by RM is between 11% and 52% lower than the baseline for low density, and between 21% and 64% for high density [13]. The number of objects included in each CPM is also significantly decreased with RM. The percentage of CPMs that contain only one object is increased from 23% (baseline) to 42% (RM with *P_Threshold*=1m) or to 74% (RM with *P_Threshold*=4m) in the low traffic density scenario, and similar trends were obtained in high density. While this reduction allows RM to reduce the channel load, the generation of CPMs with a low number of objects is not efficient, because of the large of overhead in each CPM (protocol headers, Management Container, and Originating Vehicle Container).

3.4. Look-Ahead

The Look-Ahead (LA) technique was first proposed in [12] and is now part of ETSI specifications [8][9]. LA extends the baseline generation rules to generate less frequent CPMs that contain a higher number of detected objects. Its goal is to reduce the amount of CPMs generated to reduce the communications overhead without reducing the amount of information transmitted about the detected objects. To do so, LA groups the information about the detected objects into a smaller number of CPMs of larger size. LA is triggered every time a CPM must be generated by the baseline generation rules. Then, LA adds to this CPM the objects that it predicts will be included in the next CPM, considering that the detected objects maintain their current acceleration. To this aim, LA estimates the following parameters for the objects that do not currently satisfy the baseline generation rules:

$$Next\,\Delta P = \Delta P + S\cdot T\_GenCpm + 0.5\cdot A\cdot T\_GenCpm^2 \quad (1)$$

$$Next\,\Delta S = \Delta S + A\cdot T\_GenCpm \quad (2)$$

$$Next\,\Delta T = \Delta T + T\_GenCpm \quad (3)$$

where S and A are the current speed and acceleration of the detected object. LA includes in the current CPM those detected objects that satisfy $Next\,\Delta P$>4m or $Next\,\Delta S$>0.5m/s or $Next\,\Delta T$>1s. These objects would not be initially included in the current CPM following the baseline generation rules, but are included because LA estimates that they will that satisfy the baseline generation rules in the next CPM. As a result, LA avoids that these objects generate a new CPM in the next *T_GenCpm*, and it then reduces the number of generated CPMs of small size.

The work in [12] demonstrated that LA is able to reduce the number of CPMs generated per second by 39%-44% in highway environments, and by 32%-33% in urban environments, compared to the baseline generation rules. The average number of objects included in a CPM is increased by LA by 95%-110% (highway) and 63%-71% (urban) compared to the baseline. As a consequence, the overhead in the transmitted CPMs is significantly lower with LA than with the baseline (49% vs 76%). The reduction of the overhead decreases the CBR by 11%-16% (highway) and by 22%-24% (urban). One important aspect to highlight is that LA reduces the CBR without reducing the rate at which each vehicle transmits each object. In fact, LA increases the average number of times that a detected object is reported in a CPM compared to the baseline generation rules (by 20% and 10% in the highway and urban scenarios, respectively). As a consequence, LA is able to improve the object perception ratio compared to the baseline generation rules, while reducing the channel load.

The application of LA has also an effect on redundancy. LA increases the redundancy by 21%-30% compared to the baseline generation rules [8]. This effect is produced because LA increases the number of times an object is reported in a CPM by each vehicle. The open question is whether this redundancy is needed or can be reduced through the combination of LA and RM to decrease the channel load.

4. Techniques to combine Look-Ahead and Redundancy Mitigation

RM and LA have shown to have advantages but also inefficiencies when applied individually. This section presents the 3 techniques to combine RM and LA for higher effectiveness proposed in this paper. The first two techniques sequentially combine RM and LA, and are thus referred to as LARM and RMLA. The third one, is a more advanced solution that builds on top of *e*RMLA with an enhanced selection of the objects to be included in the CPM.

4.1. LARM

LARM first applies the baseline generation rules and LA, and then applies RM to the objects selected by the baseline generation rules and LA. As a consequence, LARM removes from the CPM all the objects that are considered redundant even though they were selected for inclusion by the baseline generation rules or LA. The objective is to reduce the redundancy generated by LA, and also to benefit from the reduction in channel load achieved with RM.

Figure 1a illustrates with an example the operation of LARM, and how it selects the objects to be included in a CPM. In the example, a transmitting vehicle detects 25 objects, but only 6 of them currently satisfy the baseline generation rules. Additionally, 3 detected objects that do not currently satisfy the baseline generation rules will do so in the next *T_GenCpm*. LARM applies then RM to the 9 objects that currently satisfy the generation rules and in the next *T_GenCpm*. 3 out of these 9 objects are detected as redundant by RM and removed from the CPM. As shown in Figure 1a, RM removes the objects initially selected by the baseline generation rules and objects selected by LA. As a result, LARM finally includes 6 detected objects in the current CPM.

4.2. RMLA

RMLA first applies the baseline generation rules and RM, and then LA. As a result, RM removes first the objects that currently satisfy the baseline generation rules but are considered redundant. Then, LA is applied to the objects that do not currently satisfy the baseline generation rules (i.e., LA is not applied to the ones included in the current CPM or the ones that have been removed by RM). By applying LA after the baseline generation rules and RM, RMLA anticipates the transmission of as many objects as possible in the current CPM, but omits the ones currently considered redundant by RM. However, RMLA can anticipate the transmission of objects that may be deemed redundant because it applies LA after RM. This is one of the main differences with LARM, which applied RM at the end and then removed all objects deemed redundant from the list of objects selected by the baseline generation rules and LA. We should note that RMLA avoids the transmission of a CPM if it only contains redundant objects, because LA is only applied if at least one selected object is not redundant.

Figure 1b illustrates the operation of RMLA and how it selects the objects to be included in a CPM using the same example than in Figure 1a. With RMLA, 6 objects are selected by the baseline generation rules, and 2 of them are removed by the RM technique. Then, LA is applied to the objects that not currently satisfy the baseline generation rules, and it thus adds 3 objects to the CPM. The differences between RMLA and LARM can be clearly observed by comparing Figure 1a and Figure 1b. With LARM, part of the objects anticipated by LA were removed by RM. However, with RMLA, RM is only applied to the objects that currently satisfy the baseline generation rules, but not to the objects anticipated by LA.

4.3. eRMLA

The sequential combination of RM and LA has potential to improve the effectiveness of collective perception. However, we can already anticipate some potential inefficiencies. With RMLA, LA could add to the CPM objects that could be redundant, because RM is applied first. With LARM, RM could remove redundant objects that have been anticipated by LA, because LA is applied first; this could generate frequent CPMs with a low number of objects.

*e*RMLA is designed with two goals. The first one is to avoid the transmission of a CPM if it only contains redundant objects, i.e., when all the objects that satisfy the baseline generation rules are redundant. When this happens, *e*RMLA behaves as RM and RMLA, and it does not generate a CPM. The second goal is to include as many objects as possible in the CPM when a CPM has to be generated (e.g., when at least one object satisfies the baseline generation rules and is not redundant). When this occurs, *e*RMLA behaves as LA and includes all the objects removed by RM plus the ones anticipated by LA.

To achieve its goals, *e*RMLA first applies the baseline generation rules and then RM in order to remove all the objects included in the current CPM that are considered redundant (like in RMLA). If all the objects are removed,

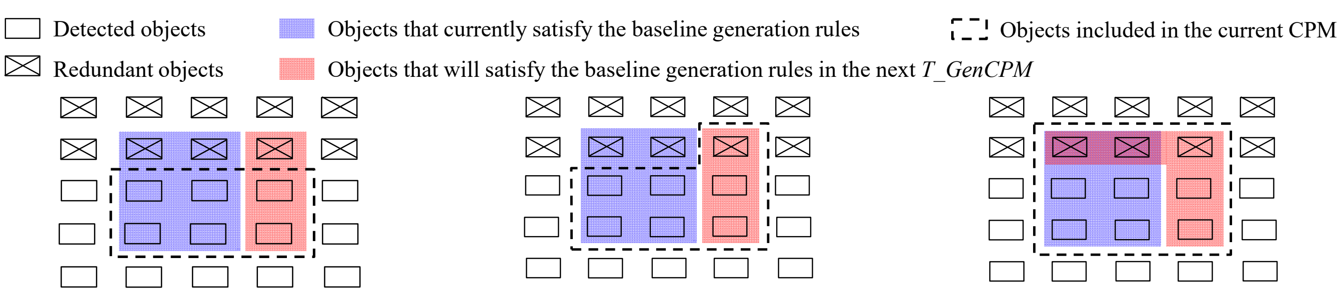

(a) LARM (b) RMLA (c) *e*RMLA

Figure 1. Example to illustrate how each proposal combines the baseline generation rules, RM and LA to build a CPM.

Gokulnath Thandavarayan, Miguel Sepulcre, Javier Gozalvez, and Baldomero Coll-Perales, "Scalable Cooperative Perception for Connected and Automated Driving." Journal of Network and Computer Applications, vol. 216, pp. 103655, May 2023. https://doi.org/10.1016/j.jnca.2023.103655

ALGORITHM I. *e*RMLA
Input: Detected objects
Output: Objects (if any) to include in CPM
Execution: Every *T_GenCpm*

1. Set *flag* = false
2. **For** every detected object **do**
3. **If** the object is a new detected object **then**
4. Include object in current CPM
5. Set flag = true
6. **Else**
7. Calculate ΔP, ΔS and ΔT since the last time the object was included in a CPM
8. **If** ΔP>4 m || ΔS>0.5 m/s || ΔT>1 s **then**
9. Include object in current CPM
10. Set *flag* = true
11. **End If**
12. **End If**
13. **End For**
14. **If** *flag* = true **then**
15. **For** every object included in the current CPM **do**
16. Calculate *ΔP_R* and *ΔS_R* since last time received in a CPM
17. **If** Δ*P_R* < *P_Threshold* && Δ*S_R* < *S_Threshold* **then**
18. Omit object in current CPM
19. **End If**
20. **End For**
21. **If** current CPM does not contain any object **then**
22. Set *flag* = false
23. **End If**
24. **End if**
25. **If** *flag* = true **then**
26. **For** every detected object not included in current CPM **do**
27. Calculate *Next* ΔP, *Next* ΔS and *Next* ΔT
28. **If** *Next* ΔP>4 m || *Next* ΔS>0.5 m/s || *Next* ΔT>1 s **then**
29. Include object in current CPM
30. **End if**
31. **If** the object is a newly detected object **then**
32. Include object in current CPM
33. **End If**
34. **End For**
35. **End If**

then the CPM is not generated. However, if at least one object satisfies the baseline generation rules and is not removed by RM, *e*RMLA applies LA to all detected objects, including those removed by RM. This is in contrast to RMLA that applies LA to all detected objects except those removed by RM. This is an important difference because all objects removed by RM currently satisfy the baseline generation rules. Therefore, when *e*RMLA applies LA to these objects, LA predicts that they will also satisfy the baseline generation rules in the next *T_GenCpm*. This is the case because e.g. the distance traveled since the last time these objects were included in a CPM increases with time. In this context, LA in *e*RMLA will anticipate their transmission in the current CPM. The only objects removed by RM that will not be anticipated by the original LA are the new detected objects. This is the case because LA is able to anticipate only the transmission of objects that have been already transmitted in a previous CPM. *e*RMLA modifies the original LA technique so that it can also anticipate in the current CPM the new detected objects that have been removed by RM.

The operation of *e*RMLA is described in Algorithm I. The baseline generation rules are first executed to identify and select for inclusion in the current CPM the detected objects that satisfy ΔP>4m or ΔS>0.5m/s or ΔT>1s (lines 1-13 of Algorithm I). Then, RM removes from the CPM the objects that are considered redundant (i.e., that satisfy *ΔP_R≤P_Threshold* and Δ*S_R≤ S_Threshold*) as specified in lines 14-23 of Algorithm I. If all objects are removed, the CPM is not generated. When a CPM must be generated (e.g. because at least one object is included in the CPM after applying RM), *e*RMLA triggers LA (lines 25-35 of Algorithm I). LA anticipates and includes in the current CPM the detected objects that satisfy *Next* ΔP>4m or *Next* ΔS>0.5m/s or *Next* ΔT>1s. LA also includes in the current CPM the new detected objects that were included by the baseline generation rules but removed by RM (lines 31-33 of Algorithm I).

Figure 1c illustrates the operation of *e*RMLA and the objects it selects for inclusion in the current CPM using the same example. The figure shows that RM removes two objects from the CPM, i.e., it removes 2 out of 6 detected objects that currently satisfy the baseline generation rules. Since RM does not remove all objects, the CPM must be generated, and LA is applied next. LA anticipates 3 additional detected objects that satisfy the baseline generation rules in the next *T_GenCpm*, plus the 2 objects initially removed by RM. As a result, the CPM generated contains 9 objects in total. The CPM generated in Figure 1c is equal to the CPM generated by LA alone. However, it is important to note that this might not be the case for all CPMs. With *e*RMLA, LA is not triggered if RM removes all the objects that currently satisfy the baseline generation rules. If this is the case, then a CPM is not generated. The objective sought with RM in *e*RMLA is to reduce the number of CPMs generated per second and increase their size compared to when using LA alone.

5. Evaluation scenario and settings

The evaluation is performed using the network simulator ns-3 and the road mobility simulator SUMO. SUMO is used to realistically generate the position and speed of all vehicles during the simulation in the considered scenario. The scenario is a 5 km highway with two driving directions. Three traffic densities (low, medium and high) are simulated as summarized in Table 2. The low and medium traffic densities consider 3 lanes in each direction (6 lanes in total) and the high traffic density considers 4 lanes in each direction (8 lanes in total). The vehicles speed for each traffic density is configured following the statistics of a typical US highways [42]. To avoid boundary effects, the statistics are only taken from the vehicles located within the 2 km around the center of the simulation scenario.

In ns-3, all vehicles are equipped with an ITS-G5 transceiver (based on IEEE 802.11p), and a CPS component that generates CPMs using an onboard sensor with 360º field of view and 150 m range. Object detection is performed in run time in ns-3, considering the sensor shadowing effect in the XY-plane that considers the occlusion caused by nearby vehicles. We assume that the sensors can detect the vehicles that are in their line-of-sight in both driving directions. The CPS component implements the baseline generation rules, as well as the LA, RM, LARM, RMLA and *e*RMLA techniques. Every generated CPM includes all the mandatory headers and containers, i.e. one ITS PDU header, one Management Container and one Originating Vehicle Container (121 Bytes in total). In addition, one Perceived Object Container (35 Bytes) is included per detected object in the CPM if it satisfies the conditions of the applied techniques. Also, one Sensor Information Container (35 Bytes) is included in the CPM every second. The transmitting vehicle dynamically computes the CPM size based on the number of containers in each CPM and their respective sizes. The value of the *T_GenCpm* parameter has been set to 100 ms as default, so the maximum CPM generation rate is 10 Hz. Following [13], the dynamics-based redundancy mitigation technique is implemented considering the following threshold values: *P_Threshold*=1m and *S_Threshold*=0.5m/s.

The transmission power of the ITS-G5 transceiver is 23 dBm and the sensing threshold is set to -85 dBm. All messages are transmitted using the 6 Mbps data rate (i.e., QPSK modulation with ½ code rate) in the 5.9 GHz band. In the simulations, the Winner+ B1 radio propagation model has been considered following the 3GPP V2X guidelines [41]. Table 3 summarizes the main communication parameters.

Table 2. Traffic scenarios

Parameter	Traffic density scenarios		
	Low	**Medium**	**High**
Number of lanes	6	6	8
Traffic density	120 veh/km	180 veh/km	240 veh/km
Speed per lane	70 km/h 66 km/h 59 km/h	50 km/h for all lanes	50 km/h for all lanes

Table 3. Communication parameters

Parameter	Values
Transmission power	23 dBm
Antenna gain (tx and rx)	0 dBi
Channel bandwidth	10 MHz
Carrier frequency	5.9 GHz
Noise figure	9 dB
Energy detection threshold	-85 dBm
Data rate	6 Mbps (QPSK 1⁄2)

6. Evaluation

The proposed techniques are first analyzed in section 6.1 without including congestion control. The goal of this first analysis is to understand well how the standalone and combination techniques behave before considering any additional influences. Then, the proposed techniques will be evaluated considering the impact of congestion control and the coexistence of CPMs with other messages in section 6.2. Congestion control can influence the generation and transmission of CPMs.

6.1. Without congestion control

To compare the different techniques from the perception point of view, we use the object perception ratio metric. This metric is defined as the probability to successfully detect an object within a given observation time window. The time window has been set equal to 300 ms, which is the time required by the baseline generation rules for a vehicle to send an update about a detected object considering the speed of the object in the considered scenarios.

Figure 2 plots the object perception ratio obtained with the different techniques under all traffic densities under evaluation. In the low traffic density scenario (Figure 2a), all the techniques achieve a very high object perception ratio up to around 300 m. For larger distances, the object perception ratio decreases due to propagation and interference effects. When the traffic density augments (Figure 2b and Figure 2c), the object perception ratio decreases in general because the channel load and interferences increase. For all the techniques, the object perception ratio is still very high up to 250 m for the medium density (Figure 2b) and up to 200 m for the high density (Figure 2c). RM and LA work as expected: LA significantly improves the perception for low and medium densities compared to the baseline, and RM improves the perception for medium and high densities. LARM and RMLA outperform the baseline technique for medium and high traffic densities, but only provide a perception higher than LA for the highest density. *e*RMLA achieves the highest perception levels for all traffic densities, improving the perception achieved by LARM and RMLA by 12%-27% (relative distance increase for an object perception ratio of 0.95 using the baseline as a reference). It also shows that *e*RMLA is scalable because the object perception ratio obtained is nearly maintained when the traffic density increases.

*e*RMLA achieves the best perception because it is able to generate the lowest channel load (Table 4) and achieves the highest transmission efficiency. More specifically, *e*RMLA is more efficient because it is able to generate longer CPMs (higher number of objects, Table 5) and less frequently (Table 6). As a consequence, *e*RMLA presents around 50% improvement in terms of channel load compared to the baseline generation rules, and 1%-12% improvement compared to LARM and RMLA (Table 4). In addition, *e*RMLA is able to reduce the redundancy achieved compared to the baseline generation rules and LA, as it can be observed in Figure 3 for medium traffic density. This figure plots the detected object redundancy, computed as the number of CPMs that a vehicle receives with information about the same object in a given observation time window (300 ms in this study). RMLA and LARM achieve the lowest detected object redundancy, and they behave nearly as RM, that focuses on controlling redundancy without trying to improve the transmission efficiency. *e*RMLA thus provides a higher redundancy than LARM and RMLA, but with lower channel load, demonstrating again its higher efficiency (it transmits more information about detected objects and less overhead). Similar trends are achieved for the other densities.

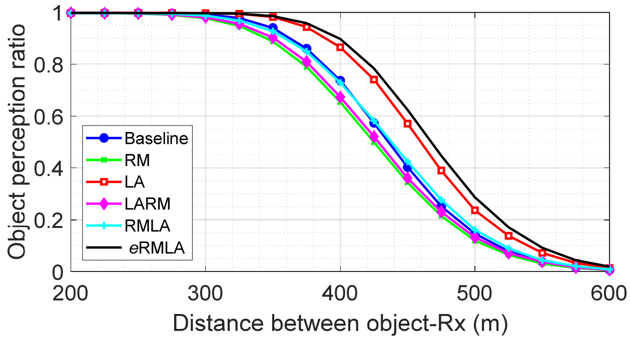

(a) Low traffic density

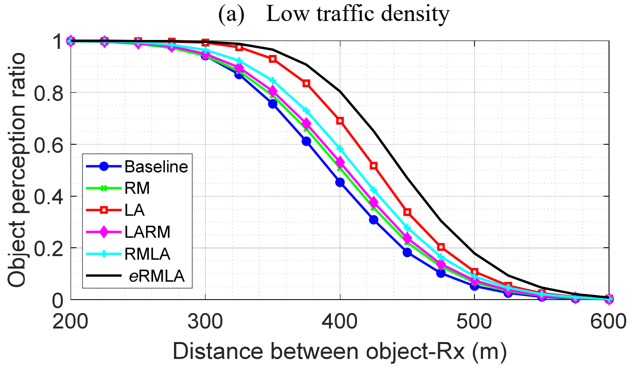

(b) Medium traffic density

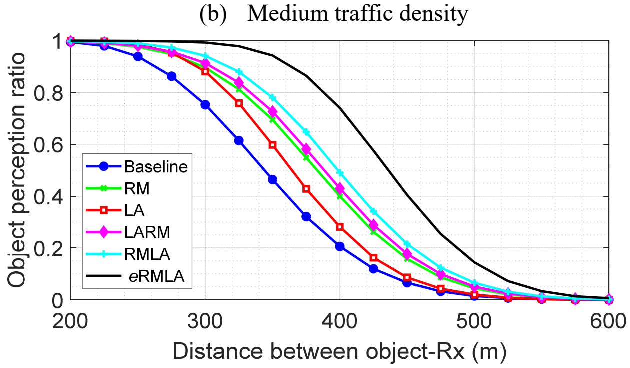

(c) High traffic density

Figure 2. Object perception ratio as a function of the distance between the detected object and the vehicle receiving the CPM.

Table 4. Average CBR (Channel Busy Ratio)

Techniques	Traffic Density		
	Low	Medium	High
Baseline	49.4%	64.4%	82.1%
RM	29.1%	35.5%	49.0%
LA	41.4%	56.5%	82.7%
LARM	27.3%	32.4%	46.0%
RMLA	25.8%	30.0%	43.0%
*e*RMLA	24.4%	29.0%	42.0%

Table 5. Average number of objects included in each CPM

Techniques	Traffic density		
	Low	Medium	High
Baseline	5.1	5.3	6.4
RM	1.9	1.8	1.9
LA	10.4	11.0	12.3
LARM	2.2	2.1	2.1
RMLA	3.4	3.1	3.2
*e*RMLA	13.8	14.1	17.4

Table 6. Average number of CPMs generated per second

Technique	Traffic density		
	Low	Medium	High
Baseline	9.6	9.4	9.6
RM	7.1	6.2	6.7
LA	5.4	5.4	6.2
LARM	6.4	5.6	6.1
RMLA	5.4	4.7	5.1
*e*RMLA	2.6	2.2	2.1

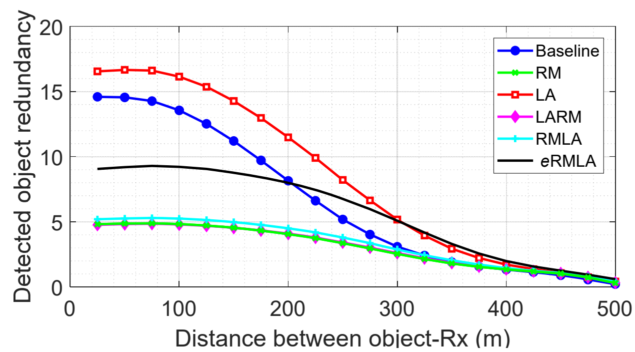

Figure 3. Detected object redundancy for medium traffic density.

The effectiveness of cooperative perception also depends on the latency experienced when exchanging the sensed data. We then evaluate the information age metric, which is defined as the difference between the time a CPM is generated and the time it is received. This metric provides information about the freshness of the received information. Figure 4 represents the mean information age obtained for all the techniques and traffic densities. The information age is highly influenced by the channel access mechanism, the channel load and also the packet length. The propagation delay can be considered negligible. Figure 4 shows that the information age is below 2 ms for all techniques under low and medium traffic densities. For the high-density scenario, the information age increases, especially for the baseline generation rules and LA because they generate higher channel load levels. Figure 4 shows that *e*RMLA results in a slightly higher information age compared to LARM and RMLA (around 0.8ms higher in the worst case). This difference is produced because *e*RMLA generates significantly larger messages. In fact, *e*RMLA has the highest average number of objects included in each CPM (Table 5), and most of its CPMs include the Sensor Information Container because it has the lowest CPM generation rate (Table 6). This increases the message size and thus the time required to receive it. This additional time is the origin of the increment of the information age observed in Figure 4 for *e*RMLA compared to LARM and RMLA.

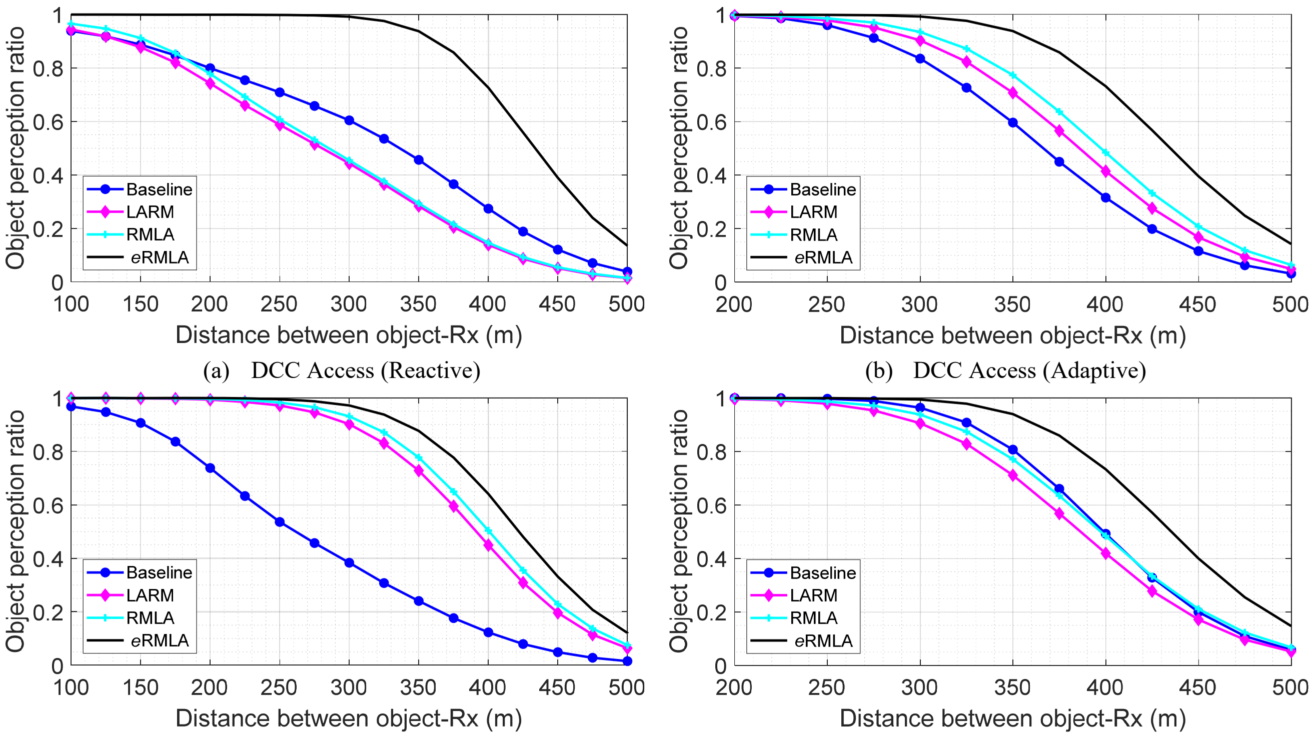

(a) DCC Access (Reactive)

(b) DCC Access (Adaptive)

(c) DCC Access (Reactive) + Facilities

(d) DCC Access (Adaptive) + Facilities

Figure 5. Object perception ratio as a function of the distance between the detected object and the vehicle receiving the CPM for different DCC configurations in the high traffic density scenario.

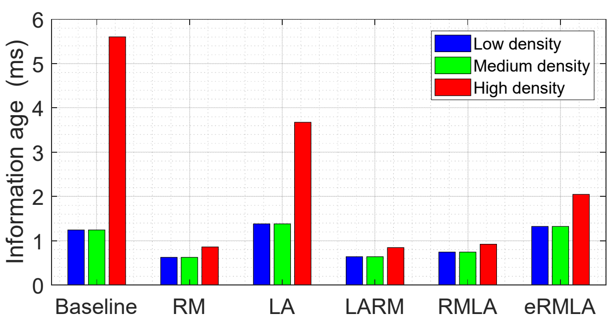

Figure 4. Average information age of CPMs for all traffic densities.

6.2. Evaluation with congestion control

This section evaluates the performance achieved with the proposed combination techniques when considering the impact of congestion control and the coexistence of CPMs with other messages. These two aspects can significantly alter the generation and transmission of CPMs and can thus influence the performance of the collective perception service. This analysis is key for the integration of the proposed techniques in the protocol stack and their potential incorporation in standards. For this analysis, this study considers the DCC framework and the impact of DCC Access [39] and DCC Facilities [40]. This evaluation focuses on the high traffic density scenario that generates the highest channel load levels. The evaluation in Section 6.2.1 considers that only CPMs are transmitted in the channel, while Section 6.2.2 considers that the transmission of CAMs and CPMs share the same radio channel, since ETSI has not decided yet the channel for the CPMs [43].

6.2.1. Only CPMs

The impact of DCC on the object perception ratio is shown in Figure 5. Figure 5a and b show the results obtained with DCC Access only, while Figure 5c and d show the results when both DCC Access and Facilities are considered. The results depicted in this figure clearly show that the highest object perception ratio is again obtained with *e*RMLA independently of the DCC configuration (relative distance improvement between 17% and 325% compared to LARM and RMLA using the baseline as a reference). Figure 5a shows that the packets dropped by DCC Access Reactive significantly reduce the object perception ratio. This degradation partially results from the well-known synchronization[1] problem [37], observed with DCC Access Reactive that increases the probability of packet collisions. Only *e*RMLA is not affected by packet dropping because of its low CPM generation rate (Table 6), that allows that CPMs do not wait in the DCC Access queue and are always transmitted. DCC Access Adaptive (Figure 5b) does not negatively impact any of the proposed techniques due to their low CBR (Table 4).

Figure 5c shows that the performance achieved with *e*RMLA with DCC Access Reactive and DCC Facilities is slightly reduced at larger distances compared to the scenario

[1] With the DCC Access Reactive approach vehicles tend to synchronize with each other and transmit nearly at the same time [45]. This effect provokes that the DCC Access Reactive approach generates a significant amount of packet collisions.

Gokulnath Thandavarayan, Miguel Sepulcre, Javier Gozalvez, and Baldomero Coll-Perales, "Scalable Cooperative Perception for Connected and Automated Driving." Journal of Network and Computer Applications, vol. 216, pp. 103655, May 2023. https://doi.org/10.1016/j.jnca.2023.103655

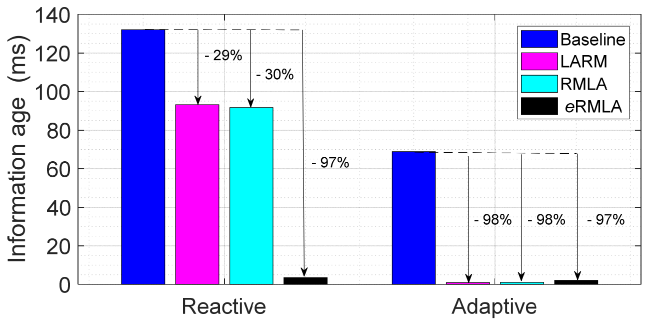

(a) DCC Access

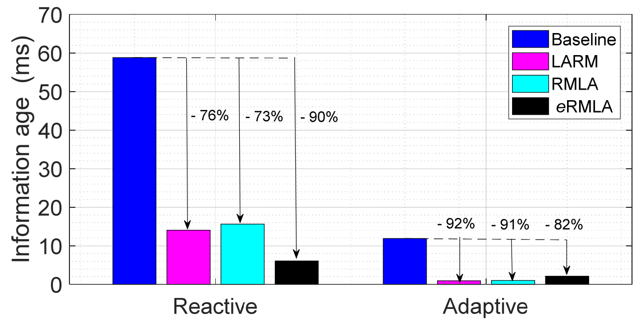

(b) DCC Access + DCC Facilities

Figure 6. Average information age for CPMs for different DCC configurations when CPMs are transmitted on the same channel in the high traffic density scenario.

with only DCC Access (Figure 5a) because of a small reduction in the CPM transmission rate. Figure 5c also shows that the object perception ratio of LARM and RMLA significantly increases compared to the scenario when only DCC Access is used. This improvement is mainly produced because LARM and RMLA generate aperiodic CPMs with DCC Facilities (LARM and RMLA sometimes omit CPMs due to the use of redundancy mitigation and Look-Ahead). This reduces the probability that all vehicles simultaneously transmit and hence combats the synchronization problem.

Figure 5d shows that the object perception ratio obtained with the three proposed techniques when jointly considering DCC Access Adaptive and DCC Facilities is close to the one obtained when only DCC Access Adaptive is used (Figure 5b). This is the case because the proposed techniques reduced the CBR below 50% and thus DCC Access Adaptive was not activated. However, the object perception ratio achieved with the baseline generation rules improves using DCC Access Adaptive and Facilities. This is the case because DCC Facilities reduces the CPM transmission rate and increases the number of objects in each CPM. This eventually reduces the CBR and improves the percentage of CPMs successfully received.

We have also evaluated the information age experienced with the baseline generation rules and the proposed techniques with all the DCC configurations evaluated. The results are shown in Figure 6. The figure shows in percentage the difference between the baseline and each of the proposed techniques. When only DCC Access is used (Figure 6a), the information age generally increases (up to 130ms) compared with the scenario without DCC (less than 6 ms) because of the waiting time of the packets at the DCC Access queues. The information age obtained with *e*RMLA and DCC Access Reactive is significantly lower than the baseline (around 130ms lower) and also lower than LARM and RMLA (around 90ms lower). This improvement is achieved by *e*RMLA because it reduces the channel load, and packets are not dropped or queued at the access layer. With DCC Access Adaptive, the proposed techniques reduce the information age compared to the baseline by around 70ms, since the channel load they generate is not sufficiently high to activate DCC Access. In this case, *e*RMLA achieves an information age that is around 1ms higher than LARM and RMLA, but, as previously discussed, this increment can be considered negligible.

When DCC Access and Facilities are used (Figure 6b), the information age is generally decreased compared to when only DCC Access is used. This is the case because DCC Facilities influences the CPM generation and CPMs tend to wait less time in the DCC Access queues. However, the information age can still be significantly higher than when DCC is not used. This is particularly the case for the baseline generation rules, and when LARM and RMLA are considered with the Reactive approach. In this case, *e*RMLA achieves the lowest information age (53ms lower than the baseline and around 8ms lower than LARM and RMLA). When DCC Facilities is combined with DCC Access Adaptive, the three proposed techniques achieve nearly the same information age than in the scenario without DCC (Figure 4).

6.2.2. CPMs and CAMs

This section complements previous evaluations with a scenario where all vehicles generate and transmit CAMs and CPMs on the same radio channel. The transmission of CAMs increases the channel load and activates DCC with higher probability. The CAMs are generated following the ETSI generation rules defined in [38] and their size is set equal to 350 bytes [44]. The default *T_GenCam* parameter has been set to 0.1 s so the maximum CAM rate is 10 Hz. However, DCC Facilities may adapt *T_GenCam*. CAMs and CPMs are configured with the same DCC profile so that they have the same priority and share the channel equally.

Figure 7 depicts the object perception ratio achieved with all DCC configurations when both CAMs and CPMs are transmitted on the same channel. The figure shows that the highest perception is again obtained withe *e*RMLA for all DCC configurations. However, the perception achieved decreases compared to the scenario without CAMs since CAMs consume part of the bandwidth and generate additional interferences. This degradation is also observed in LARM and RMLA. The baseline generation rules achieve

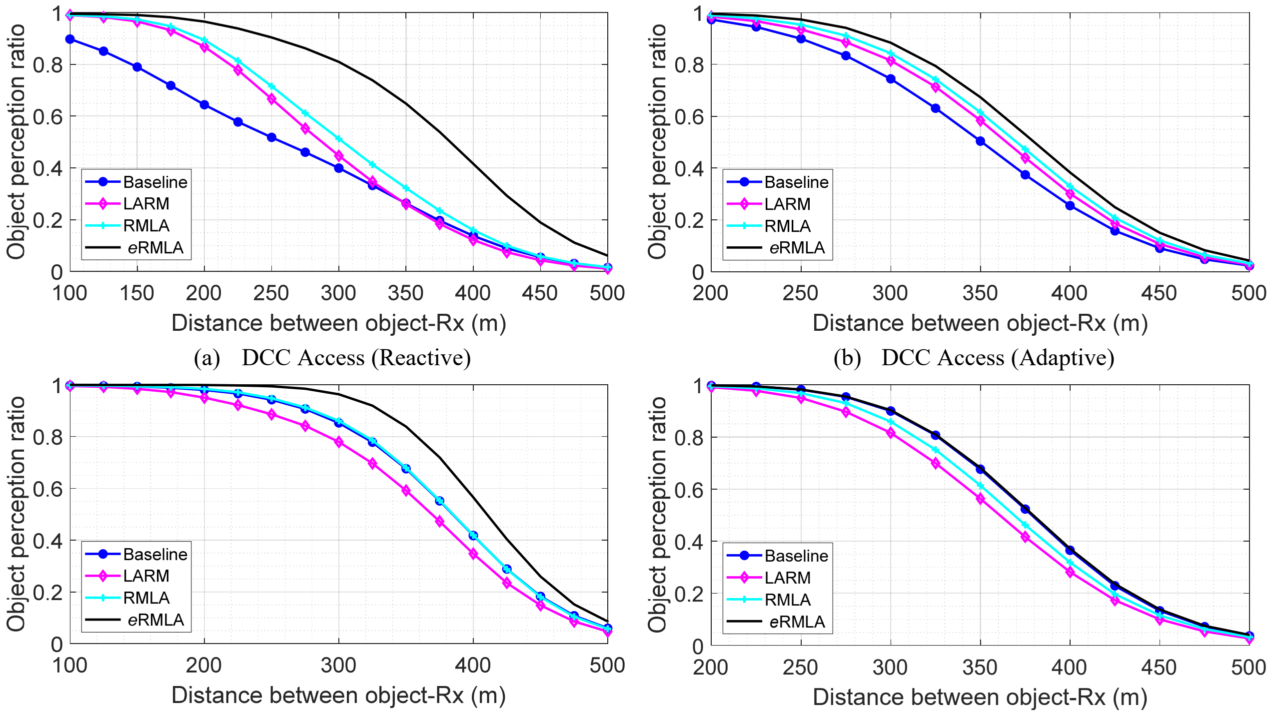

(a) DCC Access (Reactive)
(b) DCC Access (Adaptive)
(c) DCC Access (Reactive) + Facilities
(d) DCC Access (Adaptive) + Facilities

Figure 7. Object perception ratio as a function of the distance between the detected object and the vehicle receiving the CPM for different DCC configurations when CAMs and CPMs are transmitted on the same channel in the high traffic density scenario.

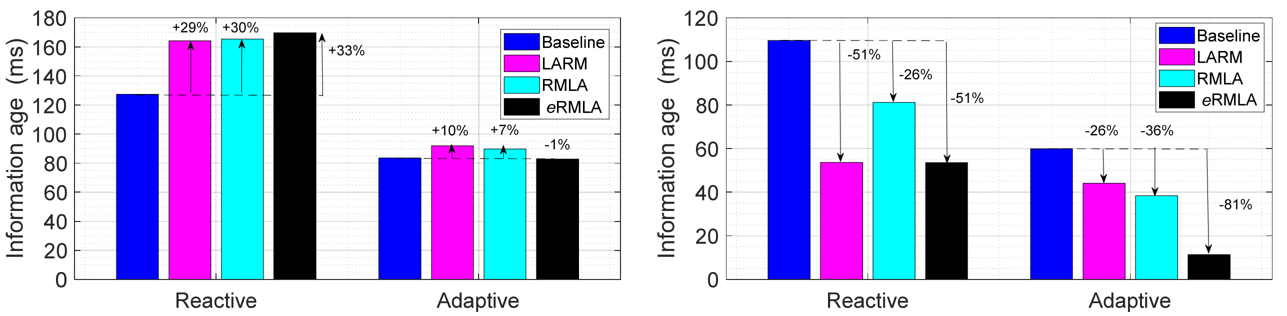

(a) DCC Access
(b) DCC Access + DCC Facilities

Figure 8. Average information age for CPMs for different DCC configurations when CAMs and CPMs are transmitted on the same channel in the high traffic density scenario.

the lowest perception ratio when only DCC Access is used (especially with DCC Access Reactive) due to the synchronization problem. The differences among the proposed techniques are particularly high with DCC Access Reactive (Figure 7a), especially because *e*RMLA is able to significantly reduce the number of CPMs generated per second. In fact, *e*RMLA is able to increase the distance at which an object perception ratio of 0.95 is achieved by 96% and 120% compared to RMLA and LARM, respectively, using the baseline as a reference. With DCC Access Adaptive (Figure 7b), the differences are reduced, but still a difference of 7% and 13% is observed.

Combining DCC Access and DCC Facilities generally improves the object perception ratio compared with only using DCC Access for all the techniques considered (Figure 7c and d). A significant increase of the object perception ratio is observed especially for LARM and RMLA when DCC Facilities is used with the Reactive approach because the lower number of CPMs dropped. DCC Facilities helps alleviating the synchronization problem observed with DCC Access Reactive even when the baseline generation rules are considered (Figure 7a and Figure 7c). On the other hand, similar results are obtained with and without DCC Facilities when using DCC Access Adaptive (Figure 7b and Figure 7d). The main difference is that the use of DCC Facilities reduces significantly the CPM generation rate of the baseline generation rules, reduces the dropped CPMs and increases the CPM size. This improves the perception achieved with the baseline generation rules (Figure 7d), which in fact achieves a higher perception than LARM when DCC Facilities is used. *e*RMLA increases the distance at which an object perception ratio of 0.95 is achieved by 44% compared to LARM and by 25% compared to RMLA with DCC Access

Reactive and Facilities. This increment is 6% and 10% with the Adaptive approach.

Figure 8 reports the average information age obtained with baseline generation rules and the proposed techniques with all DCC configurations when both CAMs and CPMs are transmitted on the same channel. The figure shows in percentage the difference between the baseline and each of the proposed techniques. The transmission of CAMs in the same channel increases the information age for all the techniques and DCC configurations, compared to the scenario without CAMs (Figure 6), because CAMs increase the channel load. When only DCC Access is used (Figure 8a), the baseline generation rules slightly reduces the information age compared with the proposed techniques, but this is achieved at the expense of a lower object perception ratio (Figure 7). In this case, the information age obtained by *e*RMLA is around 5ms (or 4%) higher than LARM and RMLA with DCC Access Reactive, but around 7ms (or 8%) lower with Adaptive. The use of DCC Access Reactive clearly shows the challenges of DCC to support the considered V2X services given that the average information age achieved is higher than 120ms.

When DCC Facilities is combined with DCC Access (Figure 8b), the information age is in general reduced, especially when using the proposed techniques. In this scenario, when the Reactive approach is considered, LARM and *e*RMLA reduce the average information age compared to the baseline generation rules by around 51%. With the Adaptive approach, *e*RMLA reduces the average information age by a factor of 5 compared to the baseline generation rules (81% reduction), and by a factor of 4 compared to LARM and RMLA, approximately. This is the case because *e*RMLA generates CPMs at a lower frequency, and most of the time the DCC Access gate is open when a CPM is generated. This result demonstrates that *e*RMLA can facilitate the transmission of CPMs with a low latency even under the presence of CAMs.

7. Discussion

The results obtained have shown important trade-offs among the proposed techniques. In the scenario without DCC, *e*RMLA provides the highest perception for all traffic densities (10%-42% higher than the baseline and 12%-27% higher than the other proposed techniques). It also reduces the channel load (around 50% lower than the baseline and 1%-12% lower than the other proposed techniques). It is true that *e*RMLA has a higher average information age compared to LARM and RMLA when DCC is not applied. But in this scenario the increase of the information age obtained by *e*RMLA is in the order of the packet duration (0.2ms-0.8ms higher), and can be considered almost negligible considering its gains in terms of perception and channel load.

The results obtained also show that DCC has an important impact on the CPM generation and transmission. The performance differences depend significantly on the DCC configuration. In terms of perception, the best results are always obtained by *e*RMLA, independently of the DCC configuration, and whether CAMs are transmitted in the same channel or not. Without CAMs, *e*RMLA increases the perception by 17%-325% compared to the other proposed techniques. With CAMs, this improvement is between 6% and 120%, depending on the DCC configuration.

The results obtained in terms of information age require a more detailed analysis. When the load is sufficiently low so that DCC is not activated (in the scenario without CAMs and with the Adaptive approach), the results obtained match with the results without DCC for the proposed techniques; that is, *e*RMLA obtains an information age that is 0.2ms-0.8ms higher than the other proposed techniques, which cannot be considered a significant increase given its gains in terms of perception. When the load increases and DCC is activated, *e*RMLA is able to achieve a lower information age than the other proposed techniques (reductions by 56%-96% without CAMs and up to 70% with CAMs), except in one DCC configuration (DCC Access Reactive and Facilities with CAMs and CPMs in the same channel). In this specific DCC configuration, *e*RMLA improves significantly the perception, but increases the information age by 4ms (3%) compared to the other proposed techniques. This DCC configuration is very specific because it generates the highest information age for all the techniques (higher than 120ms), clearly showing the need to evolve the existing DCC framework (e.g., synchronizing the message generation with the gatekeeper at the access layer to achieve low information age).

In summary, the results obtained demonstrate the *e*RMLA achieves best overall performance and is capable to enhance the perception of connected automated vehicles, independently of the traffic density, DCC configuration and coexistence with other messages in the same channel.

8. Conclusions

This paper has proposed and evaluated three techniques designed to improve the effectiveness of cooperative or collective perception services for connected and automated vehicles while ensuring their scalability. The proposed techniques combine, for the first time, baseline message generation rules for cooperative perception messages with mechanisms to control the redundancy and to efficiently organize the information about detected objects. The study has evaluated the effectiveness and scalability of the proposed techniques under different traffic density scenarios and considering the integration with congestion control mechanisms and the coexistence of cooperative perception messages and awareness messages. Their performance has been compared with baseline message generation rules and two existing techniques, all of them part of ETSI's specifications. The conducted evaluation has demonstrated that the proposed techniques improve the perception and reduce the channel load in most of the scenarios evaluated. They are thus able to improve the scalability of cooperative perception services and leave more available bandwidth for other V2X services. The proposed techniques also reduce the information age compared to the baseline in all the scenarios considered, except in one congestion control configuration that is not able to efficiently support cooperative perception messages and awareness messages.

The conducted evaluation has demonstrated that the most effective way to combine baseline collective perception generation rules with redundancy control and look-ahead mechanisms is by first applying the generation rules, then redundancy control and finally the look-ahead mechanisms to all detected objects, including those initially removed by the redundancy control scheme. This combination, referred to *e*RMLA in this study, achieves the highest perception and lowest channel load in all the scenarios and configurations evaluated thanks to a better balance between object redundancy and communications overhead. It is also able to reduce the information age in most of the scenarios and congestion control configurations; it only produces a relatively small increase in some corner cases, which it is considered acceptable given its gains in terms of perception and channel load.

As future work, one potential way to further improve the combination method proposed in this study would be to incorporate machine learning techniques. Machine learning could improve the identification of redundant objects and better support the anticipation of objects for the look-ahead mechanism. The obtained results have also demonstrated the need to improve the operation of DCC Access and Facilities to reduce the information age. This could be addressed, for example, by synchronizing the gate-opening times at the Access layer with the message generation at the Facilities layer.

Acknowledgement

This work was supported by the TransAID Project through the Horizon 2020 Framework Programme under Agreement 723390 and by the Spanish Ministry of Science, Innovation and Universities, AEI and FEDER funds through the project TEC2017-88612-R.

References

1. E. Yurtsever, J. Lambert, A. Carballo and K. Takeda, “A Survey of Autonomous Driving: Common Practices and Emerging Technologies”, *IEEE Access*, vol. 8, pp. 58443-58469, March 2020. https://doi.org/10.1109/ACCESS.2020.2983149
2. Kakkasageri, M. S., and Sunilkumar S. Manvi. “Information management in vehicular ad hoc networks: A review”, *Journal of Network and Computer Applications,* vol. 39, pp. 334-350, March 2014.
3. K. Jo *et al*.,“Development of autonomous car Part I: Distributed system architecture and development process”, *IEEE Transactions on Industrial Electronics*, vol. 61, no. 12, pp. 7131–7140, December 2014. https://doi.org/10.1109/TIE.2014.2321342
4. F. Rosique, P. J. Navarro, C. Fernández and A. Padilla, “A Systematic Review of Perception System and Simulators for Autonomous Vehicles Research”, *Sensors*, no. 19, pp. 648, January 2019. https://doi.org/10.3390/s19030648
5. J. Van Brummelen, M. O’Brien, D. Gruyer and H. Najjaran, “Autonomous vehicle perception: The technology of today and tomorrow”, *Transportation Research Part C*, vol. 89, pp. 384-406, April 2018. https://doi.org/10.1016/j.trc.2018.02.012
6. E. Marti, M. A. de Miguel, F. Garcia and J. Perez, “A Review of Sensor Technologies for Perception in Automated Driving”, *IEEE Intelligent Transportation Systems Magazine*, vol. 11, no. 4, pp. 94-108, September 2019. https://doi.org/10.1109/MITS.2019.2907630
7. SAE, Cooperative Perception System, J2945/8, 2018. https://www.sae.org/standards/content/j2945/8/ (Accessed June 2021), Online.
8. ETSI ITS, “Intelligent Transport System (ITS); Vehicular Communications; Basic Set of Applications; Analysis of the Collective Perception Service (CPS)”, ETSI TR. 103 562 V2.1.1, December 2019.
9. ETSI ITS, “Intelligent Transport System (ITS); Vehicular Communications; Basic Set of Applications; Specification of the Collective Perception Service”, ETSI TS 103 324 V0.0.20 (draft), 2021.
10. K. Garlichs, H. J. Günther and L. Wolf, “Generation Rules for the Collective Perception Service”, *Proc. IEEE Vehicular Networking Conference (VNC)*, Los Angeles, CA, USA, pp. 1-8, 4-6, December 2019. https://doi.org/10.1109/VNC48660.2019.9062827
11. G. Thandavarayan, M. Sepulcre, J. Gozalvez., “Analysis of Message Generation Rules for Collective Perception in Connected and Automated Driving”, *Proc. IEEE Intelligent Vehicle Symposium (IV)*, Paris (France), pp. 134-139, 9-12, June 2019. https://doi.org/10.1109/IVS.2019.8813806
12. G. Thandavarayan, M. Sepulcre and J. Gozalvez, “Generation of Cooperative Perception Messages for Connected and Automated Vehicles”, *IEEE Transactions on Vehicular Technology*, vol. 69, no. 12, pp. 16336-16341, December 2020. https://doi.org/10.1109/TVT.2020.3036165
13. G. Thandavarayan, M. Sepulcre and J. Gozalvez, “Redundancy Mitigation in Cooperative Perception for Connected and Automated Vehicles”, *Proc. IEEE 91st Vehicular Technology Conference (VTC2020-Spring)*, pp. 1-5, June 2020. https://doi.org/10.1109/VTC2020-Spring48590.2020.9129445
14. R. Fukatsu and K. Sakaguchi, “Automated Driving with Cooperative Perception Using Millimeter-Wave V2V Communications for Safe Overtaking”, *Sensors*, no. 21(8), pp. 2659, April 2021. https://doi.org/10.3390/s21082659
15. A. Merwaday *et al*., “Demo: Infrastructure Assisted Efficient Collective Perception Service for Connected Vehicles”, Proc. IEEE Vehicular Networking Conference (VNC), pp. 119-120, December 2021. https://doi: 10.1109/VNC52810.2021.9644671
16. H. J. Günther, B. Mennenga, O. Trauer, R. Riebl, L. Wolf, “Realizing collective perception in a vehicle”, *Proc. IEEE Vehicular Networking Conference (VNC)*, Columbus, Ohio, USA, pp. 1-8, December 2016. https://doi.org/10.1109/VNC.2016.7835930
17. C. Allig and G. Wanielik, “Dynamic Dissemination Method for Collective Perception”, *Proc. IEEE Intelligent Transportation Systems Conference (ITSC)*, Auckland, New Zealand, pp. 3756-3762, 27-30, October 2019. https://doi.org/10.1109/ITSC.2019.8917266
18. K. Furukawa *et al*., “Controlling Sensor Data Dissemination Method for Collective Perception in VANET”, *Proc. IEEE International Conference on Pervasive Computing*, Kyoto, Japan, pp. 753-758, March 2019. https://doi.org/10.1109/PERCOMW.2019.8730601
19. S. Faiz *et al*., "Increasing vehicles perception using cooperative relaying and priority-based beaconing," *Proc. IEEE 22nd International Conference on High Performance Switching and Routing (HPSR)*, pp. 1-6, June 2021. https://doi: 10.1109/HPSR52026.2021.9481799
20. T. Higuchi, M. Giordani, A. Zanella, M. Zorzi, O. Altintas, “Value-Anticipating V2V Communications for Cooperative Perception”, *Proc. IEEE Intelligent Vehicles Symposium (IV)*, Paris, France, pp. 1947-1952, June 2019. https://doi.org/10.1109/IVS.2019.8814110
21. S. Aoki, T. Higuchi, O. Altintas, “Cooperative Perception with Deep Reinforcement Learning for Connected Vehicles”, *Proc. IEEE Intelligent Vehicles Symposium (IV)*, pp. 328-

334, October 2020. https://doi.org/10.1109/IV47402.2020.9304570
22. S. M. O. Gani *et al*., "A Study of the Effectiveness of Message Content, Length, and Rate Control for Improving Map Accuracy in Automated Driving Systems", *IEEE Trans. on ITS*, vol. 20, no. 2, pp. 405-420, February 2019. https://doi.org/10.1109/TITS.2018.2812847
23. S. Li and V. A. Wolff, "Tracking Accuracy Based Generation Rules of Collective Perception Messages," *IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*, pp. 4157-4162, Macau, China, Nov 2022. https://doi.org/10.1109/ITSC55140.2022.9922147
24. H. J. Günther *et al*., "Collective perception and decentralized congestion control in vehicular ad-hoc networks", *Proc. IEEE Vehicular Networking Conference (VNC)*, Columbus, Ohio, USA, pp. 1-8, December 2016. https://doi.org/10.1109/VNC.2016.7835931
25. F. A. Schiegg *et al*., "Analytical Performance Evaluation of the Collective Perception Service in IEEE 802.11p Networks", *Proc. IEEE Wireless Communications and Networking Conference (WCNC)*, pp. 1-6, May 2020. https://doi.org/10.1109/WCNC45663.2020.9120490
26. F. A. Schiegg *et al*., "Analytical Performance Evaluation of the Collective Perception Service in C-V2X Mode 4 Networks", *Proc. IEEE Intelligent Transportation Systems Conference (ITSC)*, pp. 181-188, October 2019. https://doi.org/10.1109/ITSC.2019.8917214
27. F. A. Schiegg, I. Llatser, D. Bischoff, G. Volk, "Collective Perception: A Safety Perspective", *Sensors*, no.21(1) pp. 159, 2020. https://doi.org/10.3390/s21010159
28. G. Pereira, P. M. d'Orey and A. Aguiar, "Poster: Cooperative Perception Platform for Intelligent Transportation Systems", *Proc. IEEE Vehicular Networking Conference (VNC)*, pp. 1-2, Dec 2020. https://doi: 10.1109/VNC51378.2020.9318404.
29. M. Tsukada *et al*., "AutoC2X: Open-source software to realize V2X cooperative perception among autonomous vehicles", *Proc. IEEE 92nd Vehicular Technology Conference (VTC2020-Fall),* pp. 1-6, Dec 2020. https://doi: 10.1109/VTC2020-Fall49728.2020.9348525.
30. Shan, M. *et al*., "Demonstrations of cooperative perception: safety and robustness in connected and automated vehicle operations", Sensors, v.21(1), p.200, Dec 2020. https://doi.org/10.3390/s21010200
31. H. Huang, H. Li, C. Shao, T. Sun, W. Fang and S. Dang, "Data Redundancy Mitigation in V2X Based Collective Perceptions", *IEEE Access*, vol. 8, pp. 13405-13418, January 2020. https://doi.org/10.1109/ACCESS.2020.2965552
32. C. Ameni, P. Merdrignac, and O. Shagdar "Context-Aware Content Selection and Message Generation for Collective Perception Services", *Electronics,* v.10, no. 20: 2509, Oct 2021. https://doi.org/10.3390/electronics10202509.
33. Q. Delooz and A. Festag, "Network Load Adaptation for Collective Perception in V2X Communications", *Proc. IEEE International Conference on Connected Vehicles and Expo*, pp. 1-6, 4-8 November 2019. https://doi.org/10.1109/ICCVE45908.2019.8964988
34. B. Jung, J. Kim and S. Pack, "Deep Reinforcement Learning-based Context-Aware Redundancy Mitigation for Vehicular Collective Perception Services," *Proc. International Conference on Information Networking (ICOIN)*, pp. 276-279, Korea, Jan 2022. https://doi.org/10.1109/ICOIN53446.2022.9687254
35. H. Masuda, O. E. Marai, M. Tsukada, T. Taleb and H. Esaki, "Feature-based Vehicle Identification Framework for Optimization of Collective Perception Messages in Vehicular Networks," *IEEE Transactions on Vehicular Technology*, pp: 1-11, Oct 2022. https://doi.org/10.1109/TVT.2022.3211852
36. Q. Delooz et al., "Analysis and Evaluation of Information Redundancy Mitigation for V2X Collective Perception," *IEEE Access*, vol. 10, pp. 47076-47093, April 2022. https://doi.org/10.1109/ACCESS.2022.3170029
37. G. Thandavarayan, M. Sepulcre and J. Gozalvez, "Cooperative Perception for Connected and Automated Vehicles: Evaluation and Impact of Congestion Control", *IEEE Access*, vol. 8, pp. 197665-197683, October 2020. https://doi.org/10.1109/ACCESS.2020.3035119
38. ETSI ITS, "Intelligent Transport Systems; Vehicular Communications; Basic Set of Applications; Part 2: Specification of Cooperative Awareness Basic Service", EN 302 637-2 V1.4.1, January 2019.
39. ETSI ITS, "Intelligent Transport Systems (ITS); Decentralized Congestion Control Mechanisms for Intelligent Transport Systems operating in the 5 GHz range; Access layer part", ETSI TS 102 687 V1.2.1, April 2018.
40. ETSI ITS, "Intelligent Transport Systems (ITS); Facilities Layer; Communication Congestion Control", ETSI TS 103 141 V2.1.1, Nov 2021.
41. 3rd Generation Partnership Project (3GPP), Technical Specification Group Radio Access Network; "Study on LTE-based V2X Services,". TR 36.885 V14.0.0, 2016.
42. Caltrans Performance Measurement System (PeMS), California Department of Transportation data-set, U.S.State, http://pems.dot.ca.gov/ (Accessed: Feb. 2021), Online.
43. ETSI ITS, "Intelligent Transport Systems (ITS); Multi-Channel Operation Study; Release 2", ETSI TR 103 439 V2.1.1, Oct. 2021.
44. CAR 2 CAR Communication Consortium, "Survey on ITS-G5 CAM Statistics", Technical Report TR2052, V1.0.1, December 2018.
45. M. Sepulcre, J. Mira, G. Thandavarayan and J. Gozalvez, "Is Packet Dropping a Suitable Congestion Control Mechanism for Vehicular Networks?", *Proc. IEEE 91st Vehicular Technology Conference (VTC2020-Spring),* Antwerp, Belgium, pp. 1-5, 25-28 May 2020. https://doi.org/10.1109/VTC2020-Spring48590.2020.9128822