

Characterization of In-Vehicle Network Sensor Data Traffic in Autonomous Vehicles

L. Lusvarghi, B. Coll-Perales, J. Gozalvez, K. Aghababaiyan, M. Almela, M. Sepulcre
Uwicare laboratory, Universidad Miguel Hernandez de Elche, Elche (Alicante), Spain.
{llusvarghi, bcoll, j.gozalvez, kaghababaiyan, mario.almela, msepulcre}@umh.es

Abstract— The softwarization of vehicles and the evolution towards autonomous driving is imposing increasing flexibility and reliability demands to future in-vehicle networks (IVN). Research on 6G advocates for a seamless integration of vehicles with cellular networks for a deep edge-edge-cloud continuum that facilitates the opportunistic offloading of in-vehicle processing to the edge/cloud. Realizing this vision requires a seamless connection of IVNs with the cellular networks, which can be facilitated through the gradual adoption of in-vehicle wireless subnetworks. These subnetworks can support increasing dependable and deterministic service levels using predictive schedulers that can anticipate in-vehicle traffic flows and patterns to schedule communication resources and computing workloads. This requires an accurate characterization of IVN traffic, and this study progresses the state-of-the-art with a first characterization of IVN traffic in autonomous vehicles. The study characterizes the data captured by a full suite of sensors as well as the processed data for supporting automated driving. We also derive spatial and time correlations between the IVN data that can serve to anticipate network demands and predict traffic flows for the support of deterministic IVN services.

Keywords—*In-vehicle network, IVN, traffic characterization, subnetworks, 6G, CARLA, Autoware, Zenoh, connected and automated mobility, autonomous driving.*

I. INTRODUCTION

The European Smart Networks and Services (SNS) vision delves into the idea that future 6G networks should go beyond pure communication systems and sustainably integrate computing and data services within a network of networks ecosystem. This ecosystem includes subnetworks located at the edge of the 6G ‘network of networks’. Subnetworks are short-range and low power radio cells providing localized and cost-effective solutions to services with (potentially) extreme requirements and distributed processing capabilities for autonomous local data management. Subnetworks benefit from a seamless integration with a broader cellular network for a deep edge-edge-cloud continuum that facilitates seamless connection and the opportunistic offloading of in-vehicle processing to the edge/cloud. The edge/cloud can then act as a virtual Electronic Control Unit (ECU) that elastically extends the computing and processing capabilities of the vehicle using the 6G-based edge and cloud resources.

The Horizon Europe 6G-SHINE (6G SHort range extreme communication IN Entities) [1] and 6G flagship Hexa-X-II [2] projects envision the deployment of subnetworks in vehicles. This requires 6G-native in-vehicle wireless subnetworks to provide dependable service levels similar to those reached with cables. This must be done considering the evolution

towards zone-based Electrical/Electronic (E/E) architectures [3] and software-defined vehicles that increase the flexibility to dynamically schedule computing workloads and communication resources. In-Vehicle Networks (IVN) also need to efficiently support multiple traffic flows with varying (and increasingly demanding) requirements as we progress towards Automated Driving (AD), and this is still a significant challenge for time-safety critical systems that require deterministic service levels. Sustaining such levels can be facilitated with solutions capable of predicting traffic flows and patterns, and anticipate scheduling IVN decisions [4]. This possibility strongly depends on an accurate characterization of IVN traffic, and the open availability of in-vehicle datasets that could be used to characterize the IVN traffic is very rare for confidentiality reasons, in particular for autonomous vehicles.

An in-vehicle traffic dataset is released in [5], but the dataset was logged in traditional Controller Area Network (CAN) buses, while IVNs are transitioning towards alternatives technologies like Automotive Ethernet that can support higher bandwidth and reliability levels as demanded by AD sensors and functionalities. Recent datasets like Google-Landmarks datasets, A2D2 (Audi) or BDD100k (Berkeley) openly release data captured by sensors mounted in a vehicle, but the number of sensors is generally limited, and the dataset does not capture IVN traffic. The IEEE 802.1DG standard (version D2.4) [6] defines a Time-Sensitive Networking (TSN) profile for automotive in-vehicle ethernet communications. The standard includes an informative annex that provides a qualitative (and not quantitative) description of the types of automotive IVN traffic (cyclic, time-triggered, best effort), the dominant flows in terms of number of messages (monitor or control) or bandwidth consumption (high resolution sensors like cameras, lidars, and radars), and the flows with the highest latency constraints (acoustic sensors @ 50KHz). The study in [7] evaluates mixed traffic flows over Automotive Ethernet using the software (SW) tool RTaW-Pegase and Renault’s Ethernet prototype network. However, the evaluation focuses on 4 traffic classes (audio streams, video streams, command and control, and best-effort) characteristic of current IVN (i.e., not supporting AD) and different average bandwidth and deadline requirements. For example, the study considers audio streams with 128- and 256-byte frames and a 10ms deadline, 43 Kbytes video streams at 60fps and 30fps with 10ms and 30ms deadlines, respectively, and command and control streams with 256- to 1024- byte frames and 10 ms deadline. In [8], the authors analyze the possibility for TSN to support IVN traffic in autonomous vehicles. To this aim, the authors implement a TSN simulator and model an IVN with sensor data generated by 5 radars, 4 around view monitoring cameras, 3 mirrorless cameras, 1 forward camera, and 1 lidar. However, the focus is on testing the possibility to satisfy TSN transmission requirements, and the sensors are configured to continuously generate data with the same characteristics, which limits the capacity to capture the dynamics and patterns that IVN traffic might exhibit in different driving environments.

This work has been partially funded by the European Commission Horizon Europe SNS JU 6G-SHINE (GA 101095738) project, by MICIU/AEI/10.13039/501100011033 and FSE+ (grant RYC2022-036953-I), by the “European Union NextGenerationEU/PRTR” (TED2021-130436B-I00), by Generalitat Valenciana, and UMH’s Vicerrectorado de Investigación. The authors would also like to acknowledge ChenYing Kuo for his support in the development of the CAM platform.

This paper progresses the state-of-the-art by presenting a first analysis of IVN traffic generated by a full suite of AD sensors. To this aim, we introduce a Connected and Automated Mobility (CAM) platform that seamlessly combines CARLA’s realistic 3D modelling of the driving environment and vehicle sensing capabilities, with the Autware AD software stack to control the vehicles’ driving based on the perceived data. CARLA and Autware are integrated with a Zenoh-based bridge that has been extended by the authors to log, process, and characterize the data generated by the full suite of Level 3 (L3) AD sensors in CARLA as well as the processed data generated by Autware’s AD stack inside the vehicle¹. The logged data is first characterized, and we then identify spatial and temporal correlations between sensors logged and processed data. The obtained results pave the way for predicting IVN traffic flows that can be exploited to design future IVN, including predictive IVN wireless subnetworks for deterministic service provisioning.

II. CAM PLATFORM

Realistic IVN traffic traces are collected using an advanced Connected and Automated Mobility (CAM) platform (Fig. 1) that combines high-fidelity 3D representations of the driving environment with realistic sensing and autonomous driving capabilities. The platform integrates the open-source CARLA simulation platform [9] and Autware’s [10] AD software stack using a Robot Operating System (ROS) bridge based on Zenoh [11].

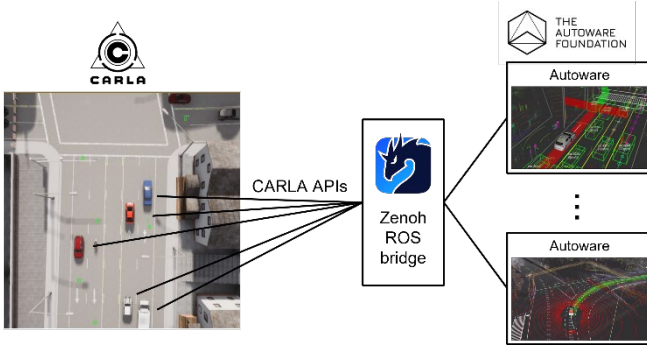


Fig. 1. CAM platform.

CARLA is a 3D simulator built on Unreal Engine 4 (UE4) for AD research that offers a realistic representation and simulation of the driving environment including different types of vehicles (cars, trucks, and motorcycles), pedestrians, buildings, light and weather conditions. CARLA also implements a complete suite of AD sensors (e.g., cameras, radars, lidars, Inertial Measurement Units – IMUs – and Global Navigation Satellite System – GNSS – receivers) that can be flexibly configured and positioned by the user on each simulated vehicle. Autware is an open-source fully functional AD software stack that implements all the perception, localization, planning, and control functionalities required for AD. Autware is built on ROS and can be deployed on a broad range of real-world vehicles and applications.

The platform integrates CARLA and Autware with an extended version of the Zenoh bridge. Zenoh is a scalable low-latency high-throughput routing protocol that allows a

bidirectional connection of multiple Autware instances (one per Connected and Automated Vehicle - CAV) with CARLA. The bridge retrieves from CARLA the raw data captured by the ego-vehicle sensors (e.g., uncompressed camera images, lidar point-clouds, raw GNSS or IMU data) and forwards it to the sensing module of the Autware instance running on the ego-vehicle. Autware uses this information to run the AD modules (e.g., perception and planning) and control the vehicle. The control commands generated by Autware (e.g., steering angle, braking, and acceleration) are sent back through the Zenoh bridge to CARLA where the ego-vehicle’s position and dynamics are updated.

The authors have expanded the integrated CARLA-Zenoh-Autware platform to enhance its sensing capabilities, including the type of supported sensors and the number of sensors that can be deployed per vehicle. This is important to realistically model the data traffic generated within a CAV and be able to accurately characterize the IVN traffic. Our current implementation reproduces the sensor deployment of a Mercedes-Benz S-class car implementing the L3 AD Mercedes Drive Pilot software [12]. The sensor deployment (Fig. 2) includes 1 Lidar, 5 Radars, 5 high-resolution (2560x1440) RGB cameras, 1 IMU, and 1 GNSS receiver. In Fig. 2, R1-R3 and C1-C4 are front-facing sensors, while R4, R5 and C5 face backwards. Table I reports the main sensor specifications in our current implementation. The deployment of this sensor setup has demanding processing and hardware requirements, in particular regarding the AI-based object detection algorithms running in Autware’s perception module. Our CAM simulation platform currently runs on a high-end desktop computer (i9-9980X 18 cores 4.4Ghz, 64 Gb RAM, SSD 960 Gb) with 2 NVIDIA Quadro P4000 GPUs, but it was necessary to implement new SW functionalities in Autware to distribute and balance the GPUs workload.




Fig. 2. L3 CAV sensor deployment.

TABLE I
MAIN SENSOR SPECIFICATIONS.

Sensor Type	Sampling Rate [Hz]	Field of View (FoV)	Range [m]
Lidar	10	Horizontal: 360° Vertical: [-30°, 10°]	100
Radar	20	Horizontal: 30° Vertical: ±30°	100
Camera	20	90°	-
GNSS	20	-	-
IMU	20	-	-

¹ In this paper, we focus on Autware’s raw and processed AD sensor data. However, the implemented CAM platform is capable to log all ROS

messages in Autware, including messages related to the perception, planning, control, and driving modules.

III. IVN TRAFFIC CHARACTERIZATION

IVN traffic generated by AD sensors has been characterized in an urban driving environment with two lanes (one per driving direction) and different vehicular densities. IVN AD traffic traces are collected from a single ego-vehicle that runs Autoware while the other vehicles are controlled by simplified car-following models implemented in the CARLA traffic manager.

Fig. 3 plots the average data rate of raw and processed sensors data considering different sensor types mounted on the ego-vehicle and a medium density scenario. The figure shows the demanding data rate requirements that characterize the transmission of raw lidar (200 kB/s), radar (800 kB/s), and camera (20 MB/s) data. Fig. 3 also shows that there is a significant (larger than 97.7%) data rate reduction for the lidar, camera, and radar sensors when moving from raw to processed sensors data. This is the case because the raw data consists of large point clouds (lidar and radar) and high-resolution images (cameras) that are transmitted without any prior compression, whereas the processed data consists of a list of detected objects for the lidar and radar, and Regions of Interest (RoI) for the cameras. Detected objects and RoIs exhibit a significantly smaller payload and, thus, average data rate as they are compactly represented by 2D points or bounding boxes with specific distance, dimension, and classification attributes. Whether the processing of the raw data is performed within the individual sensors in a distributed fashion or in a centralized high-performance CPU in the E/E in-vehicle architectures has strong implications on the dimensioning and planning of IVNs, since the use of compression techniques to reduce bandwidth requirements might be challenging for time-sensitive safety-critical functions.

Fig. 3 shows a small data rate reduction (-40%) from raw to processed data for the IMU and an increase (+53.5%) for GNSS sensors. IMU raw data undergoes a correction and denoising process that does not alter the raw data input format (like in the lidar, radar, and camera case), but only removes spurious information. In Autoware, raw GNSS data is processed to extract multiple GNSS-based information, e.g., the vehicle position and orientation, and to determine the measurement uncertainty. This additional information augments the processed data and the corresponding average data rate.




Fig. 3. Average data rate of raw and processed sensor data.

The vehicular density can have a significant impact on the number of detected objects and, hence, the data exchanged on the IVN, that should be considered when dimensioning and

designing IVN scheduling techniques. This is visible in Fig. 4, where the message size measured at the output of the Autoware's perception module is reported for low and medium densities. The message generated by the perception module contains the final list of detected objects, and this information is generated following the object detection, the sensor fusion, and the trajectory estimation algorithms. The characteristics of the traffic generated by the perception module is particularly relevant as it feeds the planning module (hence, it is extremely time-sensitive) and influences the ego-vehicle's planned trajectory.




Fig. 4. Size of the messages generated by Autoware's perception module.




Fig. 5. Cross-correlation between the sensors processed data size.

Next, Fig. 5 depicts the Spearman cross-correlation coefficient measured between the processed data size of different AD sensors. The processed data size is a particularly relevant indicator, as it reflects the number of objects detected by each sensor. The Spearman correlation coefficient is a nonparametric rank correlation indicator that measures the monotonicity degree in the relationship between two variables. A value of the Spearman correlation equal to +1 (-1) means that each variable is a perfectly positive (negative) monotone function of the other. On the other hand, a Spearman correlation coefficient equal to zero means that there is no correlation between the examined variables. Fig. 5(a) reports the cross-correlation between the ego-vehicle cameras C1-C5 (deployment in Fig. 2). The front-facing cameras (C1-C4) exhibit high Spearman cross-correlation coefficient values (or cross-correlation coefficient) because they have overlapping Field of Views (FoV). A cross-correlation coefficient close to 1 means that when one front-facing camera (e.g., C1) detects an increasing number of objects, and its processed data size augments, the size of the processed data of the other front-facing cameras (e.g., C2-C4) also increases since they detect a similar number of objects. This trend can be exploited to anticipate the characteristics of the sensors' raw and/or processed data, and schedule transmissions within the IVN accordingly. Predictive scheduling strategies are more challenging when sensors exhibit low cross-correlation coefficients. This is, for example, the case of the rear-facing camera (C5) and the front-facing cameras (C1-C4) in Fig. 5(a), due to their opposite FoVs. Similarly, Fig. 5(b) shows that high cross-correlation values are observed only between radar sensors facing the same direction (R1-R3), i.e., with overlapping FoVs. Fig. 5(c)





Fig. 6. Temporal evolution of the cross-correlation coefficient measured between the processed data size of different radars.

analyzes the cross-correlation measured between the processed data size of different sensor types, in particular cameras and radars. The figure shows that high cross-correlation values also characterize different sensor types facing the same direction, e.g., front-facing cameras (C1-C4) and radars (R1-R3). On the other hand, sensors facing opposite directions (e.g., front-facing radars and rear-facing camera C5), exhibit considerably lower cross-correlation values. These observations reinforce the potential for an IVN scheduler design that exploits IVN traffic characteristics, trends, and correlations among in-vehicle nodes (in this case, AD sensors).

Fig. 6 provides another interesting perspective on the potential information that IVN traffic characterization can provide to design future IVNs. The figure analyses the temporal evolution of the cross-correlation coefficient measured between the processed data size of the front-center radar (R1) and the other radar sensors (R2-R5). The Spearman cross-correlation coefficient reported in Fig. 6 is computed over a sliding window of 5 seconds and a selected time period (from $t = 20$ s to $t = 40$ s) that captures the moment when a vehicle driving in the opposite direction approaches and then passes by the ego-vehicle. From $t = 20$ s to $t = 30$ s, Fig. 6 shows an increasing cross-correlation between the front bumper radars (i.e., between R1 and R2, and between R1 and R3), as the vehicle driving in the opposite direction approaches the ego-vehicle and is detected by its front-facing radars. This leads to a simultaneous increase in the processed data size of R1, R2, and R3. The opposite trend is observed between the front-center radar (R1) and the rear-facing radars (R4-R5), as the latter sensors are not able to detect the approaching vehicle. The cross-correlation between sensors strongly depends on their FoV and the dynamics of the driving environment. At $t = 30$ s, the vehicle driving in the opposite direction is no longer detected by the front-right radar (R3) as its FoV is concentrated on the sidewalk and only partially covers the opposite lane, whereas it remains in the FoV of the front-center (R1) and front-left (R2) radars. The impact of the front-facing radars FoV position and orientation is visible from $t = 30$ s to $t = 35$ s, when the approaching vehicle passes by the ego-vehicle. In this simulation segment, the cross-correlation coefficient measured between radars R1 and R3 decreases, while it remains high for R1 and R2. The results depicted in Fig. 6 clearly highlight existing correlation patterns between data generated by different AD sensors mounted on vehicles. Their characterization (including their



temporal evolution) can be exploited to anticipate IVN traffic flows and predictively schedule IVN transmissions for deterministic service provisioning.

IV. CONCLUSIONS

This study has presented a first IVN traffic characterization for autonomous vehicles using a novel advanced CAM platform that allows the realistic simulation of L3+ sensor deployments and autonomous driving functionalities. The presented results offer relevant insights on the data rate requirements that raw and processed sensor data exhibit, as well as on existing spatial and temporal correlation patterns among traffic flows (e.g., of collocated sensors). Such correlations can be exploited to design future predictive IVN (wireless) schedulers capable of supporting deterministic service levels. Future work will enhance the current IVN traffic characterization study with more diverse driving and context conditions, and will further explore patterns and correlations between messages (i.e., spatial, temporal and delay tolerance) related to the sensors, perception, planning, and control AD driving modules.

REFERENCES

- [1] 6Gshine-D2.2, "Refined definition of scenarios use cases and service requirements for in-X subnetworks", Feb. 2023.
- [2] HexaXII-D1.2, "6G use cases and requirements", Dec. 2023.
- [3] T. Huck *et al.*, "The next step in E/E architectures", White Paper, 2024.
- [4] G. P. Sharma *et al.*, "Toward Deterministic Communications in 6G Networks: State of the Art, Open Challenges and the Way Forward", *IEEE Access*, vol. 11, pp. 106898-106923, Sept. 2023.
- [5] B. Lampe and W. Meng, "can-train-and-test: A Curated CAN Dataset for Automotive Intrusion Detection", arXiv:2308.04972, August 2023.
- [6] IEEE P802.1DG™/D2.4, "Time-Sensitive Networking Profile for Automotive In-Vehicle Ethernet Communications", Feb. 2024.
- [7] J. Migge *et al.*, "Insights on the Performance and Configuration of AVB and TSN in Automotive Ethernet Networks", *Proc. 9th ERTS*, Jan 2018.
- [8] J. Lee and S. Park, "Time-Sensitive Network (TSN) Experiment in Sensor-Based Integrated Environment for Autonomous Driving", *Sensors*, vol. 19, no. 5, Mar. 2019.
- [9] A. Dosovitskiy *et al.*, "CARLA: An Open Urban Driving Simulator," *Proc. 1st CoRL*, vol. 78, Nov. 2017.
- [10] "Autoware - the world's leading open-source software project for autonomous driving". [Online]. Available: <https://github.com/autowarefoundation/autoware>
- [11] "Zenoh – The Zero Overhead, Pub/Sub, Store, Query, and Compute Protocol". [Online]. Available: <https://zenoh.io/>
- [12] "DRIVE PILOT Automated Driving | Mercedes-Benz USA". [Online]. Available: <https://www.mbusa.com/en/owners/manuals/drive-pilot>