

Platoon-Local Dynamic Map: Micro cloud support for platooning cooperative perception

Original

Platoon-Local Dynamic Map: Micro cloud support for platooning cooperative perception / RISMA CARLETTI, CARLOS MATEO; Casetti, Claudio; Härrri, Jérôme; Risso, Fulvio. - (2023), pp. 405-410. (VN4RRSR'23: International Workshop on Vehicular Networks for Risk Reduction and Safety Related Systems Montreal, QC (CAN) 21-23 June 2023) [10.1109/WiMob58348.2023.10187883].

Availability:

This version is available at: 11583/2982322 since: 2023-09-28T08:39:28Z

Publisher:

IEEE

Published

DOI:10.1109/WiMob58348.2023.10187883

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2023 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

Platoon-Local Dynamic Map: Micro cloud support for platooning cooperative perception

Carlos Mateo Risma Carletti^{*†}, Claudio Casetti^{*}, Jérôme Härrri[†], Fulvio Risso^{*}

^{*}Politecnico di Torino - Department of Control and Computer Engineering, Turin, Italy

[†]EURECOM - Communication Systems Department, Sophia-Antipolis, France

E-mail: {carlos.rismacarletti, claudio.casetti, fulvio.risso}@polito.it
{risma, jerome.haerri}@eurecom.fr

Abstract—Modern vehicles are equipped with onboard perception systems that generate rich context data, which can be shared with nearby vehicles as the number of vehicles with communication capabilities grows. Platooning is a popular vehicular application for autonomous driving on which the Platoon Leader (PL) manages all maneuvers using context information from Vehicle-to-Vehicle (V2V) messages. However, redundant context information from nearby vehicles in the platoon can increase computational costs for the PL. To solve this issue, vehicular micro-clouds can be formed to enable collective data processing and aggregation, thus reducing the PL's perception workload. The proposed solution, called Platoon Local Dynamic Map (P-LDM), creates a single database of context information, distributing the data aggregation load among all members of the platoon. Simulation results evaluate the effectiveness of the proposed solution and compare it to typical Cooperative Perception mechanisms.

Index Terms—V2X, Autonomous Vehicles, Platooning, Vehicular Micro Cloud, Local Dynamic Map, LDM, ITS, Vehicular Networks, 5G

I. INTRODUCTION

The exchange of data between vehicles is one of the key enablers for the realization of automated driving in the near future. Vehicle-to-Everything (V2X) communications allow for this data exchange under telecommunication standards such as IEEE 802.11p/bd and 3GPP C-V2X, with its latest version referred to as NR-V2X introduced in 3GPP Release 16.

As envisioned by the European Telecommunication Standards Institute (ETSI), data is shared among vehicles through the periodic broadcast of standardized messages such as Cooperative Awareness Messages (CAMs), containing geographical and dynamic information about the vehicle sending the message. In recent years, ETSI has been working on the standardization of a service aimed at sharing data about objects detected by vehicles' perception systems through the exchange of the so-called Collective Perception Messages (CPMs).

These messages are intended to overcome onboard sensors' need for line-of-sight as well as to increase the vehicle perception range with the information coming from neighboring vehicles. However, the distributed nature of CPMs can result in lots of redundant information being exchanged, hence incurring an unnecessary channel load if the inclusion of information is not controlled [1]. Moreover, the computational

power needed to match all locally sensed objects with the data coming from other vehicles increases the delay in the whole process, translating into stored perceptions being outdated [2].

One of the most popular ITS applications studied by the research community is platooning, which is one of the cooperative Automated Driving Systems (ADS) applications to be sooner deployed on highways around the world [3].

Platoons are, in essence, clusters of vehicles that organize each other to travel close behind one another in order to save fuel, reduce traffic and increase road safety. Vehicles belonging to a platoon perform all platooning control and maneuvers with the aid of context information available through V2V messages exchanged by all members of the platoon.

A problem arises for vehicles in a platoon that, due to their close proximity, gather and share highly redundant data that needs to be processed to match with the local one by each platoon member. Although this enhances the view of the road for each member, the Platoon Leader (PL) needs a complete view of the surroundings to manage platoon control. The shared data from Platoon Members (PMs) needs to be further processed by the PL, resulting in an elevated computational cost added to the already computationally demanding platoon management. To address this issue, a vehicular micro-cloud can be employed, which shares the computing resources of the PMs. This enables them to cooperatively process the redundant data, resulting in a complete vision of the road while reducing the overall computational cost. Furthermore, this approach improves the efficiency of platooning and enables the PL to make more informed control decisions.

The solution presented in this paper is a micro-cloud service creating a single aggregated dataset of context information relevant to a platoon, through the exchange of V2X messages, where all PMs share their data with the PL, which, in turn, assigns PMs the task to aggregate and merge a subset of the platoon database. With this paradigm, the collective perception of vehicles inside the platoon is enhanced, thus improving the safety of the platooning maneuvers while distributing the computational load.

The remainder of the paper is organized as follows: in Section II, a description of related works studying the development of cooperative perception schemes is followed by a detailed description of the presented service and architecture

in Section III. Section IV describes the implementation of the service on a custom simulation framework and outlines the obtained results showing the benefits of our solution. Finally, Section V concludes the paper.

II. RELATED WORKS

The ability to perceive and react to changes in the environment is crucial for autonomous vehicles. The concept of Local Dynamic Map (LDMs) has emerged as a powerful facility to represent the dynamic world around a vehicle and has been standardized by ETSI [4]. The latter defines the interfaces for context data providers (e.g., other facilities) and consumers (e.g., vehicular applications) as well as the format for such data to be stored and retrieved. In addition to storing data provided by sensors, ETSI envisions that the LDM is populated by data coming from the Cooperative Awareness Service [5] and Collective Perception Service [6], currently under standardization.

Many works have explored the challenges of exchanging sensor data through CPMs, such as the elevated amount of redundant data being exchanged when no mitigation control mechanism is applied. Indeed, ETSI defines redundancy mitigation rules in [6], similar to the ones defined for CAM generation. However, these rules are not sufficient to limit the redundancy and channel load, as proven by the authors in [1], [7].

Furthermore, the exchange of CPMs with redundant information not only represents a problem in terms of channel load but in the age of the information being stored in the LDM due to elevated computational load as well [2]. As vehicles need to match the objects from received CPMs with the already stored ones from onboard sensors and previous CPMs, the repeated reception of data about the same objects (from different sources) overloads the matching process, introducing high end-to-end delay, rendering the additional information useless.

Several works have addressed the use of the infrastructure to collect and aggregate the different perceptions of vehicles into a single LDM, through the high computational power available at edge servers, as outlined in [8]. Although edge servers can aggregate and merge perceived objects faster than vehicles, the number of received objects is much higher compared to the ones received in a distributed approach. Furthermore, vehicular applications such as platooning operate mostly in rural areas where infrastructure deployment is scarce and network coverage is not ensured at all times.

The authors in [9] introduced a new concept leveraging idle resources present in vehicles belonging to platoons to aid edge servers when needed, proposing to offload vehicle tasks to platoon members. In [10], the authors proposed a federated learning mechanism leveraging the steady network conditions within a platoon despite the elevated absolute speed of worker nodes.

Very few works have investigated the effects of legacy non-connected vehicles in the execution of platooning maneuvers, as outlined in [11], where the authors explore the safety

and stability benefits that cooperative perception provides for platooning. Even though cooperative perception is proven to be an asset for platooning, in the considered scenario all PMs send their raw sensed information to the PL directly. This not only represents a challenge in terms of channel load, especially considering V2V communications, but also in terms of the additional computational overhead brought upon the Platoon Leader, which, given the redundancy of the information, might not be proportional to the perception enhancement. In this work, instead, we propose a lightweight service working with object-level perception sharing, which offloads and distributes the PL's perception tasks needed for the execution and planning of platooning maneuvers, to all Platoon Members.

III. SERVICE DESCRIPTION

A. System Architecture

In the considered architecture all vehicles belonging to a platoon are equipped, without loss of generalities, with NR-V2X Sidelink interfaces for the exchange of V2V messages (e.g., CAMs and CPMs). All NR-V2X-enabled vehicles, referred to in this paper as Connected Vehicles (CVs), are equipped with a GNSS receiver as well as a LiDAR sensor and a camera. The PL manages the platoon by controlling the speed and inter-vehicle distance of all PMs. Additionally, the PL is in charge of all platooning maneuvers such as those needed to let vehicles join or leave the platoon, controlling them according to the context information available to it. For the correct execution of these maneuvers, accurate knowledge of platoon surroundings is paramount, especially when there are legacy non-connected vehicles present. Indeed, the awareness of non-connected vehicles relies entirely on the CVs' onboard sensors being able to detect them. In our use case, all CVs' Perceived Objects (POs) represent the non-connected vehicles currently being detected by their sensors.

All the information gathered by the sensors of a CV is merged to generate an object list that is then used to update the Vehicle's Local Dynamic Map (V-LDM). The V-LDM contains a list of objects and their path history periodically updated with data coming from onboard sensors together with data received from other CVs through the reception of CAMs and CPMs. The information stored in the V-LDM is used by all CVs for the generation of CPMs, following the specifications outlined in [6] for perceived object inclusion.

In our proposed system, the PL runs the P-LDM service, with the aim of aggregating the context information perceived by all PMs, as depicted in Figure 1. In order to not overload the PL with the processing load necessary for aggregating all context data of Platoon Members, in our approach the aggregation process is distributed along the entire platoon, assigning only a subset of all Platoon Perceived Objects (PPOs) to each PM. One of the main aspects of our solution is the synchronization of the context data, reducing the additional processing load needed to aggregate perceptions from multiple vehicles. Thus, each PM runs a local copy of the P-LDM enabling them to be aware of the platoon context

as well, necessary for the correct association between local and platoon-wide perceptions. Additionally, with the aim of further enhancing platoon awareness, for the case in which messages from a CV outside the platoon are not reaching all PMs, a PM is assigned to relay the messages of such CV for all PMs that *subscribed* to it, i.e., PMs that wish to receive the messages of that CV. The P-LDM hence contains information about all PPOs, i.e., objects currently being perceived by one or more PMs, as well as information about all CVs in the platoon vicinity, i.e., CVs from which messages (e.g., CAMs) are currently being received by one or more PMs. Finally, the P-LDM contains a *PM state map* with information of each PM, including which PPO they are currently detecting, which PPO they are assigned with, and their available CPU.

In addition to CAMs and CPMs, the P-LDM service relies on two types of messages, namely, *Platoon Leader Update* (PLU) and *Platoon Member Update* (PMU). The PLU is sent periodically by the PL, carrying the following information containers:

- **Platoon Leader Container:** containing geographical and dynamic information of the PL, optionally extended with platoon-related information (e.g., platoon speed and inter-vehicle gap).
- **Platoon Members:** list of all PMs' ITS Station Identifier (ITS-S ID) and Platoon ID.
- **Connected Vehicles:** list of each CV not belonging to the platoon, containing their ITS-S ID together with the PM assigned with it and the PMs subscribed to it.
- **Platoon Perceived Objects:** a list containing geographical and dynamic information about all the PPOs present in the P-LDM, specifying the PPOs' Platoon ID and the PM assigned with each of the PPOs.

The PMU, instead, is sent by each PM after the reception of a PLU carrying the following information containers:

- **Assigned Perceived Objects:** a list containing geographical and dynamic information about all the PPOs that have been assigned to this PM.
- **New Perceived Objects:** a list containing information about all the Perceived Objects 'seen' by this PM (i.e. appearing in its V-LDM) but not yet included in the P-LDM.
- **New Connected Vehicles:** list of ITS-S IDs of the CVs that are 'seen' by this PM (i.e. appearing in its V-LDM) but not yet included in the P-LDM.
- **Connected Vehicles Subscriptions:** list of ITS-S IDs of the CVs that are found in the P-LDM but not yet seen by this PM (i.e. not messages received from that CV), for which a subscription is requested.
- **PM State:** containing the available CPU of the PM that is dedicated to the P-LDM, and the list of PPO IDs of all the PPOs this PM is currently detecting (assigned to this PM or not).

B. P-LDM within ETSI stack

The presented service is designed to work as an extension of ETSI's LDM facility, working together with facilities such as

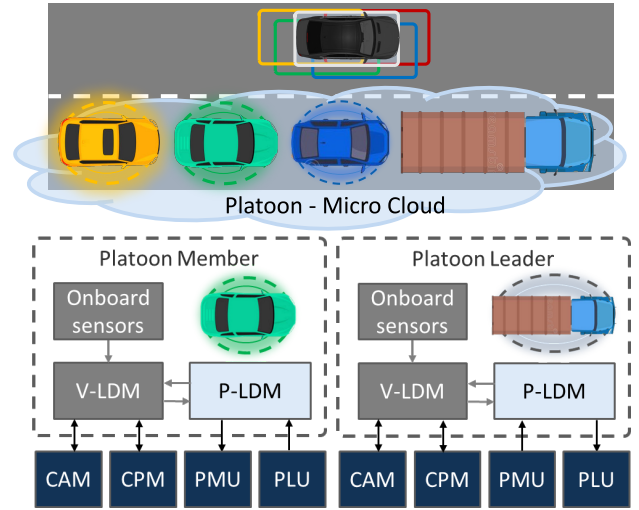


Fig. 1. P-LDM architecture and perception aggregation scheme.

the Cooperative Awareness and Collective Perception services. Indeed, the P-LDM relies on these services for the main purpose of taking full advantage of the available information shared by all vehicles in ETSI's vision of ITS. As already hinted in the previous section, the P-LDM service uses the exchanged information between all connected vehicles through CAMs and CPMs.

For the management of the information of all the CVs outside the platoon, the PL stores all the incoming information found in their CAMs, and whose ITS-S ID is used as the Platoon ID. In this way, all platoon management messages refer to CVs through their IDs. When a PM receives a CAM from a CV assigned to it by the PL, and there is at least one PM subscribed to receive information on that CV, then the PM in turn re-broadcasts the received CAM. On the other hand, since relayed CAMs can translate into duplicates being received by some PMs, every PM checks the timestamp of all received CAMs before storing its information into its own local P-LDM. Furthermore, as soon as a PM receives a duplicate CAM with an ID to which it was previously subscribed, said subscription is canceled since it is now within the communication range of such CV.

On the other hand, for the management of PPOs, each PM updates the PL with the context information of its assigned PPOs, for which not only the sensor information of that PM is used, but the sensor information of other PMs as well, through the exchange of CPMs. The information from both local and remote perceptions is exploited for an enhanced perception fusion of the assigned PPO, both in terms of perception confidence and perception age as well. All received CPM perceptions of a PPO are then stored by the PM assigned with it, in its local copy of the P-LDM for which a *path history* is created for each assigned PPO, containing the last 10 received perceptions. As it is the goal of the presented service, in order to reduce the computational overhead and latency introduced by the matching and fusion of objects from inbound CPMs,

PMs only process CPM objects with an ID assigned to it, leveraging the ID synchronization described in Section III-A. Moreover, as the ID synchronization is only ensured for CPMs from other PMs, all CPMs coming from CVs outside the platoon, are discarded. This policy is further justified given the goal of the service to provide a continuous vision of the environment, which is only possible with information from vehicles continuously being in communication range, i.e., Platoon Members.

Finally, each PM aggregates all the perceptions of an assigned PPO before sending a PMU, computing a weighted average of all geographical positions received within a 100ms window, considered in our case as the sensing period of the on-board sensing module. Perceptions are jointly weighted according to their age and confidence to compute the PPO position at time t according to the following expression:

$$PPO_t = \frac{\sum_{i \in T} \frac{C_i}{A_i} \cdot P'_i}{\sum_{i \in T} \frac{C_i}{A_i}} \quad (1)$$

where C_i is the perception confidence, A_i the perception age and P'_i the estimated position at time t using a Constant Velocity model.

C. P-LDM management

As mentioned in Section III-A, for the management of the P-LDM service, the PL periodically sends a PLU to all PMs, every 100 ms. Upon reception of a PLU, each PM matches all locally Perceived Objects with the ones that might be already stored in the P-LDM, in this way synchronizing all PPO IDs to the local ones found in their V-LDM. The goal of this ID synchronization step is to simplify the matching of objects found in the following PLUs, as well as incoming CPMs from other PMs. Indeed, CPMs specify an object ID arbitrarily generated by the forwarding vehicle when it was first stored in its V-LDM, usually resulting in the same object being referenced by different IDs, on CPMs from different vehicles. Thus, by synchronizing the IDs of the V-LDM with the ones in the P-LDM, all objects in CPMs will be referenced by the same ID as the one in the P-LDM, substantially simplifying perception matching. In order to perform the ID synchronization, each object of the V-LDM and P-LDM is translated into a 2D box defined by 4 corners and a central point. Each object of the V-LDM is compared to each object of the P-LDM computing the Intersection-over-Union (IoU) and the distance between the central points of each object. Considering the possible time difference between the detection of two compared objects, the 'oldest' object position is recomputed, considering its speed and heading, to the predicted position at the time of the 'newer' object. For an IoU higher than 0.7 or a central point distance lower than 2 m, the two objects are considered to be the same, and the ID of the V-LDM object is synchronized. The reason behind the usage of the central point in addition to the IoU as a threshold to match the objects is due to the shape of vehicles, often being long rectangles in their 2D representation, thus resulting in low IoU values for small lateral misalignments

resulting in a high number of false negatives during matching. It is worth mentioning that the additional object-matching process needed for ID synchronization is performed only once for each new PPO in the P-LDM, undertaking the same computational overhead as matching an object from a received CPM, which would instead be needed for all PPOs when not leveraging the P-LDM.

Once all the new PPOs received by a given PM are matched with its locally Perceived Objects (if they are in perception range), each PM checks which of the PPOs 'in sight' have been assigned to it by the PL. For all locally Perceived Objects stored in the PM's V-LDM that were not matched to any of the PPOs in the P-LDM, the PM shares their information with the PL to be appended to the P-LDM, with a new Platoon ID generated excluding currently used ones.

For what regards CVs, each PM reads all the ITS-S IDs received in the PLU, checking if they are all stored in its V-LDM, i.e., checking if CAMs were received from all these ITS-S. In case of an ITS-S ID missing in the PM's V-LDM, the PM subscribes to receive CAMs from that ITS-S, relayed by the PM assigned with that CV. On the other hand, in the case of ITS-S IDs present in the PM's V-LDM but missing in the PLU, the PM will notify the PL about the awareness of those CVs. Finally, for all CVs present in both P-LDM and V-LDM, each PM checks which ones have been assigned to them.

After the PLU reception procedure is finished, each PM returns a PMU (described in Section III-A, with the updated information.

On each PMU reception, the PL updates the P-LDM database with the updated perceptions extracted from the *Assigned Perceived Objects* container and checks for any new PPOs in the *New Perceived Objects* to be added to the P-LDM. Before storing a new PPO in the P-LDM, the PL checks for duplicates that might be found in other PMUs received during the same update period, for which it selects the one with the highest perception confidence to be stored in the P-LDM.

Before starting a new update period, the PL begins the *assignment selection procedure* in which it will update, if necessary, the PPOs assignments. The *assignment selection procedure* is a two-step process aiming to distribute the computational load on all PMs as well as to increase the *perception duration* of each PPO inside the P-LDM, i.e., increasing the time during which a PPO is continuously 'seen' by the platoon. The reason for this process being performed in two steps is due to the fact that a higher priority is given to the distribution of the computational load along the platoon, with respect to the PPOs perception duration. Thus, a preliminary assignment selection is computed looking to obtain the longest perception duration from which a definitive assignment selection is computed prioritizing the best distribution of the computational load in terms of the number of PPOs assigned to each PM.

In the first step of the assignment selection depicted in Algorithm 1, for every PPO, the PL checks which of the PMs that are currently perceiving it, is going to be the closest one to the predicted position after one second, considering a constant

velocity model. It is worth mentioning that since all PMs travel at the same speed, the PPO's predicted position relative to the platoon is computed considering all PMs static and the PPO traveling at its relative speed. When the preliminary selection is finished, the PL then distributes the assignments depending on the current load in each PM. Algorithm 2 depicts the mechanism used, where for each PM a new assignment is given depending on its current number of assigned PPOs and available CPU, extracted from the PM State map. More in detail, starting from the most loaded PM in terms of the number of assigned PPOs, for each PM_i , the algorithm checks if there is another less loaded PM_k , currently perceiving one of the PPOs assigned to PM_i , and assign that PPO to PM_k only if it has enough CPU available. In order to compute if a PM has enough CPU available, the average CPU load per PPO is computed in line 14 to then check if an additional assigned PPO can be sustained by PM_k . Finally, a switch of PPO assignments is considered only if the difference in the number of assigned PPOs between PM_i and PM_k is higher than one, evaluated in line 8.

Algorithm 1 Perception Duration improvement selection

```

1: t = current time
2: Get PPOs list
3: for each  $PPO_i$  in P-LDM do
4:   Get all PMs perceiving  $PPO_i$ 
5:   for each  $PM_j$  perceiving  $PPO_i$  do
6:     Compute predicted  $PM_j$  distance from PPO at  $\tau = t + 1s$ 
7:   end for
8:   Pre-assign PPO to closest PM at time  $\tau$  in PMs state list
9: end for

```

Algorithm 2 Computational load distribution selection

```

1: Get PMS = PMs State map  $\{|PMS| = PMS \text{ length}\}$ 
2: Read all PPOs assignments
3: Order PMs states from most to least loaded
4: for  $i=0$  to  $|PMS|$  do
5:   Read Assigned PPOs ( $PM_i$  A-PPOs)
6:   for  $k=|PMS|$  downto  $i$  do
7:     Read Assigned PPOs ( $PM_k$  A-PPOs)
8:     if  $PM_i$  A-PPOs <  $PM_k$  A-PPOs + 1 then
9:       Break
10:    end if
11:    for each  $PPO_j$  assigned to  $PM_i$  do
12:      if  $PM_k$  is perceiving  $PPO_j$  then
13:        Read Available CPU (A-CPU), Total CPU (T-CPU)
14:         $PPOload = (T-CPU - A-CPU)/A-PPOs$ 
15:        if  $PM_k$  A-CPU >  $PPOload$  then
16:          Assign  $PPO_j$  to  $PM_k$ 
17:          Update PMS list
18:        end if
19:      end if
20:    end for
21:  end for
22: end for
23: Update P-LDM

```

IV. PERFORMANCE EVALUATION

A. Simulation Framework

In order to assess the performance of the P-LDM service, an implementation has been developed on an open-source vehicular network simulation framework coupling SUMO (Simulation of Urban MObility) and ns-3. This framework, called ms-van3t, has been designed for the validation of vehicular applications offering different access technologies possibilities, such as NR-V2X in our use-case, together with an implementation of a full ETSI ITS-G5 stack for different services, such as the Cooperative Awareness Service.




Fig. 2. ms-van3t simulated sensor implementation.

For this work, an implementation of the Collective Perception Service (CPS) and Local Dynamic Map (LDM) facilities has been developed and incorporated into the framework, enabling the deployment of both services for all simulated vehicles. Similarly to other available solutions for the analysis of collective perception [12] [13], each connected vehicle in the simulation is equipped with a simulated radar sensor (placed on top of the vehicle) detecting objects in line of sight, as depicted in Figure 2, within a configurable range set to 50 meters by default. For every detection period of 100 ms, the sensor module queries SUMO about all vehicles within sensor range, discards blocked vehicles and introduces gaussian noise with a gain proportional to the object distance from the sensor. Lastly, a box-like representation is stored in the V-LDM, with additional measurements such as speed and heading, for which the gaussian noise is added according to the following expressions:

$$d_s = d_{gt} + \delta \cdot \frac{d_{gt}}{r_s} \quad (2)$$

$$[w_s, l_s] = [w_{gt}, l_{gt}] \cdot \left(1 + \delta \cdot \frac{d_{gt}}{10 \cdot r_s}\right) \quad (3)$$

$$h_s = h_{gt} + \theta \cdot \frac{d_{gt}}{r_s} \quad (4)$$

$$v_s = v_{gt} + \nu \cdot \frac{d_{gt}}{r_s} \quad (5)$$

where d_s , $[w_s, l_s]$, h_s and v_s represent, respectively, the stored distance, width, length, heading, and speed, while d_{gt} , $[w_{gt}, l_{gt}]$, h_{gt} and v_{gt} are the respective ground truth values of the above quantities as provided by SUMO's Traffic Control Interface (TraCI). For each of the ground truth measurements the added Gaussian noise, denoted by δ , θ and ν , is generated with zero mean and a standard deviation of 1.0 m, 0.01 rad and 0.5 m/s respectively [13]. Lastly, d_{gt}/r_s represents the noise gain, normalized by the sensor range r_s .

The LDM keeps track of all neighbouring connected vehicles through the reception of CAMs from which the information is stored by the implemented LDM API. It is worth mentioning that all sensor module's detected objects go through a matching process with the already stored elements of the LDM, in order to keep track of objects already detected and to avoid perceiving a connected vehicle as a detected object. The LDM API implementation enables the CPS to query the available information about detected objects to be used for the creation of CPMs. Furthermore, the implemented CPS follows the generation rules outlined in [6] for redundancy mitigation.

Finally for both P-LDM and V-LDM implementations, in order to keep only updated information and avoid a higher object matching complexity, with a periodicity of one second, the database deletes all perceived objects detected more than one second ago.

B. Simulation Setup

In accordance with the investigated use-case for the P-LDM (i.e., platooning), all the simulations have been performed on a highway scenario, for which a custom 3-lane highway loop has been designed in SUMO with a total circumference of 3 km. The platoon formation protocol is out of the scope of this work, hence all simulation measurements have been performed on an already-formed platoon. Platoon sizes of 10 and 20 PMs have been considered to evaluate its impact on the service performance. Beside platoon members, a variable percentage of simulated vehicles has been considered to be equipped with NR-V2X interfaces according to a pre-configured market penetration rate. Indeed, the penetration rate has an impact on the density of vehicles without V2X capabilities which in our scenario translates to an impact on the number of potential objects to be detected and shared by connected vehicles. As our service has been envisioned to be used by a platooning application deployed in the near future, i.e., at low market penetration rates, we consider values of 10% and 50%. All communications are configured to take place using NR-V2X mode 2 on a 20 MHz channel at a central frequency of 5.9 GHz, with a subcarrier spacing of 30 kHz, using the Modulation and Coding Scheme (MCS) index 14, and a transmission power of 23 dBm.

With the goal of evaluating the benefits of P-LDM service, its KPIs were compared with two different baselines:

- Full Cooperative Perception, representing a completely distributed scenario in which the Platoon Leader manages the platoon according to the information received from all neighboring vehicles, including CVs not belonging to the

platoon, and stored in its own V-LDM; such information is processed entirely by the PL itself.

- Platoon Cooperative Perception, representing an intermediate scenario between the Full Cooperative Perception baseline and the use of the P-LDM service, in which the PL and PMs only process the information of detected objects, which comes from other PMs only. This baseline implies that PMs only process a portion of all incoming information, stored in their own V-LDM; since it comes from fellow PMs, its context will be more relevant to the platoon, and with a longer availability as well.

The considered KPIs for each scenario are the following:

- Total Number of PPOs, meaning the total number of Platoon Perceived Objects stored in the LDM: the higher the number, the greater the awareness.
- Number of Perceived Objects handled by each PM, representing the number of Perceived Objects for which each PM has to periodically perform object matching and fusion, proportional to their computational load.
- Platoon Perception Duration, as introduced in Section III-C, it represents the time duration for which a given object is continuously perceived by the platoon. This metric translates into the robustness of the aggregated perception, overcoming problems due to individual vehicle perception range and sensor blockage.
- Perception Age, representing the time between the sensor detection of a PPO and the present time.

For all the results showcased, 10 simulations of 600 seconds have been considered with a vehicle density of 30 vehicles per kilometer for each scenario, where measurements for all KPIs have been taken every 100ms. Finally, all the results are shown from the point of view of the PL, since it uses the context data for platoon maneuvers. It is worth mentioning that in the case of the P-LDM scenario, the same data is available for all PMs in addition to the PL.

C. Simulation Results

Regarding the Total Number of Perceived Objects in the LDM, the Cumulative Distribution Function (CDF) is depicted in Figure 3 for each of the scenarios. The obtained results show that when working with Full Cooperative Perception (Full CP), the PL has a higher number of perceptions stored in its LDM for all scenarios, except for the case of a platoon size of 20 and a 10% market penetration rate, for which the P-LDM provides almost identical results as the aggregated platoon perception range increases and gets close to communication range. The higher performance of the Full CP is due to objects perceived by vehicles outside the platoon being included in the LDM, as opposed to only including ones perceived by PMs, which is more evident when considering a higher penetration rate as there are more connected vehicles outside the platoon.

However, in both the Full CP and Platoon CP scenarios, all perceived objects need to be processed by each of the PMs as opposed to the P-LDM where PMs only process the information about perceived objects assigned to them. As can be seen from the CDF of the number of perceived

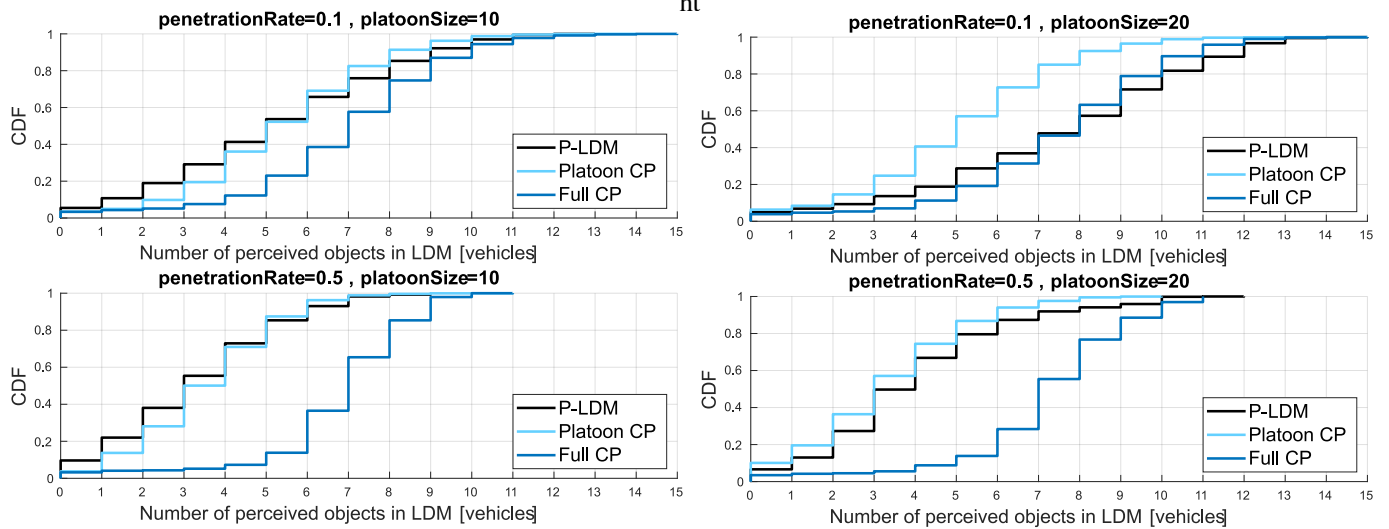


Fig. 3. Cumulative Distribution Function of the Number of Perceived Objects stored in the LDM at a given point in time, for different platoon sizes and market penetration rate.




Fig. 4. Cumulative Distribution Function of the Average Number of Perceived Objects handled by each PM, in the LDM at a given point in time, for different platoon sizes and market penetration rate.

objects handled by each PM, depicted in Figure 4 the P-LDM substantially reduces the computational load on each of the PMs, handling no more than 2 PPOs each for 99% of the time. Additionally, when analyzing the average Platoon Perception Duration depicted in Figure 5 for a platoon size of 10 vehicles and a market penetration rate of 50%, it can be seen that when relying on the P-LDM, perceptions are held on the LDM for 10 times the duration of both Full CP and Platoon CP. When comparing these results with the ones depicted in Figure 3, even though there are more perceptions 'seen' by the PL on the Full CP case, these perceptions are very sporadic due to perceptions received from vehicles outside the platoon, who fail to update the positions of objects after getting out of communication range.




Fig. 5. Cumulative Distribution Function of the Average Platoon Perception Duration of Objects in the LDM, for platoon size of 10 and market penetration rate of 50%.

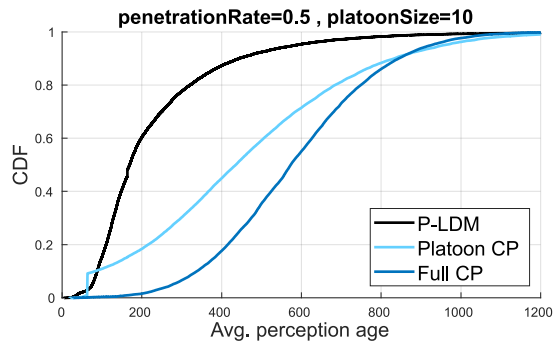


Fig. 6. Cumulative Distribution Function of the Average Age of Perceived Objects in the LDM, for platoon size of 10 and market penetration rate of 50%.

Finally, when analyzing the age of the perceptions stored in the LDM, again the P-LDM outperforms the two baselines, as can be seen from the CDF of the average value for all perceptions at a given point in time depicted in Figure 4. The results show that the P-LDM perception on average are not older than 400 ms on 90% of the time, compared to 800 ms for the other 2 baselines. It is worth mentioning that the depicted values represent the average over all perceptions at a given point in time, thus including perceptions that have stopped being updated, e.g., have gone outside the platoon perception range, and stored until the database is cleaned as described in Section IV-A.

V. CONCLUSIONS AND FUTURE WORK

This paper presented a novel approach to the concept of cooperative perception in a platooning use-case offering an enhanced vision of the platoon surroundings for the better planning and execution of platoon maneuvers while distributing the computational overhead due to redundant context data, for all members of the platoon. We have showcased how the service succeeds in providing a more stable, more up-to-date view of the road, compared to the regular cooperative perception scheme, while turning the existence of redundant context data into an advantage as it is aggregated without overloading all platoon members. Indeed, each platoon member only performs the aggregation of a subset of all the context data, instead of the entirety of it in the regular cooperative perception scheme.

The focus of future work will be twofold, analyzing the benefits of an enhanced perception for the actual execution of platooning maneuvers and a more in-depth evaluation of the computational load of the service on real hardware.

REFERENCES

- [1] 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, 2020.
- [2] G. Volk, Q. Delooz, F. A. Schiegg, A. Von Bernuth, A. Festag, and O. Bringmann, "Towards realistic evaluation of collective perception for connected and automated driving," in *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, 2021, pp. 1049–1056.

- [3] ITS Deployment Evaluation Executive briefing. [Online]. Available: <https://www.itskrs.its.dot.gov/briefings/executive-briefing/truck-platooning>
- [4] ETSI, "ETSI EN 302 895 V1.1.1 (2014-09) - Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Local Dynamic Map (LDM)," European Telecommunications Standards Institute, Standard ETSI EN 302 895 V1.1.1, 2014.
- [5] —, "ETSI EN 302 637-2 V1.4.1 (2019-04) - Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Part 2: Specification of Cooperative Awareness Basic Service," European Telecommunications Standards Institute, Standard ETSI EN 302 637-2 V1.4.1, 2019.
- [6] —, "ETSI TR 103 562 V2.1.1 (2019-12) - Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Analysis of the Collective Perception Service (CPS); Release 2," European Telecommunications Standards Institute, Standard ETSI TR 103 562 V2.1.1, 2019.
- [7] G. Thandavarayan, M. Sepulcre, and J. Gozalvez, "Redundancy mitigation in cooperative perception for connected and automated vehicles," in *2020 IEEE 91st Vehicular Technology Conference (VTC2020-Spring)*, 2020, pp. 1–5.
- [8] Z. Zhang, S. Wang, Y. Hong, L. Zhou, and Q. Hao, "Distributed dynamic map fusion via federated learning for intelligent networked vehicles," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 953–959.
- [9] M. Nasimi, M. A. Habibi, and H. D. Schotten, "Platoon-assisted vehicular cloud in vanet: Vision and challenges," 2020. [Online]. Available: <https://arxiv.org/abs/2008.10928>
- [10] X. Ma, J. Zhao, and Y. Gong, "Joint scheduling and resource allocation for efficiency-oriented distributed learning over vehicle platooning networks," *IEEE Transactions on Vehicular Technology*, vol. 70, no. 10, pp. 10894–10908, 2021.
- [11] R. Xu, Y. Guo, X. Han, X. Xia, H. Xiang, and J. Ma, "Openca:an open cooperative driving automation framework integrated with co-simulation," 2021. [Online]. Available: <https://arxiv.org/abs/2107.06260>
- [12] H.-j. Gunther, O. Trauer, and L. Wolf, "The potential of collective perception in vehicular ad-hoc networks," in *ITST 2015*, 2015, pp. 1–5.
- [13] C. Allig and G. Wanielik, "Extending the vehicular network simulator artery in order to generate synthetic data for collective perception," *Advances in Radio Science*, vol. 17, pp. 189–196, 09 2019.