SCUOLA di DOTTORATO

Dottorato in Ingegneria dell'Informazione

Direttore della Scuola di Dottorato
Prof. Felice Arena

Collegio dei docenti
Dottorato di Ricerca in Ingegneria dell'Informazione
XXXVIII ciclo

Tommaso ISERNIA (Coordinatore)
Giuseppe ARANITI
Francesco Antonio BUCCAFURRI
Claudia CAMPOLO
Riccardo CAROTENUTO
Giuseppe COPPOLA
Mariantonia COTRONEI
Lorenzo CROCCO
Dominique DALLET
Claudio Roberto Maria DE CAPUA
Francesco Giuseppe DELLA CORTE
Loreto DI DONATO
Giuliana FAGGIO
Gioia FAILLLA
Sofia GIUFFRÈ
Voicu GROZA
Gianluca LAX
Aimè LAY EKUAKILLE
Elena Simona LOHAN
Pietro MANZONI
Giacomo MESSINA
Antonella MOLINARO
Andrea Francesco MORABITO
Francesco Carlo MORABITO
Rosario MORELLO
Gabriel-Miro MUNTEAN
Giuseppe MUSOLINO
Fortunato PEZZIMENTI
Filippo Gianmaria PRATICÒ
Sandro RAO
Maria ROMANO
Domenico ROSACI
Giuseppe RUGGERI
Mariateresa RUSSO
Alexey VINEL

Acknowledgment

This work was supported by a PhD scholarship funded by the Italian National Recovery and Resilience Plan (PNRR), Mission 4 'Education and Research', Component 1, Investment 4.1 (Action B), pursuant to Ministerial Decree No. 351/2022, funded by the European Union – NextGenerationEU.

SCUOLA DI DOTTORATO
Università Mediterranea di Reggio Calabria

DIPARTIMENTO
Dipartimento di Ingegneria dell'Informazione,
delle Infrastrutture e dell'Energia Sostenibile (DIIES)

DOTTORATO DI RICERCA
Ingegneria dell'Informazione

S.S.D. IINF-03/A
XXXVIII CICLO

Advanced Solutions for the CCAM ecosystem: Design, Prototyping and Assessment

Dottorando
Domenico Mario Zappalà

Tutor
Prof.ssa Claudia Campolo
Prof.ssa Antonella Molinaro

Coordinatore del Dottorato
Prof. Tommaso Isernia

Finanziato dall'Unione europea NextGenerationEU
Italiadomani

Abstract

The evolution toward Cooperative, Connected and Automated Mobility (CCAM) is transforming road transportation into a cooperative system in which vehicles, Vulnerable Road Users (VRUs), (including pedestrians, cyclists, etc,) and infrastructure exchange information through heterogeneous wireless technologies and integrate Artificial Intelligence (AI) algorithms for prediction and decision support. This paradigm introduces several challenges while enabling safety-oriented services such as cooperative awareness, cooperative perception, and remote assistance. In particular, these services impose performance, interoperability, and context-awareness requirement, which must be met even in the presence of heterogeneous (and potentially resource-constrained) platforms in terms of computational capabilities, communication interfaces, and supported messaging protocols.

In this context, the thesis addresses the design and validation of integrated sensing, communication, perception, and control solutions to enable safe and efficient cooperative mobility services. In particular, the work combines theoretical investigation, architectural design, and experimentation on real prototypes for vehicles and electric-bicycles (e-bike).

The first research line focuses on VRUs, with particular interest on e-bikes. Through the development of on-board hardware prototypes for e-bike, data acquisition and the analysis of mobility patterns jointly with the dynamics of VRU Awareness Message (VAM) generation, new message triggering mechanisms were proposed. These triggers rely on cyclist stability, in line with standardization bodies recommendations, with the objective of increasing road safety in mixed traffic scenarios (vehicles and e-bikes) that characterize CCAM settings.

The thesis also includes the design, implementation, and validation of an open sensing, perception, and communication platform for connected vehicles. The architecture combines low-cost devices for embedded sensing and perception (e.g., cameras) with a multi-radio On-Board Unit (OBU) that supports short-range connectivity via ITS-G5 and long-range connectivity via 5G. Machine learning (ML)

models are appropriately adapted to the operating context to improve surrounding-environment detection and to estimate the driver's attentional state, ultimately enhancing driving safety.

The final part of the thesis considers Teleoperated Driving, a crucial service in the CCAM ecosystem, given the pressure it puts on the network infrastructure, the tight delivery constraints and the presence of human-in-the-loop. The thesis proposes and validates a solution for remote vehicle control that leverages vehicle proximity to dynamically form a platoon and exploits the contextual perception of adjacent vehicles, achieving benefits in terms of reduced radio resources required for video streaming toward the remote operator.

Overall, the thesis shows that context-driven solutions, embedded perception, and communication are a key enabler to achieve scalable CCAM services, particularly when human-in-the-loop operation and high bandwidth demanding services must be supported under realistic network constraints.

Sommario

L'evoluzione verso la mobilità cooperativa, connessa e automatizzata (CCAM) sta trasformando il trasporto stradale in un sistema cooperativo in cui veicoli, Utenti Vulnerabili della Strada (VRU), (tra cui pedoni, ciclisti, etc.), e infrastrutture scambiano informazioni tramite tecnologie wireless eterogenee e integrano algoritmi di Intelligenza Artificiale (IA) per la previsione e il supporto decisionale. Questo paradigma introduce diverse sfide, consentendo al contempo servizi orientati alla sicurezza quali la consapevolezza cooperativa, la percezione cooperativa e l'assistenza remota. In particolare, questi servizi impongono requisiti di prestazione, interoperabilità e consapevolezza del contesto, che devono essere soddisfatti anche in presenza di piattaforme eterogenee (e potenzialmente con risorse limitate) in termini di capacità computazionali, interfacce di comunicazione e protocolli di messaggistica supportati.

In questo contesto, la tesi affronta la progettazione e la validazione di soluzioni integrate di rilevamento, comunicazione, percezione e controllo per consentire servizi di mobilità cooperativa sicuri ed efficienti. In particolare, il lavoro combina indagine teorica, progettazione architettonica e sperimentazione su prototipi reali per veicoli e biciclette elettriche (e-bike).

La prima linea di ricerca è dedicata agli utenti vulnerabili della strada (VRU), con particolare focus sulle e-bike. Tramite lo sviluppo di prototipi hardware a bordo e-bike per l'acquisizione dati e l'analisi dei pattern di mobilità congiuntamente alla dinamica di generazione dei messaggi VRU Awareness Message (VAM), sono stati proposti nuovi meccanismi di trigger per l'invio dei messaggi. Tali trigger si basano sulla stabilità del ciclista, in linea con le raccomandazioni degli enti di standardizzazione, con l'obiettivo di aumentare la sicurezza stradale in scenari di traffico misto (veicoli ed e-bike) tipici dei contesti CCAM.

La tesi include inoltre la progettazione, implementazione e validazione di una piattaforma aperta di sensing, percezione e comunicazione per veicoli connessi. L'architettura combina dispositivi a basso costo per sensing e percezione integrata,

(ad esempio videocamere), e una On-Board Unit (OBU) multi-radio che supporta la connettività a corto raggio tramite ITS-G5 e a lungo raggio tramite 5G. Modelli di Machine Learning (ML) sono opportunamente adattati al contesto operativo per migliorare la rilevazione dell'ambiente circostante e stimare lo stato di attenzione del conducente al fine di migliorare la sicurezza di guida.

La parte finale della tesi prende in esame la guida teleoperata (ToD), un servizio cruciale nell'ecosistema CCAM, data la pressione che essa esercita sull'infrastruttura di rete, i rigidi vincoli di consegna e la presenza dell'uomo nel loop di controllo (human-in-the-loop). E' quindi proposta e validata una soluzione per il controllo remoto dei veicoli che sfrutta la prossimità tra veicoli per formare dinamicamente un convoglio, sfruttando anche la percezione del contesto dei veicoli adiacenti, con benefici in termini di riduzione delle risorse radio utilizzate per lo streaming video all'operatore remoto.

Nel complesso, la tesi mostra che soluzioni context-driven, percezione embedded e comunicazione rappresentano un abilitatore chiave per ottenere servizi CCAM scalabili, in particolare quando operazioni human-in-the-loop e servizi ad alta richiesta di banda devono essere supportati sotto vincoli di rete realistici.

Contents

List of Figures

Abbreviations

3GPP 3rd Generation Partnership Project

4G Fourth Generation

5G Fifth Generation

5GAA 5G Automotive Association

AI Artificial Intelligence

ADAS Advanced Driver Assistance Systems

API Application Programming Interface

BLE Bluetooth Low Energy

C-V2X Cellular-V2X

CABA Context-Aware Bitrate Adaptation

CAM Cooperative Awareness Message

CAN Controller Area Network

CACC Cooperative Adaptive Cruise Control

CCAM Cooperative, Connected and Automated Mobility

CAV Connected and Autonomous Vehicle

CC Cruise Control

CNN Convolutional Neural Network

COCO Common Objects in Context

CPM Cooperative Perception Message

CSMA/CA Carrier Sense Multiple Access with Collision Avoidance

DCC Decentralized Congestion Control

DENM Decentralized Environmental Notification Message

DSRC Dedicated Short Range Communications

DT Digital Twin

NCNN Nihui Convolutional Neural Network

NCAP European New Car Assessment Programme

FPS Frames per Second

eNB evolved NodeB

ETSI European Telecommunications Standards Institute

FT Fine-Tuning

gNB gNodeB

GNSS Global Navigation Satellite System

GPS Global Positioning System

IGG Inter Generation Gap

IMEI International Mobile Equipment Identity

ITS Intelligent Transportation System

JSON JavaScript Object Notation

LiDAR Light Detection and Ranging

LTE Long Term Evolution

MAC Medium Access Control

ML Machine Learning

MQTT Message Queue Telemetry Transport

NR-V2X New Radio Vehicle-to-Everything

NLOS Non-line-of-Sight

OBD On-Board Diagnostics

OBU On Board Unit

PCM Platoon Control Message

PDU Packet Data Unit

PHY physical

PoC Proof of Concept

PL Platoon Leader

PM Platoon member

QoE Quality of Experience

QoS Quality of Service

RAT Radio Access Technology

ROI Region of Interest

RAN Radio Access Network

ROS Robot Operating System

RSU Road Side Unit

SA Standalone

SAE Society of Automotive Engineers

SBC Single Board Computer

SPS Semi-Persistent Scheduling

SU Sensing Unit

SUMO Simulation of Urban MObility

ToD Tele-Operated Driving

UDP User Datagram Protocol

UE User Equipment

V2X Vehicle-to-Everything

V2I Vehicle-to-Infrastructure

V2N Vehicle-to-Network

V2N2X Vehicle-to-Network-to-Everything

V2V Vehicle-to-Vehicle

V2P Vehicle-to-Pedestrian

VAM VRU Awareness Message

VOC Volatile Organic Compounds

VRU Vulnerable Road User

WAVE Wireless Access in Vehicular Environments

XR extended Reality

YOLO You Only Look Once

List of publications

Conferences

1. C. Campolo, A. Molinaro, B. Pizzimenti, M. Schrapel, A. Vinel and **D. M. Zappalà**, *"Understanding Vulnerable Road Users Awareness Messages Generation Dynamics: The Case of e-Bikes,"* 2024 IEEE Vehicular Networking Conference (VNC), Kobe, Japan, 2024, pp. 336-342, doi: 10.1109/VNC61989.2024.10575948.
2. M. Schrapel, T. Graf, **D. M. Zappalà**, C. Campolo, A. Molinaro, and A. Vinel. *When the E-bike Takes Over: Speed Precision and Perception of Cruise Control for Cyclists*. In Proceedings of the 16th International Conference on Automotive User Interfaces and Interactive Vehicular Applications (AutomotiveUI '24). Association for Computing Machinery, New York, NY, USA, 2024, 272–282. https://doi.org/10.1145/3640792.3675714.
3. C. Campolo, N. Mammone, A. Molinaro, B. Pizzimenti, G. Singh and **D. M. Zappalà**, "*Poster: Digital Twins for Personalized and Safer Automated Driving,"* 2024 IEEE Vehicular Networking Conference (VNC), Kobe, Japan, 2024, pp. 277-278, doi: 10.1109/VNC61989.2024.10575986.
4. C. Campolo, A. Molinaro, B. Pizzimenti and **D. M. Zappalà**, "*Poster: Multi-RAT Multi-Protocol OBUs for Vehicle-to-Everything Communications*," 2025 IEEE Vehicular Networking Conference (VNC), Porto, Portugal, 2025, pp. 1-2, doi: 10.1109/VNC64509.2025.11054147.
5. C. Campolo, A. Molinaro, B. Pizzimenti, G. Singh and **D. M. Zappalà**, *"Digital Twin-assisted Context-aware Object Detection for Connected and Automated Vehicles*," 2025 IEEE Vehicular Networking Conference (VNC), Porto, Portugal, 2025, pp. 1-4, doi: 10.1109/VNC64509.2025.11054167.
6. D. Gioffrè., B. Pizzimenti, H. Rashid, **D. M. Zappalà**, C. Campolo, A. Molinaro, and G. Ruggeri, "*Traffic Condition Dissemination through 5G-based Vehicle-to-Network-to-Everything Communications*", 2025 IEEE Conference on Standards

for Communications and Networking (CSCN), Bologna, Italy, 2025, pp. 1-1, doi: 10.1109/CSCN67557.2025.11230751.

Journals

1. C. Campolo, G. Genovese, A. Molinaro, B. Pizzimenti, G. Ruggeri and **D. M. Zappalà**, "*An Edge-Based Digital Twin Framework for Connected and Autonomous Vehicles: Design and Evaluation*," in IEEE Access, vol. 12, pp. 46290-46303, 2024, doi: 10.1109/ACCESS.2024.3382001.

1

Introduction

Over the last decades, transportation systems have undergone deep transformations, driven by the convergence of technological innovation, societal needs, and environmental constraints. The continuous growth of urban populations, combined with increasingly complex mobility patterns, especially in densely populated metropolitan areas, has amplified the demand for safer, more sustainable, and more efficient mobility solutions.

Today, according to the World Health Organization (WHO), road traffic crashes result in more than 1.19 million deaths each year worldwide, with human error identified as the leading cause [5]. Congestion represents another major challenge, particularly in urban areas: the 2024 Global Traffic Scorecard report by INRIX estimated that traffic congestion costs the U.S. economy approximately $74 billion annually in lost productivity and wasted fuel [6].

Beyond safety and efficiency, transportation also has profound environmental implications: road transport is responsible for a significant share of global CO_2 emissions, making it a critical sector in the fight against climate change.

Traditional approaches to addressing these issues have primarily focused on improving standalone solutions for individual vehicles, for example, through safer chassis design, more efficient powertrains, or Advanced Driver Assistance Systems (ADAS). However, such approaches are no longer sufficient to tackle challenges such as road congestion, accidents caused by human inattention, and the carbon footprint of wheeled transportation systems. As a result, both research and industry have progressively shifted their focus toward paradigms that transcend the boundaries of the single vehicle, emphasizing cooperation, communication, and system-level optimization.

Within this context, the concept of **Cooperative, Connected and Automated Mobility (CCAM)** has emerged as one of the most promising paradigms for shaping the future of mobility [7]. Unlike the concept of Connected and Autonomous Vehicles (CAVs), which mainly emphasizes the role of cars equipped with connectivity

and automation features, CCAM provides a more inclusive ecosystem. It integrates a diverse set of actors, such as cars, commercial vehicles, bicycles, pedestrians, and road infrastructure, into a unified framework in which all participants share information and intents to coordinate their mobility patterns.

1.1 Towards CCAM

The transition from conventional mobility to automated transportation has often been framed in terms of CAVs. The CAV concept envisions vehicles equipped with advanced sensors, such as Light Detection and Ranging (LiDAR), radar, and cameras, high-performance on-board computing units, and connectivity interfaces that enable limited data exchange. While CAVs represent a significant technological advancement, they remain inherently vehicle-centric. Their primary focus is on empowering individual vehicles to perceive their surroundings, make autonomous decisions, and interact with the environment, with only limited levels of cooperation with other road users.

In contrast, CCAM builds on the idea that mobility should not be viewed as a collection of independent entities sharing the same physical space, but rather as a dynamic, interconnected system in which information exchange enables coordination among all actors on the road. This cooperative vision is enabled by V2X communications, which allow vehicles to exchange information with other vehicles, infrastructure, and road users.

The availability of timely and cooperative information exchange opens the way to a wide range of advanced CCAM services and safety applications. Cooperative collision avoidance, coordinated lane merging, and collective perception are examples of applications that exploit shared data from multiple sources to improve decision-making and coordinate actions. Moreover, by explicitly incorporating Vulnerable Road Users (VRUs) [8], such as pedestrians and cyclists, into cooperative scenario, CCAM enables benefits that cannot be achieved by isolated CAVs. These include:

- **Safety**: Through cooperative collision avoidance, vehicles and VRUs can share early warnings, enabling proactive responses.
- **Efficiency**: Coordinated traffic flow and platooning can reduce congestion and increase roadway capacity.
- **Environmental sustainability**: Smoother and more predictable traffic flows minimize fuel consumption and emissions, contributing to global climate objectives.
- **Accessibility**: By accounting for the needs of all road users, CCAM promotes inclusive mobility systems that improve both safety and usability.

This broader vision aligns with major policy initiatives worldwide, including the European Commission's Horizon Europe CCAM Partnership [7], which brings together academia, industry, and public stakeholders to accelerate the research, development, and deployment of cooperative mobility solutions.

1.2 CCAM Enabling Technologies

The deployment of CCAM relies on the convergence of multiple technological domains.

1.2.1 V2X Communications

Connectivity is a fundamental enabler of CCAM. V2X communications allow vehicles to exchange data not only with other vehicles (Vehicle-to-Vehicle (V2V)), but also with roadside infrastructure (Vehicle-to-Infrastructure (V2I)), vulnerable road users (Vehicle-to-Pedestrian (V2P)), and backend services through the network (Vehicle-to-Network (V2N)). Through this pervasive exchange of data, V2X communications support cooperative awareness, coordinated maneuvers, and advanced safety and efficiency services.

At present, V2X communications are supported by two main radio access technologies: **IEEE 802.11**, representing a short-range, ad hoc communication paradigm, and **Cellular-V2X (C-V2X)**, which leverages cellular network technologies to enable both direct and infrastructure-assisted communication.

IEEE 802.11

The IEEE 802.11p standard [9] is a Wi-Fi–based radio access technology specifically designed for vehicular communications. It operates in the licensed Intelligent Transportation System (ITS) band at 5.9 GHz (5.85-5.925 GHz) and defines the physical (PHY) and Medium Access Control (MAC) layers used by both the European Telecommunications Standards Institute (ETSI) ITS-G5 and the North American IEEE Wireless Access in Vehicular Environments (WAVE) standards.

A key characteristic of IEEE 802.11p is its ability to support direct, connectionless communication between highly mobile nodes without the need for prior association or centralized coordination. This feature makes it particularly suitable for safety-critical broadcast-based services, relying on **Cooperative Awareness Messages (CAMs)** and **Decentralized Environmental Notification Message (DENMs)**, which require low latency and predictable communication behavior in rapidly changing topologies.

Medium access in IEEE 802.11p relies on a contention-based mechanism derived from Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA). While this approach provides simplicity and low access latency under moderate network load, it may suffer from performance degradation in dense traffic scenarios. To mitigate this issue, ETSI ITS-G5 introduced the Decentralized Congestion Control (DCC) mechanism, which dynamically regulates network load to prevent channel saturation in dense traffic scenarios [10].

To address evolving CCAM requirements, including higher throughput, extended communication range, dense deployment scenarios, and improved support for long packets under high mobility, IEEE 802.11p is being superseded by its next-generation evolution, IEEE 802.11bd [11]. IEEE 802.11bd introduces advanced PHY features such as midambles, higher-order modulation schemes, and channel bonding, while maintaining backward compatibility with IEEE 802.11p to ensure a smooth transition.

Cellular-V2X

C-V2X is standardized by the 3rd Generation Partnership Project (3GPP) and represents a complementary approach to vehicular communications. Initially based on Fourth Generation (4G) cellular networks, C-V2X has progressively evolved toward Fifth Generation (5G) systems through New Radio Vehicle-to-Everything (NR-V2X). By leveraging cellular infrastructure, C-V2X offers wide-area coverage, improved scalability, and seamless integration with existing mobile networks.

The first specifications and use-case definitions for C-V2X were introduced in 3GPP Release 14 [12] and further expanded in Release 15 [13], which focused on LTE-based support for V2X services. Release 14 defined two distinct modes for V2V resource allocation:

- **Mode 3**, where communication resources are centrally assigned by the evolved NodeB (eNB) when vehicles operate under cellular coverage;
- **Mode 4**, where vehicles autonomously select communication resources using a Semi-Persistent Scheduling (SPS) mechanism, enabling direct communication even in the absence of network coverage.

C-V2X supports two main communication interfaces:

- the **PC5 interface**, enabling direct sidelink communication (V2V, V2P, V2I) without infrastructure support;
- the **Uu interface**, allowing vehicles to communicate with the eNB, enabling V2N services such as traffic management, cloud access, and software updates.

With 3GPP Release 15, C-V2X expanded its scope to support both safety-critical services, such as cooperative driving and platooning, and non-safety services, including high-data-rate infotainment and real-time high-definition map updates.

A major evolution occurred with 3GPP Release 16 [14], which introduced **NR-V2X** as the 5G-based evolution of C-V2X. NR-V2X brings significant enhancements at the PHY and MAC layers, including flexible numerology, lower latency, higher reliability, and support for advanced use cases such as cooperative maneuvers and sensor sharing with stringent Quality of Service (QoS) requirements.

In 2021, 3GPP Release 17 [15] consolidated NR-V2X capabilities and introduced further refinements to improve energy efficiency, coverage extension, and positioning accuracy. Notably, it expanded support for **Cooperative Perception Messages (CPMs)**, enabling the exchange of raw or processed sensor data among vehicles and infrastructure. Furthermore, Release 17 addressed extended Reality (XR) applications in vehicular scenarios, paving the way for immersive services such as **ToD**.

The most recent 3GPP Release 18 [16] laid the foundation for a transition toward **6G-V2X**, focusing on ultra-reliable and low-latency communications, enhanced sidelink capabilities, and Low Power High Accuracy Positioning (LPHAP). These advancements aim to further extend the performance envelope of V2X communications and support future CCAM applications with extreme reliability, precision, and scalability requirements.

1.2.2 Heterogeneous Sensing and Embedded Perception

CCAM services build upon heterogeneous sensing, where vehicles, infrastructure, and VRUs contribute observations with different patterns, quality, and update rates. Modern vehicles typically integrate complementary sensors such as cameras, radar, LiDAR, and inertial measurements, enabling robust perception under diverse environmental conditions. Conversely, VRUs often rely on lightweight consumer devices, where sensing capabilities and energy budgets are inherently constrained. Non trivial design trade-offs are introduced by this heterogeneity.

To accommodate embedded platforms, on-board perception pipelines must continue to be precise, low-latency, and resource-efficient. Cooperative services must simultaneously deal with inconsistent data granularity across actors, partial observability, and variable sensing coverage. Lightweight embedded perception solutions and context-aware strategies that adjust sensing and dissemination to device capabilities and operational constraints are motivated by these factors.

1.2.3 Edge Intelligence

Many CCAM applications impose stringent requirements in terms of latency, reliability, and real-time responsiveness. Traditional approaches that rely exclusively on on-board computing resources or centralized cloud infrastructures often fail to meet these demands. On-board units are typically constrained in processing power and energy availability, while cloud-based solutions may introduce excessive communication delays, making them unsuitable for safety-critical and time-sensitive vehicular applications.

To address these limitations, **edge computing** has emerged as a key enabling technology able to bridge the gap between vehicles and the cloud. By relocating part of the computational workload closer to data sources, typically within network nodes located in proximity to vehicles, edge computing offers a favorable trade-off between processing efficiency and latency reduction. Beyond reducing end-to-end delay, edge computing also enables localized data processing, which enhances contextual awareness and allows services to adapt dynamically to the rapidly changing conditions of the vehicular environment.

Building on this infrastructure, the concept of **edge intelligence** arises from the integration of Artificial Intelligence (AI) capabilities within edge computing platforms. Rather than executing complex AI models solely on resource-constrained vehicles or in distant cloud data centers, edge intelligence distributes learning and inference tasks across vehicles and nearby edge nodes. This approach enables the execution of advanced AI algorithms with low latency and high reliability, while limiting backhaul traffic and preserving data locality.

In CCAM systems, edge intelligence plays a central role in vehicle perception tasks such as object detection, lane recognition, and trajectory prediction, which increasingly rely on deep learning techniques. At the cooperative level, edge-enabled AI supports collective perception, whereby data from multiple sources, including vehicles, VRUs, and infrastructure sensors, are aggregated and processed at the edge to generate a more comprehensive and consistent view of the environment. By fusing heterogeneous sensor information beyond the capabilities of individual vehicles, edge intelligence significantly enhances situational awareness, robustness, and safety in cooperative mobility scenarios. Beyond sensing and computation, deployable CCAM services also require communication stacks that explicitly embrace heterogeneity, combining short-range V2X technologies with cellular connectivity and application-layer protocols to ensure service continuity across actors and network conditions.

1.2.4 Multi-Interface and Multi-Protocol Communication for CCAM

The intrinsic heterogeneity of the CCAM ecosystem calls for communication solutions that go beyond a single protocol or radio access technology. CCAM applications differ widely in terms of latency, reliability, bandwidth, and energy constraints, while the actors involved, such as vehicles, infrastructure, and VRUs, exhibit different capabilities and operational requirements. As a result, effective CCAM deployments require both a **multi-protocol** communication framework, capable of selecting the most appropriate protocol for each application, and **multi-interface** support, which ensures flexibility, resilience and seamless connectivity under dynamic conditions.

This heterogeneity also extends to end devices employed by different actors in the CCAM ecosystem. Vehicles can be equipped with powerful on-board computing platforms and multiple communication interfaces, benefiting from large energy and processing resources. In contrast, pedestrians and cyclists typically rely on lightweight, battery-powered devices with limited computational capabilities. Furthermore, CCAM services differ in transmission frequency and computational load, ranging from periodic safety beacons to sporadic event-driven messages. These constraints require energy-aware communication strategies that efficiently manage limited computational and battery resources while still meeting application-specific performance requirements.

Multi-Protocol Communications to Improve Reliability and Interoperability

The CCAM ecosystem involves a wide variety of actors that generate and consume data with different characteristics and temporal dynamics. In this context, adopting a **multi-protocol** communication approach is essential to ensure interoperability and to support efficient information exchange across heterogeneous applications, users and operational scenarios.

Different road users generate different types of messages depending on their role and the application requirements. For example, VRUs transmit VAMs [17], which convey kinematic and contextual information according to ETSI specifications [18]. Similarly, vehicles periodically broadcast CAMs to share their kinematic state with neighboring entities, also following ETSI standard [19]. Both message types are central to safety-critical applications, which impose stringent requirements in terms of latency, reliability, and availability. These applications typically rely on V2X-specific communication technologies, such as ITS-G5 or C-V2X, which are explicitly designed to support real-time, highly dependable interactions in vehicular scenarios.

In safety-oriented CCAM systems, the design of messages like CAMs and VAMs is intrinsically tailored to localized, broadcast-based dissemination within the vehicular domain. These messages are optimized for short-range, low-latency communication, where information relevance is spatially and temporally bounded and delivery is confined to proximate actors.

However, safety-related information increasingly needs to be propagated beyond the local vehicular environment. Road operators, traffic management centers, cloud-based analytics platforms, and third-party safety services require access to selected safety messages for purposes such as situation awareness, incident management, system monitoring, and post-event analysis. This introduces a fundamental change in communication scope: from proximity-based broadcast to infrastructure-assisted, wide-area data dissemination.

V2X communication stacks are not designed to natively support this form of remote dissemination. Their broadcast-oriented operation, absence of session management, and lack of mechanisms for scalable data aggregation make them ill-suited for interfacing with backend systems and external applications. Attempting to extend V2X-native communication mechanisms beyond their intended domain would lead to inefficiencies, unnecessary protocol complexity, and limited interoperability with existing information systems.

As a result, CCAM deployments commonly complement V2X communications with application-layer protocols when safety information must be conveyed to remote backends. Lightweight publish/subscribe solutions, such as Message Queue Telemetry Transport (MQTT) [20] running over cellular networks, are frequently adopted to bridge the vehicular and backend domains. These protocols provide asynchronous delivery, decoupling between data producers and consumers, and seamless integration with cloud infrastructures, while allowing safety-relevant information to be selectively forwarded without disrupting local V2X operations.

Multi-Interface Communication for Continuous Connectivity

Complementing protocol diversity, **multi-interface** communication capabilities are crucial to guarantee continuous connectivity in CCAM environments. In highly dynamic and often unpredictable mobility scenarios, network availability cannot be taken for granted: cellular coverage may be temporarily unavailable in certain areas, while short-range communications may suffer from congestion, interference, or unfavorable propagation conditions. To cope with these challenges, CCAM devices must be capable of dynamically selecting and switching among multiple communication technologies, thereby maintaining continuous data exchange regardless of the underlying infrastructure.

Beyond providing redundancy, multi-interface devices also enable simultaneous communication with multiple actors using different technologies. For example, a single device can transmit safety-critical messages via dedicated short-range V2X interfaces, while forwarding non-critical or high-volume data over cellular networks. This capability significantly improves service continuity, optimizes resource utilization, and enhances the overall resilience of CCAM communication systems.

1.3 Challenges and Open Issues

Despite significant advances in communication technologies, sensing platforms, and automation capabilities, the realization of large-scale, deployable CCAM systems remains hindered by several unresolved challenges. These challenges stem from the intrinsic heterogeneity of vehicular environments, the stringent performance requirements of safety-oriented services, and the growing scale and complexity of cooperative applications.

Interoperability and standardization.

Vehicular environments are intrinsically heterogeneous, characterized by the coexistence of multiple Radio Access Technologys (RATs), vendors, and protocol stacks, often exhibiting different deployment maturity levels and coverage characteristics. Short-range V2X technologies and long-range cellular connectivity are typically developed within distinct standardization frameworks and operational assumptions, complicating their seamless integration.

Achieving interoperability across these heterogeneous communication domains, while preserving efficiency and avoiding excessive system complexity, remains a central challenge for deployable CCAM services. In particular, the lack of unified mechanisms for coordinating data dissemination, service continuity, and protocol selection across interfaces limits the ability of CCAM systems to operate consistently under varying network conditions.

Quality of Service (QoS) requirements.

CCAM systems must support heterogeneous traffic classes, ranging from event-driven safety messages to cooperative perception data and high-rate video streams. These traffic classes impose diverse and often conflicting requirements in terms of latency, reliability, throughput, and jitter. Guaranteeing bounded latency and high delivery reliability under conditions of high mobility, radio interference, and network congestion remains a fundamental challenge.

Meeting these requirements necessitates adaptive QoS strategies that operate across multiple layers, spanning data generation, encoding, transmission scheduling, and edge-assisted processing. In particular, ToD introduces exceptionally stringent real-time constraints and sustained uplink bandwidth demands. Ensuring predictable performance for such services under fluctuating network conditions highlights the limitations of static resource allocation and motivates the need for context-aware optimization mechanisms.

Scalability and resource efficiency.

As the number of connected CCAM participants increases—including vehicles, VRUs, and roadside infrastructure—the cumulative communication load can rapidly saturate available radio resources. Fixed-rate data dissemination and continuous high-bandwidth streaming, if naively applied, do not scale with participant density and may severely degrade system performance.

Scalability therefore depends on mechanisms that reduce communication overhead while preserving situational awareness. Promising directions include event-driven communication strategies for safety messaging, context-aware perception pipelines that prioritize relevant information, and adaptive rate control mechanisms driven by scene dynamics and relevance. Efficient resource utilization is especially critical in cooperative scenarios, where redundant data streams can be selectively suppressed or adapted without compromising safety objectives.

Heterogeneous users and operational constraints.

CCAM systems must accommodate actors with widely different sensing, computing, and energy capabilities. Vehicles typically feature powerful on-board computing platforms and multiple communication interfaces, whereas VRUs rely on resource-constrained, battery-powered devices with limited sensing and processing capabilities. Moreover, interaction patterns and operational constraints vary significantly across roles, contexts, and usage scenarios.

Designing solutions that remain robust across this heterogeneity requires modular system architectures and adaptive policies capable of tailoring behavior to device capabilities, user roles, and environmental context. Without such flexibility, CCAM solutions risk either excluding certain actors or imposing excessive overhead that undermines usability and sustainability.

These challenges highlight the need of designing CCAM as a integrated system. Such a system must combine communication, sensing, and cognitive capabilities, while accounting for real-world limitations. Table 1.1 maps the identified challenges to specific research goals, which this thesis will address, as detailed in the following section.

1.4 Thesis Objectives

Despite extensive standardization and research efforts, existing CCAM solutions typically address communication, sensing, and intelligence as largely independent design problems, often optimized for isolated use cases or homogeneous deployment assumptions. In practice, however, CCAM systems operate across heterogeneous actors, communication technologies, and operational contexts, where strict performance requirements coexist with resource constraints and scalability limitations. Current approaches frequently lack the flexibility required to seamlessly integrate short-range V2X communication with long-range connectivity, to adapt sensing and perception pipelines to dynamic network and context conditions, and to efficiently support high-demand services such as ToD.

This gap motivates the need for integrated, adaptive, and experimentally validated CCAM platforms that explicitly embrace heterogeneity rather than abstracting it away. In particular, there is a lack of practical solutions that combine multi-interface and multi-protocol communication, context-aware data dissemination, and resource-efficient AI-based perception within a unified system design. Addressing this gap requires not only architectural contributions, but also prototyping and real-world evaluation to assess trade-offs between latency, reliability, scalability, and energy efficiency.

In response, this thesis contributes a set of flexible and adaptive CCAM solutions that jointly address communication interoperability, perception efficiency, and network scalability. By integrating heterogeneous communication technologies with lightweight AI-based sensing and context-aware data dissemination strategies, the proposed work advances the state of the art toward deployable CCAM systems capable of operating reliably across diverse actors, environments, and service demands.

From a *communication perspective*, the thesis focuses on the design and deployment of multi-interface and multi-protocol platforms capable of interoperating across heterogeneous networking contexts. The goal is to enable cooperation among vehicles, VRUs, and infrastructure elements independently of the underlying communication technology, effectively supporting an “always connected” paradigm. This includes the investigation of dissemination workflows that combine short-range V2X communications for localized situational awareness with long-range cellular connectivity toward edge services, remote backends, and third-party stakeholders. Particular attention is devoted to preserving efficiency, reliability, and service continuity while traversing heterogeneous communication domains.

From an *AI and sensing perspective*, the thesis investigates low-latency perception solutions for CCAM applications, including embedded object detection and driver

state monitoring. Emphasis is placed on runtime efficiency and on lightweight personalization strategies that enhance robustness across different users and driving behaviors. Edge computing resources are considered as a complementary enabler to support model adaptation, deployment, and inference offloading, allowing advanced perception capabilities to be integrated with limited operational overhead and without compromising real-time constraints.

Finally, the thesis addresses *highly demanding services* characterized by stringent performance delivery requirements, with a specific focus on ToD.

This CCAM use case puts a high pressure on the network infrastructure for the uplink delivery of the video streams captured by on board cameras to the remote operator to whom a proper situational awareness should be provided. Moreover, it entails a stable and low end-to-end latency to guarantee safe maneuvering.

Such features make the use case quite unique and challenging. Indeed, other use cases, like cooperative awareness and perception, mainly rely on data dissemination over sidelink among road entities in close proximity. Even when data dissemination over the cellular infrastructure is needed, the transmission of detected objects and not of raw data is handled, with a limited strain on the network. Moreover, unlike for ToD, human-in-the-loop is not considered.

In this context, the work investigates strategies to reduce network resource consumption in cooperative multi-camera streaming scenarios. By exploiting context awareness, the proposed solutions adapt source rates and streaming policies to reduce bandwidth consumption while preserving operator situational awareness. The resulting design includes the CABA mechanism and platoon-level policies that couple video streaming with geometric safety constraints on inter-vehicle spacing.

Overall, the proposed solutions are validated through both simulation-based evaluations and real-world prototyping and testbed deployments. Performance is assessed in terms of latency, reliability, energy consumption, scalability, and interoperability.

1.4.1 Thesis Outline

The remainder of this thesis is organized as follows.

- **Chapter** 2 investigates the integration of VRUs into the CCAM ecosystem, with a specific focus on the e-bikes and the exchange of standardized VAMs. After reviewing relevant ETSI specifications and the state of the art, the chapter presents a prototype platform specifically designed for data collection and motor control. The prototype is used for:
  - *user-centric* evaluation of cooperative assistance and perceived workload, and

  - *communication-centric* analysis of VAM generation based on real-world traces, including an extension of stability trigger conditions derived from 5G Automotive Association (5GAA) recommendations.

- **Chapter 3** presents an integrated on-board sensing, perception, and communication platform for CCAM. The chapter describes a modular architecture combining low-cost sensing, embedded perception of both road environment and driver status, and heterogeneous short- and long-range communication technologies. Experimental results from multiple testbeds are presented, including on-device perception, lightweight personalization for driver monitoring, and multi-protocol dissemination to both local and remote consumers.
- **Chapter 4** analyzes network efficiency in the context of ToD in vehicular platooning applications. This chapter introduces the CABA mechanism for multi-camera streaming.CABA utilized context awareness and platoon constraints (including geometric considerations of inter-vehicle spacing related to the platoon policy) to reduce uplink bandwidth consumption while maintaining remote operator situational awareness. Finally, the chapter presents a preliminary PoC implementation over a real 5G-Standalone (SA) testbed, which validates the feasibility of the end-to-end teleoperation pipeline under realistic conditions.
- **Chapter 5** concludes the thesis by summarizing the main findings and providing future research directions, with specific focus on enhanced cross-layer integration between sensing, perception, communications, and edge computing for scalable CCAM deployments.

Table 1.1: Mapping of CCAM challenges to thesis objectives and chapters

Challenge	**Research Objective**	**Addressed in Chapter(s)**
Interoperability and standardization across heterogeneous communication technologies	Design and prototype standard-compliant dissemination workflows and platforms enabling cooperation among vehicles and VRUs, and interoperability across short-range V2X and cellular connectivity	Chapters 2, 3
Improved QoS under mobility and congestion	Investigate adaptive dissemination and communication workflows that preserve low latency and high reliability across heterogeneous networking contexts, including support for real-time safety services such as uplink-intensive ToD streaming	Chapters 3, 4
Resource-efficient embedded perception under real-time constraints	Develop and evaluate lightweight on-device AI perception pipelines (e.g., object detection and driver state monitoring), including personalization support for low-latency operation	Chapter 3
Scalability and radio resource usage efficiency in dense cooperative scenarios	Develop context-aware adaptive streaming mechanisms that reduce communication overhead while preserving situational awareness for ToD	Chapter 4
Heterogeneous users and device constraints (vehicles vs. VRUs)	Design modular architectures and lightweight mechanisms that adapt sensing, communication, and processing to device capabilities and operational constraints	Chapters 2, 3
High uplink demand and stringent real-time requirements of ToD	Optimize multi-camera streaming for ToD through context-aware bandwidth adaptation and cooperative policies that respect safety constraints	Chapter 4
Lack of experimental validation in real-world CCAM deployments	Validate proposed solutions through simulation and real-world prototyping, evaluating latency, reliability, energy consumption, scalability, and interoperability	Chapters 2, 3, 4

2

Enabling VRUs awareness in the CCAM ecosystem: the case of e-bikes

This chapter investigates how Vulnerable Road Users (VRUs), with a specific focus on e-bikes, can be made observable within the CCAM ecosystem through the exchange of standardized awareness messages. We first recall how VRUs are defined in the ETSI framework and review the state of the art on VRU detection and communication. We then introduce the e-bike use case and present a prototype platform and data collection campaign. Finally, we report and discuss two complementary studies: a user-centric analysis of cooperative assistance on e-bikes and a communication-centric evaluation of the ETSI standard VAMs generation rules fed with real data and extended with an additional stability-based trigger.

2.1 Context

VRUs are one of the most critical categories of users to be integrated in the CCAM ecosystem, since they are central to the design of cooperative safety and control strategies. Movement behavior of VRUs is not easily predictable, and the level of attention during road crossing is often limited due to the variety of unexpected situations that may occur in urban environments. As a result, pedestrians, cyclists, powered two-wheeler riders, and users of personal mobility devices account for almost 70% of total road fatalities [21].

Differently from CAVs, which use advanced sensing and communication technologies, VRUs remain largely unprotected by current safety systems. Today, VRUs are partially covered by the protection of active ADAS deployed in modern vehicles, such as Automatic Emergency Braking (AEB) and Pre-Collision Systems with Pedestrian Recognition (PCR). However, the effectiveness of these systems strongly depends on the timely recognition of hazardous situations and, in many cases, the available reaction time is insufficient to avoid the accident.

In this context, improving the detection and protection of VRUs has become a priority for both the research and industrial communities, as it represents a key step toward the realization of fully cooperative mobility systems.

The remainder of this chapter is organized as follows. Section 2.2 recalls how VRUs are defined and classified within the ETSI framework and introduces the standardized VAMs used to disseminate VRU awareness. Section 2.3 reviews the state of the art on VRU detection and on network-based/hybrid approaches for cooperative awareness. Section 2.4 motivates the e-bike use case and discusses the main challenges for integrating e-bikes in the CCAM ecosystem, while Section 2.5 summarizes the original contributions of this work. Section 2.6 then presents the proposed framework, including the prototype platform and the data collection campaign. Finally, Section 2.7 reports the user-centric study on cooperative assistance and Quality of Experience (QoE), and Section 2.8 presents the communication-centric analysis of ETSI VAMs generation rules and the proposed stability-based trigger extension. Section 2.9 concludes the chapter by summarizing the main findings and outlining implications and limitations.

The main achievements of the two conducted studies have been published, respectively, in [1] and [2].

2.2 VRUs in ETSI specifications

2.2.1 Definition and classification of VRUs

ETSI has formally included VRUs within its ITS specifications. These users are categorized into specific profiles based on their speed, the environment they are in, and their similarities [17], as shown in Tab 2.1:

Profile	**Category**	**Examples / Characteristics**
1	Pedestrians and similar users	Walkers, joggers, prams, road workers; no mechanical assistance, low speed, highly unpredictable trajectories.
2	Light vehicle riders	Cyclists, e-scooters, skaters, wheelchair users, e-bikes; up to 25/32 km/h; medium speed
3	Motorized two-wheelers	Motorcycles, side-cars; higher speeds and longer ranges;
4	Animals	Domestic or wild animals interacting with traffic (e.g., dogs, horses, wildlife); mostly detected indirectly, unpredictable behavior.

Table 2.1: ETSI VRU Profiles [17].

Among these categories, Profile 2 is of particular relevance for this study. This group typical includes cyclists and other light vehicle users, possibly equipped with an electric motor. It includes cyclists on conventional or electrically assisted bicycles,

wheelchair people, horseback riders, skaters, e-scooters riders, and people using personal transport devices.

2.2.2 VRU awareness messages and generation rules

One of the key enabling requirements for integrating VRUs into the CCAM ecosystem derives from their ability to exchange messages with other entities on the road.

ETSI has standardized a specific category of messages for VRUs, called VAMs [18]. VAMs are generated following event-based triggers combined with periodic refresh mechanisms defined in [18]. They are structured as ITS Packet Data Unit (PDU)s composed of mandatory and optional containers.

The mandatory containers are:

- **ITS PDU Header VAM**: it includes fields such as *protocolVersion*, *messageID* (identifying the VAM type) and *stationID* of the originating ITS station.
- **Generation Delta Time**: a data element encoding the generation time of the reference position in the VAM, derived from the global *TimestampIts* and wrapped modulo 2^{16}.
- **Basic Container**: it contains information such as the *stationType* of the VAM originating device and the *referencePosition* with its accuracy.
- **VRU High Frequency Container**: it contains kinematic information such as *heading* and *headingAccuracy* with respect to true north, *speed* and *speedAccuracy*, *longitudinalAcceleration*, and optional fields such as *curvature*, *yawRate*, *lateralAcceleration*, and *vruLanePosition*.

Other containers, such as the *VRU Motion Prediction Container*, *Cluster Information Container*, and *Cluster Operation Container*, are optional and are used to provide higher level information related to VRU motion intent, group dynamics, or cluster management.

Generation rules can distinguish between changes caused directly by the VRU and those related to the dynamics of clustering conditions. According to the ETSI definition [17], a cluster is a group of VRUs sharing similar behavior and moving within a defined area. The concept of clustering is introduced to reduce channel load and improve scalability in dense scenarios, as the cluster can be managed as a single logical entity when generating VAMs.

ETSI event-based triggers include:

- **Time-based rule**: a new VAM must be transmitted at least every 5 s ($T_{\text{GenVAMMax}}$) since the transmission of the last VAM.

- **Kinematic changes**: a new VAM is triggered if the VRU position changes by more than 4 meters, heading changes by more than 4°, or speed changes by more than 0.5 m/s compared to the last transmitted VAM.
- **Proximity and collision risk**: a VAM is triggered when trajectory intersections or dangerously close encounters with other road users are detected.
- **Clustering**: if a VRU joins or leaves a cluster, a new VAM will be generated to update the neighboring entities.

The minimum time elapsed between the start of consecutive VAM generation events shall be equal to or greater than T_{GenVAM}, which must satisfy:

$$T_{\text{GenVAMMin}} \leq T_{\text{GenVAM}} \leq T_{\text{GenVAMMax}}. \tag{2.1}$$

Here, $T_{\text{GenVAMMin}}$ and $T_{\text{GenVAMMax}}$ are, respectively, the minimum and maximum time intervals between consecutive VAM generation events. Their recommended values are 100 ms and 5000 ms, respectively.

The triggers related to VAM generation must be checked periodically, with a period defined as $T_{\text{CheckVamGen}}$. The parameter $T_{\text{CheckVamGen}}$ shall be equal to or less than $T_{\text{GenVAMMin}}$.

2.3 Analysis of the state of the art

Making VRUs observable to other road users and to the infrastructure can follow three complementary approaches [22]:

- **Passive detection and "other" vision based-perception solutions**: VRUs are not able to actively share information; they are detected by other actors on the road, such as vehicles, through on-board sensors (e.g., camera, radar, LiDAR) and traffic cameras.
- **Active detection and network-based perception**: VRUs themselves become network participants, broadcasting their presence and kinematics using standardized messages like VAMs and thus, connecting directly to the Cooperative Adaptive Cruise Control (CACC) via wearable and portable devices.
- **Hybrid detection**: in this case the VRUs'identity and state are inferred from both self-reported data (active messages) and external detection from other road users. This approach uses sensor data from multiple sources to improve reliability and accuracy, resolve ambiguities, and reduce the occurrence of false positives and negatives in complex urban environments.

In the context of passive detection, one of the main challenges is the timely recognition of VRUs and the prediction of their movements, which are often difficult to

anticipate. Current approaches heavily rely on Machine Learning (ML) techniques applied to real-time video streams using computer vision. These techniques combine complementary tasks, such as object recognition, localization, tracking across frames, pose estimation, and image segmentation [23]. In particular, pose estimation and segmentation allow not only to detect the presence of VRUs, but also to anticipate their possible actions, a capability that is crucial in safety-critical situations.

These ML tasks can be deployed on the infrastructure side or on board the vehicles.

Infrastructure based approach solutions typically involve either structural elements, such as dedicated road design elements, road signs, or traffic calming measures and the use of intelligent components [24]. The last one includes smart traffic light controllers, Road Side Units (RSUs), and network edge servers which can detect VRUs and adapt traffic flows to reduce collision risks.

Vehicle-based approaches require instead more specific AI algorithms due the limited computational resources mounted on board of the vehicles. In this context, the authors in [25] proposed a real-time framework for the detection, tracking, and vulnerability assessment of pedestrians and cyclists, leveraging dashboard-mounted cameras, deep learning techniques for accurate recognition, and unsupervised learning methodologies for tracking purpose. Similarly, the authors in [26] investigate a data-driven approach for VRU trajectory prediction, integrating elements such as VRU intention, heterogeneity, and interactions with both vehicles and other VRUs to improve the reliability of predictions in autonomous driving contexts.

However, this type of passive VRU detection performed by other actors along the road suffers from significant limitations, particularly under Non-line-of-Sight (NLOS) conditions. This type of systems needs cameras to visually detect VRUs, which loose their effectiveness in complex urban environments or under non perfect weather conditions.

In the CACC framework, where every participant is able to perceive other users through network-based approaches, VRUs become active participants capable of broadcasting their kinematic parameters and even their intentions to surrounding vehicles and infrastructure via V2X communications.

Several studies have investigated such systems. The work in [27] integrated laser-based perception with cooperative detection, facilitating the identification of pedestrians in NLOS situations; the assessment was restricted to presence detection, excluding risk evaluation. Likewise, the authors in [28] highlighted the advantages of information fusion between perception and V2P communications, demonstrating that perception enables accurate classification, whereas communication provides precise positioning, thereby complementing each other effectively.

Further studies have explored standardized testing protocols defined by European New Car Assessment Programme (NCAP) [29] to evaluate the performance of cooperative safety systems. The authors in [30] compared camera-based and cooperative collision avoidance systems, following the NCAP AEB-VRU protocol [31], and showed the benefit of using 5G and Dedicated Short Range Communications (DSRC) based cooperation for VRUs in NLOS conditions. Extending this approach, the work in [32] included contextual information from user devices like smartphones and smartwatches. This showed that adding activity recognition, such as detecting when a pedestrian is crossing a curb, could significantly improve collision detection performance.

The effectiveness of network based and hybrid solutions depends heavily on the accuracy and update frequency of Global Positioning System (GPS) data in VAMs. Positioning inaccuracies and infrequent updates continue to represent significant obstacles, potentially compromising the timely situational awareness needed for collaborative safety. To address these problems, different mitigation strategies have been suggested to mitigate their impacts. The authors of [33] applied a filtering approach to mitigate GPS errors, thus improving predictive accuracy within a simulated environment. Other research efforts have focused on overcoming the limitations of Global Navigation Satellite System (GNSS) updates. For instance, [34] utilized a smartphone-based accelerometer classifier to identify pedestrian movements within a 500 ms window representing a notable advancement compared to the 4-s interval generally associated with GNSS. Likewise, the work presented in [35] integrated GPS and accelerometer data, resulting in recognition accuracies exceeding 93% and a reduction in detection delays of up to 1.5 s.

In addition to positioning, the exploitation of contextual information has also been investigated. The authors in [34] demonstrated that including context information linked to VRU, such as location, activity, or distraction level, can improve efficiency, optimize warning delivery, and improve prediction accuracy. Likewise, the study in [32] analyzed contextual risks including curb crossing behavior, demonstrating demonstrating that cooperative systems may substantially enhance collision detection, even despite communication delays or positioning mistakes.

Recent research has focused on the analysis of VAMs, examining message generation dynamics, transmission reliability, and the impact of delay and channel load on cooperative awareness performance.

The study in [36] reports Long Term Evolution (LTE)-V2X field tests with a bicycle and an e-scooter, showing that with the default heading threshold of 4° heading-based triggers dominate and many VAMs are generated on straight segments; by

raising the heading threshold to 7°–10° they demonstrated a substantial reduction of such spurious messages without degrading the safety relevant update distance.

The work in [37] implemented a full ETSI VRU Basic Service in both simulation and field tests with an equipped stroller, confirming that higher heading thresholds (e.g., 10°) are preferable in realistic conditions affected by GNSS drift, and showing that redundancy mitigation can drastically reduce redundant VAMs while keeping the spatial update distance small.

In [38] the authors analyzed the current ETSI VAM generation rules on a large pedelec data set, quantifying the importance of each trigger through an information loss metric, and showed that the distance rule is indispensable, the heading rule is second in importance, and the speed rule contributes the least additional information; they also suggested a minimum sampling rate of 4 Hz for GNSS and Controller Area Network (CAN) data and reported average VAM generation frequencies up to 2.5 Hz for pedelecs.

The work in [39] analyzed the effect of different VAM sampling rates (from 1.25 to 10 Hz) on inter-generation gaps, packet delivery ratio and channel occupancy for e-bikes, showing that moderate rates can maintain high reliability while keeping the channel load well below congestion thresholds and highlighting the dominant role of position-based triggers.

Overall, these works indicate that VAM generation is often more frequent than originally assumed, that standard trigger thresholds may require tailoring to specific VRU types (e.g., e-bikes), and redundancy mitigation mechanisms must be carefully parameterized to balance channel load, positional accuracy, and awareness timeliness.

Such premises motivate the study reported in this chapter.

2.4 The use case of e-bikes

The use case of e-bikes is emblematic in the VRU context. Their rapid adoption, driven by sustainability policies and growing demand for flexible commuting, positions them at the center of micromobility.

Findings in the literature cover pedestrians [40, 41, 42] and bicycles [43, 44] in V2X settings. However, e-bikes exhibit distinct dynamics [45], behavioural patterns [46], higher mass and speeds [47]. ETSI trigger rules are general; tuning their thresholds and frequencies to cyclists' mobility is therefore essential to understand e-bike-specific VAM generation.

Unlike traditional bicycles or pedestrians, e-bikes and e-scooters can automatically reduce speed when a risk is detected, independently of rider attention.

2.4.1 Challenges

To insert e-bike in the CCAM ecosystem is essential to understand the main challenges to complete this task.

- **Integration of new external hardware**: The first major challenge is the integration of a new external platform within the electrical system of the e-bike. Interoperability between the two entities is not natively supported. It required different new connections and software development to address goals. This integration must guarantee both reliability and safety, while minimizing the impact on the native functioning of the e-bike and its factory specifications.
- **Energy efficiency and device constraints**: The design of VRU-oriented devices is constrained by size, weight, and energy efficiency. Pedestrians or cyclists cannot carry heavy hardware. Instead, communication relies on smartphones [48], wearables, or lightweight embedded platforms [49]. The continuous sensing and frequent message transmission risk to rapidly consume batteries. This makes energy-aware communication strategies essential. In the case of e-bike the energy can be taken directly from the on-board battery but also here the energy plays a crucial role in the development of the framework.
- **QoE**: In addition to technical challenges, it is crucial to account for the QoE perceived by users. Workload studies show effects related to cognitive and physical effort and demanding [50], [51], [52], comfort [53], [54], [55] and safety perception [56], [57] of external system with the users. However, cooperative features must remain **non-intrusive**: excessive alerts, overly complex feedback mechanisms, or poorly timed interventions may distract riders rather than protect them, undermining trust and usability.
- **Rider's confidence with e-bike**: beyond differences in experience levels [58], e-bikes present different riding dynamics compared to traditional bicycles. The electrical powertrain alters acceleration, braking, and cornering, supporting the rider with varying levels of torque depending on the selected mode. As a result, experienced and inexperienced riders may react differently to changes in cycling dynamics caused by the motor, and to future assistance strategies that modulate the progression and cut-off of motor power.

2.5 Contributions

Given the aforementioned premises, the objective of the study presented in this Chapter is the design and development of a unified framework, implemented on a lightweight prototype, for integrating VRUs into CCAM, with a particular focus on e-bikes. The main original contributions can be summarized as follows:

- **Hardware prototyping for data collection and control**: a multi-interface prototype platform was designed and implemented to collect realistic datasets from e-bike riders in real-world conditions and to modulate motor assistance. The device integrates GNSS, speed, cadence, and gyroscope measurements, while also logging motor assistance states and injecting trigger signals to adjust speed limits.
- **Profiling methodologies for riders**: field trials with participants of different skill levels enabled the analysis of heterogeneous riding styles and their impact on mobility traces and VAM generation patterns.
- **QoE evaluation**: we assessed user perception, workload, and perceived control to assure that cooperative assistance remains not intrusive to riders.
- **Analysis of ETSI VAM triggers under adverse GNSS**: starting from real world traces, we quantified irregularities in VAM generation along urban routes with degraded GNSS signals, and evaluated the resilience provided by local sensing systems such as wheel speed.
- **Gyroscope-based stability trigger**: building on recommendations from 5GAA working groups, we introduced a stability driven trigger derived from gyroscope data. This trigger captures sudden instabilities and near-stop maneuvers not covered by current rules, and we evaluated its impact on the timeliness of VAM exchange and channel usage.

2.6 Proposed Framework

2.6.1 Prototype

To address the defined objectives, the first step consisted in the development of a prototype system. An Arduino Uno Wi-Fi Rev. 2, embedded within the e-bike and equipped with Bluetooth Low Energy (BLE), was used to monitor key operational parameters. Specifically, the prototype records the current speed, cadence, motor assistance state (on/off), adjusted speed limit, latitude, longitude (using U-BLOX NEO-7M GPS module), and timestamp, storing all data on an SD memory card. Moreover, the system collects motion information from the accelerometer and gyroscope module integrated in the Arduino board.

In addition to passive monitoring, the prototype enables active motor control, similar to Cruise Control (CC). By injecting trigger signals into the embedded motor control unit, the system enforces dynamic speed limits without bypassing regulatory constraints. This feature allows the study of adaptive motor disengagement and its role in workload reduction, speed perception, and cooperative traffic scenarios. An

Android BLE application was developed to adjust the assistance threshold in real time.

In this context, trigger injection refers to the insertion of an electrical signal into the sensor data line, reproducing the effect of a magnet passing by the sensor.

The prototype was integrated into HNF Nicolai UD4 All-Terrain 27.5" e-bikes (sizes M and L), equipped with the 4th Generation Bosch Smart System Performance Line CX motor (Fig. 2.1).

Fig. 2.1: The developed prototype [1].

This motor delivers a maximum torque of 85 Nm and provides multiple assistance modes to regulate torque output. For instance, the Eco mode prioritizes energy efficiency and conversely, Turbo mode maximizes motor power output. According to the manufacturer specifications, the force of the manual pedal increases by 60% in the Eco mode and by up to 340% in Turbo mode. The code used for device tuning and the dataset are available at: https://github.com/M-Schrapel/E-bike-Cruise-Control/tree/main

2.6.2 Data Collection and Profiling of Participants

Field tests were conducted in the city of Karlsruhe (Germany) during autumn 2023, the first period spent abroad during the PhD course.

A total of 15 participants (13 males, 2 females), aged from 23 to 40 years (mean = 27.7, standard deviation, SD, = 4.7), with average height 1.80 m (SD = 0.11) and average weight 80 kg (SD = 18.5), took part in the studies. Recruitment was done via university mailing lists and messaging groups, and experiments were carried out exclusively in daytime.

The data collection occurred into two complementary scenarios:

- **Urban trip** (1.1 km route, mixed environment): Participants' trajectories, speeds, and stability were collected into dataset to be analyzed and to understand the VAMs generation rate according to ETSI thresholds, later extended with gyroscope-based stability triggers.
- **Controlled circular track** (2.1 km across multiple trials): Participants cycled under varying motor assistance modes (None, Eco, Turbo) and target speed limits (16, 19, 22, 25 km/h). After the trials, participants compiled different questionnaires, to obtain a qualitative feedback about personal perceptions related to workload, and usability of dynamic motor assistance.

Tracks used for the studies are presented in the map in Fig. 2.2.

Fig. 2.2: Map of the considered trips [1].

This dual methodology combined *realistic traffic conditions*, captured during preliminary journeys in the city center (blue path in Fig. 2.2) to account for GPS limitations caused by NLOS areas and tree-induced interference while preserving the natural cycling behavior of participants, and the *controlled testing environments* (path in red in Fig. 2.2), designed to enable a reproducible assessment of speed accuracy, cognitive workload, and perception of the new e-biking technology.

In the following two main studies are discussed: one focusing on cyclist perception, workload, and QoE; the other addressing VAM generation dynamics and a proposed stability-based trigger extension.

2.7 E-bikes in the CCAM ecosystem: a user-centric perspective

The objective of this first study was to quantify the accuracy with which cyclists maintain target speeds under active motor control, i.e., a cruise-control-like assistance for e-bikes. The study further aimed to characterize the interaction between the assistance mechanism and the rider, and to assess its impact within the human–machine loop.

2.7.1 Settings

On a circular track, the target speed was set to 16, 19, 22 or 25 km/h via an experimental BLE application, while the motor support was configured in *Eco* or *Turbo* mode in a counterbalanced order across participants. The selected targets cover a representative range of urban cycling speeds: 16 km/h approximates the median speed on shared path [59], 19 km/h matches the average speed of e-bike users in urban areas [47], and 22 km/h has been adopted to study cooperation with traffic lights and modified e-bikes [60]. The limit of 25 km/h has used as a final benchmark to better understand the human feeling with the system. A *None* condition (no electrical support) was included to observe behaviour without assistance.

After each trial, participants reported the perceived maximum assisted speed and completed the NASA-Task Load Index (TLX) questionnaire [61] to assess workload. They also provided a one-word descriptor of the riding experience (prompted by examples such as *dynamic*, *lively*, *quiet*, *rough*, *sporty*, *elegant*, *comfortable*, *strenuous*, *stressful*). Perception of motor disengagement and the tendency to counteract it were collected on a five-point Likert scale. At the end of the session, participants filled in the *Affinity for Technology Interaction* (ATI) [62] and *System Usability Scale* (SUS) [63] questionnaires, and responded to general items with an optional free-text field for comments.

2.7.2 Results

Figure 2.3 summarizes the distribution of actual cycled speeds across all motor modes and targets.

Measured speed increases with higher targets. Motor mode also matters: Turbo tends to produce slightly higher speeds with tighter dispersion than Eco.

A mode–by–target interaction is apparent, with assistance effects more visible at lower targets (e.g., 16 km/h), where fluctuations are more pronounced.

Non parametric analyses confirmed that both target speed and motor mode significantly affect measured speed, and that their interaction is significant. Most pair-

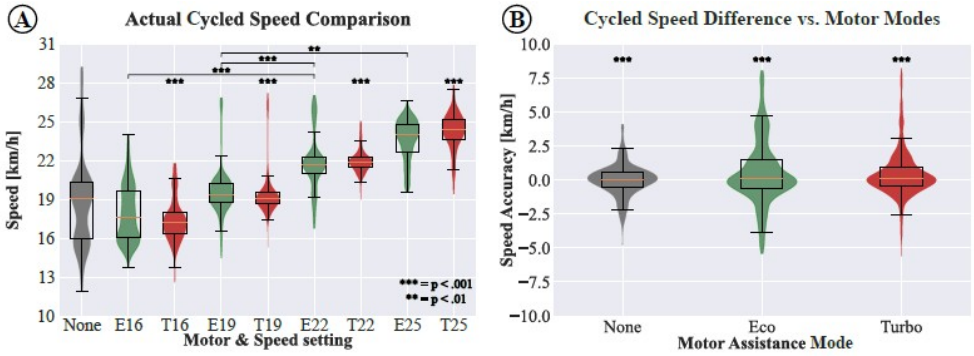

Fig. 2.3: Actual cycled speed comparison. Figure A shows the cycled speeds in each trial and figure B the normalized cycled speed distribution for each motor mode. The violin plots marked with asterisks but without lines show significant differences to all possible combinations [2].

wise comparisons differ, with a few adjacent target pairs showing no clear separation (e.g., E16 vs. E19/E25; E22 vs. E25).

To compare stability across modes, speeds were centered to the target; for *None*, samples were centered to each rider's preferred speed. After centering, differences between modes remain: Turbo shows the smallest fluctuations around zero, Eco exhibits wider variability, and *None* reflects individual pacing.

Without assistance, riders' preferred speed averaged about 18.95 km/h (range from 15.4 to 27.3 km/h), indicating substantial inter-individual diversity.

Cadence remained close to 60 rpm across all modes (Eco and *None* slightly higher than Turbo), consistent with low physiological demand and steady pedaling mechanics. This suggests that assistance mainly supports torque requirements while riders naturally maintain their comfortable pedaling rate. This benefits also reflect on self-estimating speed perception of the riders (Fig. 2.4).

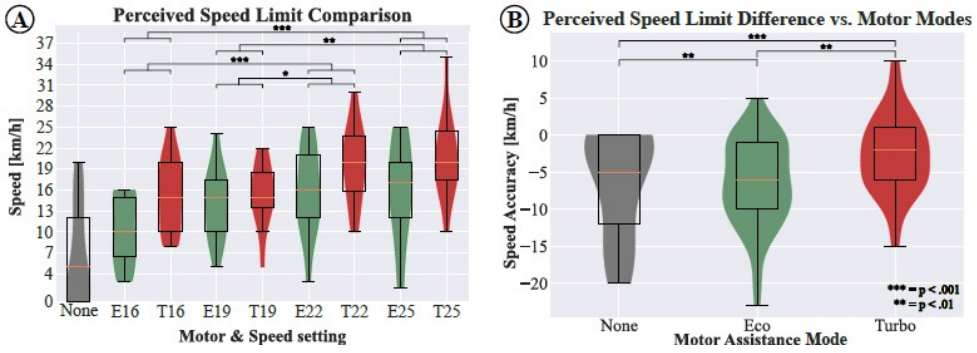

Fig. 2.4: Perceived speed comparison. A shows the reported speeds after each trial and B the normalized reported speed distribution for each motor mode [2].

Data analysis showed significant main effects of *motor mode* (Eco vs. Turbo) and *target speed* on reported speed. Post-hoc tests confirmed differences between Eco and Turbo, and between most target pairs (16 vs. 22/25; 19 vs. 22/25). Power analysis indicated adequate sensitivity for the main effects (but not for the interaction), explaining the non-significant mode target term.

To incorporate *None*, Eco/Turbo reports were centered to their targets and *None* treated as 0 km/h; a *rank-based repeated-measures analysis across modes* indicated significant differences in *how* users perceive the system in different motor modes. Overall, participants *underestimated* their speed by $M = -4.61$ km/h ($SD = 2.40$) across modes; the bias was smaller in *Turbo* (-2.70 ± 1.45 km/h) and larger in *Eco* (-6.52 ± 1.47 km/h). Six participants correctly reported 0 km/h in *None*; four indicated a brief initial assistance (5 km/h), consistent with perceived motor behavior (Fig.2.4A–B).

As regard the workload and cycling behavior, the results are presented in Figure 2.5, according to NASA-TLX Scores.

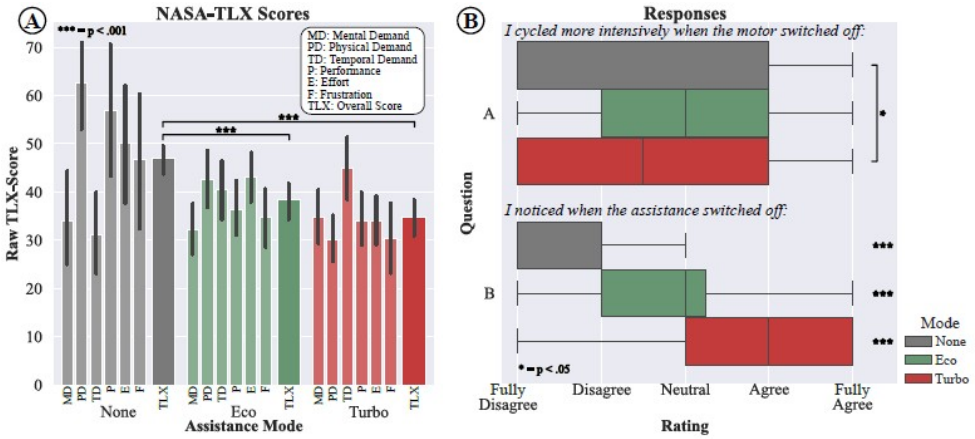

Fig. 2.5: Raw NASA-TLX scores (left) and reported behavior and perception of the motor disengagement (right). Turbo mode achieves the lowest workload and can be easily noticed. Questions A and B in the right figure are mentioned above the corresponding responses [2].

A *rank-based repeated-measures analysis across modes* showed significant differences among *None*, *Eco*, and *Turbo*. Comparisons indicated that *None* differed from both assisted modes, whereas *Eco* vs. *Turbo* was n.s. ($p = 0.569$). Mean TLX decreased with assistance: *None* $M = 46.89$ ($SD = 11.88$), *Eco* $M = 38.22$ ($SD = 15.40$), *Turbo* $M = 34.67$ ($SD = 15.14$). Higher scores in *None* were mainly driven by *Physical Demand* (PD), *Effort* (E), *Frustration* (F), and reduced overall *Performance* (P), as visible in Fig 2.5A.

Perception of assistance switch-off showed a strong trend and all pairs differed, with *Turbo* most detectable (Fig. 2.5B). Taken together, assistance reduced workload and improved the legibility of control actions, supporting steadier pacing and more regular self-reports.

To enhance our comprehension of cruise control for e-bikes and to investigate prospective future applications, we incorporated supplementary inquiries into the concluding questionnaire, as depicted in Figure 2.6.

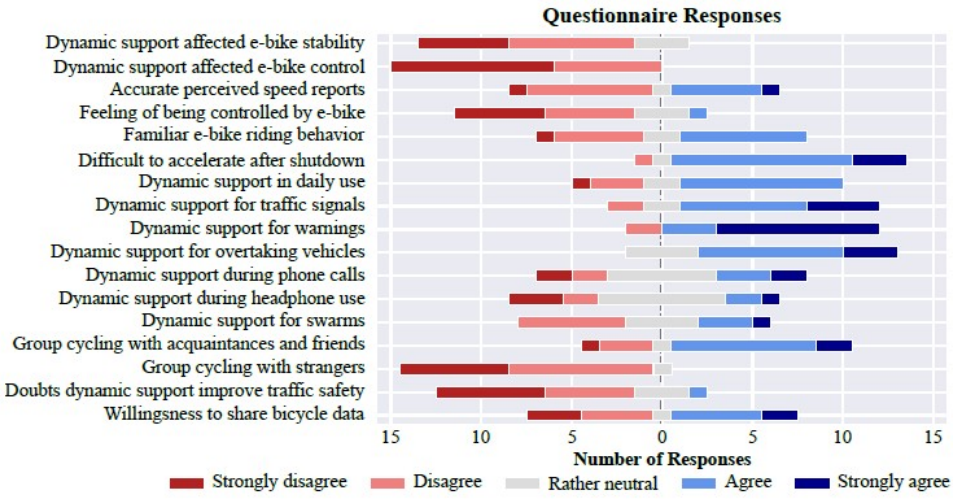

Fig. 2.6: Questionnaire responses collected through a five-point scale set of questions regarding future use cases and the e-bike cycling behavior [2].

From the results obtained by the QoE of participants, they perceived that dynamic support had *little impact* on e-bike stability and control (skew towards disagree/neutral), and did *not* feel controlled by the e-bike, aligning with the goal of non-intrusive assistance. Perceived accuracy of self-reported speeds showed mixed opinions, which is consistent with the systematic underestimation found in the objective/subjective analyses.

Many participants reported *difficulty accelerating after shutdown* (13 agreed), indicating that motor disengagements are salient and can momentarily raise effort to predictable transitions. Acceptance for daily use was positive with strongest interest in *traffic-signal coordination* and *warning* use-cases; support for *overtaking* alerts was also favorable. In contrast, triggers during *phone calls* or *headphone* use were less appealing.

Social context mattered: riders favored *group cycling with acquaintances and friends*, while *swarms with strangers* were less welcome. Participants generally believed dynamic support can *improve traffic safety*; only a small minority expressed doubts.

Willingness to *share bicycle data* was polarized, revealing notable privacy concerns in roughly half of the sample.

2.8 E-bikes in the CCAM ecosystem: a communication-centric perspective

From a communication perspective, the analysis focused on the generation of VAMs following ETSI rules and on the time between two consecutively generated VAMs, referred to as the *inter-generation time* (T_{GenVAM}).

Results are reported for two different participants representative of an experienced cyclist and an inexperienced one. The distinction is based on the preliminary questionnaire filled out by each participant. Specifically, the survey also asks about the perceived average speed during trips, the number of rides and kilometers covered per week, and the type of bike usually employed. The experienced cyclist reports higher values to the aforementioned questions compared to the inexperienced one.

Traces collected during the 1.1 km urban trip revealed that both experienced and inexperienced riders generated VAMs with highly variable inter-generation times, though remaining within the 5-s bound imposed by the standard.

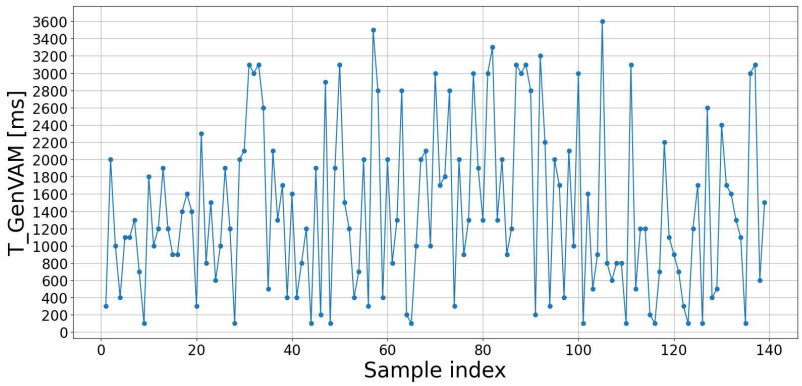

Fig. 2.7: Inter-generation time, T_{GenVAM}, for the experienced rider along the 1.1 km urban route [1].

Figure 2.7 reports the T_{GenVAM} sequence for an experienced rider: inter-generation times between two consecutive VAMs fluctuate between 200 and 2200 ms, with occasional peaks up to about 3.5 s. Short intervals correspond to phases where several kinematic triggers are activated in quick succession, whereas longer intervals ap-

pear when the rider maintains nearly constant speed and heading or when the route crosses segments with degraded or missing GNSS coverage.

A critical aspect of the campaign was the presence of segments with partial or no GNSS coverage, replicating real world NLOS conditions. Since speed, heading and position in VAMs are typically derived from GNSS, satellite loss inflates uncertainty on triggers. Local wheel speed sensing mitigated this by sustaining VAM generation until GNSS was restored.

Results show that the experienced cyclist generated 140 VAMs with an average T_{GenVAM} of 1428 ms, while the inexperienced cyclist generated 151 VAMs with average T_{GenVAM} of 1687 ms, on the same route and conditions. The more regular riding style of experienced users led to fewer speed change triggers; novice riders produced additional VAMs due to more frequent speed variations (Fig. 2.8).

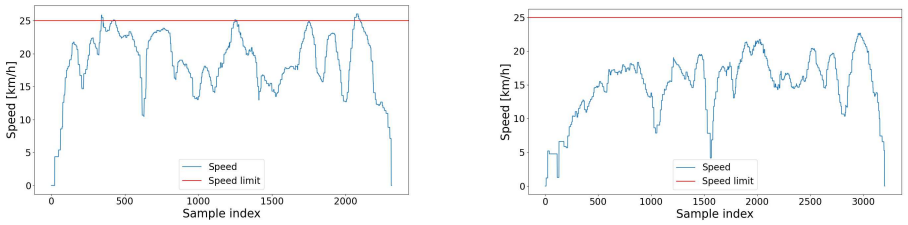

(a) Experienced (avg speed=18.68 km/h). (b) Inexperienced (avg speed=15.06 km/h).

Fig. 2.8: Average speed for the experienced (a) and inexperienced cyclists (b) [1]

Trigger composition for the two cyclists is reported in Fig. 2.9. In both cases, position changes dominated, while heading-based triggers were marginal due to GNSS noise and NLOS segments. In normal conditions, heading should be a common trigger for VAM generation [37], but here GNSS noise contradicts this expectation in the urban environment.

The duty cycle, defined as the percentage ratio of the transmitter's total "on" time on single carrier frequency over a 1-s period [64], remained very low (≈0.07%), well below the 3% threshold defined by ETSI [64].

To capture additional conditions that are difficult to derive exclusively from GPS-based heading and position data, we introduced, following 5GAA recommendations, a new stability trigger based on rider dynamics. This trigger uses gyroscope data to evaluate variations in angle along the y-axis (roll), as shown in Fig. 2.10.

Following ETSI, trigger polling used $T_{\text{CheckVamGen}} = 100$ ms. The stability trigger was evaluated for thresholds $\Delta_g = 4°$, $7°$ and $10°$ relative to the last VAM, and combined with the standard rules.

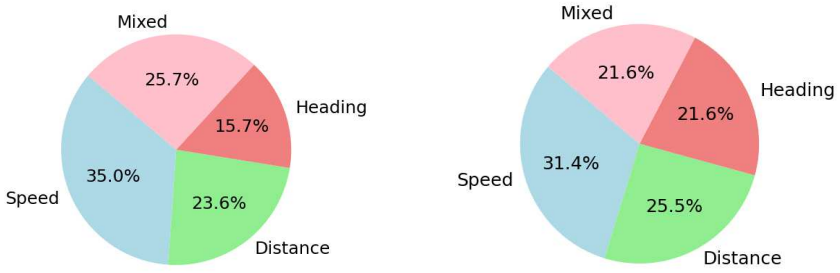

(a) Experienced (b) Inexperienced

Fig. 2.9: Distribution of VAM triggers with ETSI rules only [1].

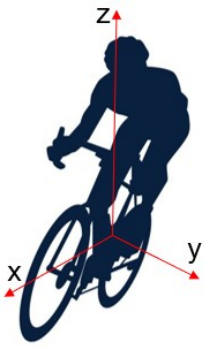

Fig. 2.10: Reference axis for stability-based triggering.

The stability trigger substantially altered trigger composition and rate (Fig. 2.11). For the experienced rider with $\Delta_g = 4°$, T_{GenVAM} decreased from ~1428 ms to ~338 ms, with stability responsible for approximately 81.6% of the triggers; for the inexperienced rider, T_{GenVAM} was ~340.8 ms with 83.9% stability triggers. The duty cycle stayed compliant (≈0.3%), and the additional overhead was minimal (+6 Bytes per message), for an overall generation rate of about 2.3 kbps per cyclist. Table 2.2 reports the metrics for the experienced rider for different thresholds.

Table 2.2: VAM metrics for the experienced cyclist [1].

Parameter	**ETSI only**	**ETSI+Δ_g=4°**	**ETSI+Δ_g=7°**	**ETSI+Δ_g=10°**
T_{GenVAM} [ms]	1427.88	338.24	546.18	743.13
Heading trigger [%]	15.7	1.9	2.7	5.2
Gyroscope trigger [%]	–	81.6	68.6	55.4
Duty cycle [%]	0.07	0.30	0.18	0.13
Total generated VAMs	140	591	366	269

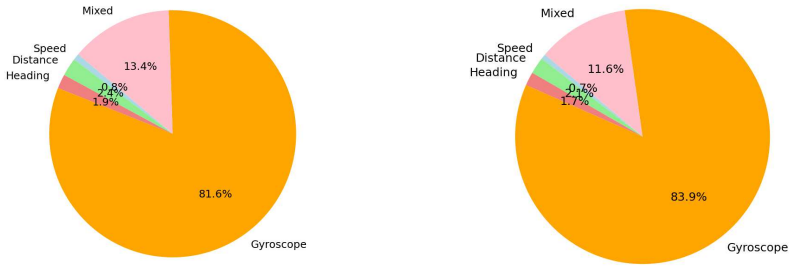

(a) Experienced (b) Inexperienced

Fig. 2.11: VAM trigger distribution with additional stability trigger [1].

2.9 Main findings

In this section, we report the evidence gathered from the two study pillars and highlight actionable takeaways. We first summarise the outcomes from the user-centric study, followed by the results of the communication–centric study, and conclude with a discussion on the implications for CCAM, along with limitations and future outlook.

User–centric study

- Motor assistance improves control around the target: *Turbo* exhibits the tightest dispersion and the smallest centred deviations, while *Eco* remains more variable (Fig. 2.3).
- Cadence stays close to 60 rpm across modes, indicating that assistance primarily covers torque demand without altering cycling mechanics.
- Participants systematically underestimate their speed; the bias is smaller in *Turbo* than in *Eco* (Fig. 2.4).
- Workload decreases with assistance: NASA–TLX is highest in *None* and lower in *Eco*/*Turbo*, driven by Physical Demand, Effort and Frustration (Fig. 2.5).
- Motor disengagement is most noticeable in *Turbo* and several riders report short-lived difficulty accelerating after cut-off, suggesting the need for smoother re–engagement.
- Acceptance is highest for predictable and non-intrusive cooperative features (e.g., warnings), while context-mismatched cues (e.g., during phone calls/headphones) are less appealing; privacy attitudes are polarized, motivating privacy-by-design controls (Fig. 2.6).

Communication–centric study

- With ETSI triggers only, T_{GenVAM} remains within the prescribed [100 ms, 5 s] interval along the city route but exhibits high variability. Experienced riders generate fewer VAMs (140) than inexperienced riders (151) due to smoother speed profiles (Fig. 2.8).
- In GNSS-degraded segments, heading-based triggers become marginal, while position changes dominate (Fig. 2.9); local wheel speed sensing sustains VAM generation even when GNSS is unavailable. The measured duty cycle remains very low ($\approx 0.07\%$), well below the 3% limit.
- Introducing a gyroscope–based stability trigger (inclination on the y-axis) markedly reduces T_{GenVAM} while preserving compliance: for the experienced rider, T_{GenVAM} decreases from 1428 ms (ETSI only) to 338 ms at $\Delta_g = 4°$, with stability responsible for 82% of triggers; similar figures hold for the inexperienced rider (341 ms). The duty cycle stays low ($\approx 0.3\%$) and message overhead is minimal (+6 Bytes per message), for an overall generation rate of about 2.3 kbps per cyclist (Table 2.2, Fig. 2.11).

Implications for CCAM

- **Assistance shaping:** a *Turbo* like torque profile improves pace regularity and reduces workload without masking speed perception; cooperative speed support should ensure smooth cut–off and re–engagement.
- **Sensing robustness:** fusing wheel speed with GNSS mitigates VAM blind spots in NLOS; stability sensing offers a GNSS independent trigger that captures maneuvers and near stops.
- **Standards evolution:** adding a gyroscope-based stability trigger to ETSI VAM generation, with thresholds in the 4–10° range as a practical compromise between reactivity and channel load, improves VAM timeliness while keeping the channel duty cycle comfortably within limits. Table 2.2 suggests a simple operative guideline for selecting the gyroscope-based threshold. Lower thresholds (e.g., $\Delta g = 4°$) maximize reactivity and message freshness but increase the number of stability-triggered VAMs, which may be undesirable in dense scenarios. Higher thresholds (e.g., $\Delta g = 10°$) reduce the trigger rate and the resulting overhead, at the cost of less frequent updates during moderate rider motions. Intermediate values (e.g., 7°) represent a practical compromise when the goal is to improve awareness timeliness while keeping the duty cycle low.

Limitations and outlook

The study involves a moderate sample with a male prevalence, one motor platform, and two assistance profiles; large–scale traffic interactions and real radio performance were not considered.

The user-centric evidence is based on a sample of 15 participants. The findings should be interpreted as exploratory and may not fully generalize to broader micromobility populations and contexts. Larger and different trials are needed to strengthen external validity and to quantify interaction effects more robustly.

3

Sensing, Perception and Communication Platform for CCAM

This chapter presents an integrated on-board platform that combines low-cost sensing of the vehicle surroundings, embedded perception of the road environment and driver status, and heterogeneous short-range/long-range communication technologies for data dissemination to support CCAM services. We first discuss the conceived platform's design and then report experimental results from testbeds where the prototype collects vehicular and environmental data, runs enhanced and specialized detection models on embedded hardware, and promptly delivers messages towards nearby and remote users through the configuration of proper workflows and protocols. Finally, we discuss the main findings and implications for future vehicular platforms, focusing on performance, scalability, and deployability.

3.1 Background and motivations

CCAM services build on CAVs retrieve, process, store, and share data about their status and surrounding environment to enable advanced driving applications with different autonomy levels, according to the SAE J3016 taxonomy [65]. Modern vehicles utilize on-board sensing technologies including cameras, GNSS, radar, and LiDAR, alongside embedded ML models for perception and decision-making processes. Furthermore, they exchange information with nearby entities and remote systems via V2X communications [66, 67, 68]. This integration of communication, computing and automation opens the way for applications ranging from cooperative perception and maneuvering to environmental sensing and remote diagnostics [69, 70].

However, relying only on on-board sensors limits coverage and detection accuracy, especially in complex traffic scenarios with VRUs and heterogeneous objects [67, 66]. To improve situational awareness, vehicles are increasingly expected to act as probe nodes, combining kinematic data from the CAN bus with supplementary environmental measurements, such as air quality and noise levels. These data un-

dergo preliminary processing and inference, and standardized information is then disseminated to nearby vehicles and infrastructure via both short- and long-range interfaces [71, 72]. At the same time, vehicular data are considered a valuable asset for many stakeholders (e.g., service providers, public authorities), which calls for interoperable and efficient data models and interfaces [73].

From a connectivity perspective, different V2X technologies coexist and are expected to interoperate. IEEE 802.11p/bd and 3GPP C-V2X and NR-V2X support short range communications among vehicles, RSUs and VRUs [74, 68], while cellular 4G/5G links enable V2N and more generally Vehicle-to-Network-to-Everything (V2N2X) interactions with remote entities such as mobility service providers and infrastructure operators [75].

Today, commercial On Board Units (OBUs) and RSUs available on the market (e.g., Cohda, Commsignia, Movyon, Acetel) typically feature closed-source ITS stacks, limited radio configurations and high costs [76]. Academic open-source prototypes address part of these limitations, but they predominantly focus on single-technology IEEE 802.11p platforms for ITS message dissemination [77, 78, 79, 80], often overlooking the need for tight integration with long-range 5G connectivity and application-level data models.

In parallel, the authors in [81] highlight that next generation vehicles will expose multiple heterogeneous radio interfaces (5G V2I, sidelink V2V, Wi-Fi, legacy cellular) and that exploiting them jointly is crucial to achieve low-latency, reliable synchronization between on-board processing and external entities. The works in [81, 82] emphasize the need for Single Board Computer (SBC)-based gateways capable of running containerized services for data collection, multi-interface connectivity and AI-based perception on commodity hardware.

Standard ETSI CAMs and DENMs allow vehicles and infrastructure to share kinematics and event-driven alerts in real time. However, these messages are primarily designed for local broadcast in the vehicular domain; when information has to reach remote backends and third-party applications, application layer solutions are typically adopted, often relying on lightweight publish/subscribe protocols such as MQTT running over cellular links [83], [84], [85].

These considerations motivate the design of an integrated and open-source on-board platform for CCAM that:

- collects data from heterogeneous on-board and external sensors for local processing and/or sharing with other entities;
- is able to exploit short-range IEEE 802.11p/ITS-G5 V2X communication technologies and long-range 5G Uu connectivity (jointly, if needed) to support both cooperative interactions with entities in proximity and remote data exchange;

- manages standard-compliant ITS messages jointly with lightweight application layer protocols to ensure interoperability and an efficient use of radio resources, while satisfying application requirements;
- embeds a perception module able to run real-time object detection models on off-the-shelf embedded hardware, leveraging context-aware and fine-tuned Convolutional Neural Network (CNN)-based architectures tailored to specific driving scenarios such as highway and rural roads and to driver monitoring pipelines (e.g., drowsiness and attention detection).

Achievements related to the aforementioned topics are, respectively, discussed in the following papers [86, 87, 88, 89, 4].

The remainder of this chapter is organized as follows. Section 3.2 describes the platform architecture and implementation. Section 3.3 presents the improvements introduced in the perception module, addressing context-aware object detection and driver monitoring, with relevant results. Section 3.4 reports the experimental set-up and results on data collection and dissemination over short-range and long-range communication interfaces. Finally, Section 3.5 summarizes the main findings.

3.2 Platform description

The experimental platform adopts a modular architecture that integrates sensing, perception, and communication functionalities into a single on-board system. As illustrated in Fig. 3.1, the platform combines a heterogeneous sensing suite, an embedded perception module, and a multi-interface, multi-protocol communication system, enabling the real-time generation and dissemination of messages.

The main functionalities of the platform are described in the following.

3.2.1 Sensing module

The sensing module is designed to acquire, timestamp, and expose a coherent stream of geo-referenced measurements that fuses environmental variables (air quality indicators and noise levels) with vehicle kinematics under portability, reproducibility, and low cost constraints.

In our platform, this is achieved through a physically separate SU based on an open-source ESP32 microcontroller (Fig. 3.2), which can be mounted externally on the vehicle to preserve air-quality measurement fidelity and reduce wiring.

The hardware configuration includes a Sensirion SEN55 for particulate matter ($PM_{1/2.5/4/10}$), Volatile Organic Compounds (VOC) index, temperature, and humidity; an Adafruit GNSS Ultimate Breakout V3 for geo-referencing; an MQ-7 sensor for

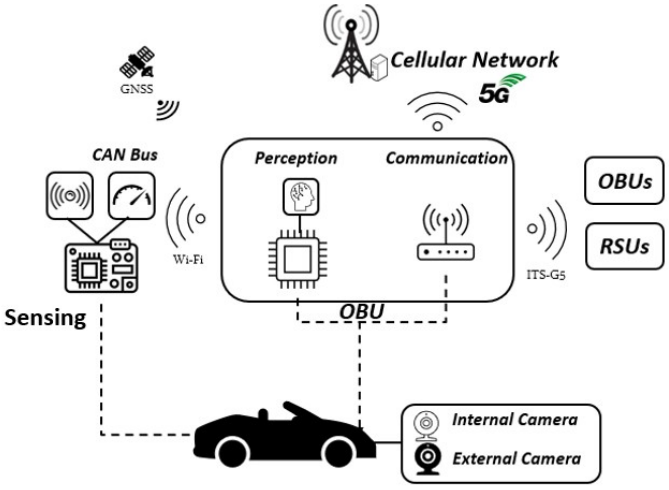

Fig. 3.1: Overview of the integrated sensing, perception, and communication platform deployed on the test vehicle.

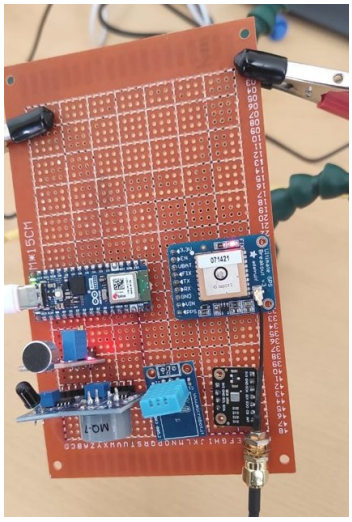

Fig. 3.2: External Sensing Unit (SU).

Carbon Monoxide (CO) operated with alternating heater phases; and a microphone-based noise sensor for ambient sound levels. In the PoC configuration, the GNSS module operates at 2 Hz while environmental sensors sample at 1 Hz. Timestamps are attached at acquisition and later reconciled on the on-board computing module to balance temporal accuracy, network footprint, and the SU energy budget.

The SU forwards lightweight JavaScript Object Notation (JSON) payloads over Wi-Fi using User Datagram Protocol (UDP) to the on-board computing module, while local SD storage enables store-and-forward operation for resilience in case of temporary connectivity loss.

Additionally, vehicle telemetry is collected through the On-Board Diagnostics (OBD) port via the CAN bus and integrated by the on-board compute using *Sock-*

etCAN and *python-can*. These data are temporally aligned with GNSS and environmental measurements to ensure coherent and synchronous streams before dissemination.

3.2.2 Perception module

The perception module is responsible for transforming raw visual input into high-level semantic information that can be used by CCAM applications and by the communication module for alert generation and dissemination.

At the hardware level, the module employs two embedded cameras connected to the on-board compute, implemented on a Raspberry Pi 5 that acts as the central OBU of the platform. The outward-facing camera is mounted on the front glass of the vehicle to cover the road ahead and is time-aligned with the GNSS reference. A second inward-facing camera observes the vehicle cabin and the driver.

Two services can run in parallel:

- **External detection**: object detection on the outward-facing stream is performed using a lightweight one-stage detector based on YOLOv11 [90], selected to balance accuracy and computational cost on embedded hardware.
- **Driver monitoring** : the inward-facing camera enables driver state assessment (e.g., fatigue/attention proxies) through lightweight vision-based pipelines that can be personalized via fine-tuning [86].

For both pipelines, the perception module performs video pre-processing, inference, and event extraction, producing structured outputs, such as lists of detected objects or driver state estimates. Each event is formatted as a compact JSON message containing timestamp, frame identifier, GNSS coordinates, and class/confidence metadata. These messages are exposed to the communication subsystem through RESTful APIs, which allow the communication module to access high-level perception outputs and to trigger the transmission of safety-critical messages.

3.2.3 Communication module

The communication module ensures seamless, low-latency, and standard-compliant data exchange between the on-board sensing and perception modules and external CCAM entities. Its design follows a multi-interface and multi-protocol approach, integrating short-range communications and long-range connectivity and different messaging protocols.

Short-range communications rely on the ITS-G5 (IEEE 802.11p-compliant) communication technology and enable the exchange of ETSI ITS messages.

Complementing ITS-G5, long-range communications rely on a 5G gateway that provides connectivity to the cellular infrastructure. Hence, the gateway also enables V2N2X data dissemination towards remote stakeholders and supports application-layer data delivery (e.g., via MQTT) over 5G.

The specific hardware components, dissemination workflows and the related experimental evaluation are described in the following.

3.3 Improving perception capabilities

The perception module must deliver accurate, timely and context-aware detections while respecting the computing and energy constraints of embedded on-board hardware. Hence, it represents a crucial component of the conceived architecture which deserves proper improvements. In particular, we extend the perception module along two main directions:

- **Context-aware, energy-efficient road object detection**: we specialize lightweight detectors for different driving contexts (e.g., highway vs. rural) and enable dynamic model selection at runtime.
- **Driver state monitoring with lightweight personalization**: we evaluate a compact vision-based driver monitoring pipeline and quantify the benefits of driver-specific fine-tuning.

Context-awareness and personalization are achieved through fine-tuning techniques, properly applied to achieve the trade-off between accuracy and complexity, as discussed in the following.

3.3.1 Context-aware energy-efficient on-board object detection

Motivations

Deploying on-board perception in CCAM requires operating under strict and often conflicting constraints. First, the detector must sustain real-time throughput on embedded hardware where compute, memory, and thermal headroom are limited. Second, energy consumption is a first order design variable: prolonged operation on SBC class devices must avoid thermal throttling and excessive power draw, especially when perception runs continuously alongside sensing and communications.

A single monolithic detector trained on generic datasets is typically suboptimal across heterogeneous scenes. The mix of relevant classes and their appearance differs between road contexts; consequently, compute is often wasted on irrelevant classes, while accuracy may degrade on safety-relevant road users. These limitations motivate a smart object detection strategy based on *(i)* context-aware fine-tuned models

specialized for specific deployment scenarios and *(ii)* an efficient inference runtime for resource-constrained devices.

Proposal and methodology

We adopt a workflow that starts from You Only Look Once (YOLO)v11 [90], a well-known one-stage convolutional object detector, and specializes it for in-vehicle operation through a reproducible PoC and a context-aware inference stack integrated with the sensing and communication modules. The prototype runs on a Raspberry Pi 5 with a Raspberry Camera V3 and is powered directly from the vehicle. Frames are captured and inferred on-board, with outputs aligned to the GNSS clock derived from the sensing unit.

The model is fine-tuned starting from pre-trained YOLO weights and leveraging Common Objects in Context (COCO) [91] as a baseline dataset, then specializing towards road relevant classes. To optimize the inference pipeline, we adopt a context-aware specialization strategy: two dedicated models are trained for *rural* and *highway* contexts, each focusing on the subset of classes most representative of its operational domain (Tab.3.1). Training performed using fine-tuning methodology preserves the generalization capabilities of the pre-trained backbone while specializing the detector to the operational domains (highway and rural) and to relevant classes emphasized by the proposed selection. This specialization reduces the effective label space and improves throughput, while increasing accuracy for objects in the target context.

Scenario	**COCO Classes**
General training	Bicycle, Bus, Car, Cat, Cow, Dog, Horse, Motorbike, Person, Sheep, Truck, Fire hydrant, Traffic light, Train
Highway	Bus, Car, Motorbike, Person, Truck
Rural	Bicycle, Bus, Car, Cat, Cow, Dog, Horse, Motorbike, Person, Sheep, Truck

Table 3.1: Selected classes for the considered driving scenarios.

This specialization reduces the label space and inference complexity, improving throughput and lowering energy consumption while improving accuracy relevant to each scenario.

Fine-tuning is applied to the lightweight YOLOv11(n)-*nano* configuration with the following setup: 200 epochs; initial learning rate $\texttt{lr}_0 = 10^{-4}$ decaying to a final $\texttt{lrf} = 10^{-2}$; training image size $\texttt{imgsz} = 640$; batch size = 16; dropout = 0.15. This approach adapts a pre-trained model to a specific operational domain, improving accuracy without training from scratch [92, 93].

To reduce training cost and the risk of overfitting on the target domain data, we adopt partial fine-tuning: early backbone stages are kept frozen, while only the higher-level layers are updated during training.

Legacy and fine-tuned models are exported to, and executed with, the Nihui Convolutional Neural Network (NCNN) inference framework [94] to accelerate inference on resource constrained devices. At runtime, the selection of the active detection model is driven by contextual cues from the sensing module (e.g., speed and location). When the operating context changes, the perception service switches to the most suitable optimized model, keeping the accuracy and efficiency trade-off balanced across heterogeneous scenarios with minimal overhead.

Results

The performances of the following solutions are compared:

- YOLOv11s (baseline, native runtime);
- YOLOv11n (lighter baseline, native runtime);
- YOLOv11s exported to NCNN;
- YOLOv11n exported to NCNN;
- YOLOv11n exported to NCNN and fine-tuned for the target context (proposed solution).

The *Mean Average Precision* mAP50–95 [95] is used as the main accuracy metric for object detection. It averages precision over *Intersection Over Union* (IoU) thresholds from 0.50 to 0.95 in steps of 0.05, thus reflecting the model's ability to localize and classify objects at different overlap levels.

As expected, accuracy decreases when moving from *YOLO v11s* to the lighter *YOLO v11n* (Table 3.2).

However, the fine-tuned *YOLOv11n-NCNN-FT* outperforms the more complex *YOLOv11s*: mAP50–95 increases by 39.9% in the highway scenario (from 0.464 to 0.649) and by 31.75% in the rural scenario (from 0.441 to 0.581). This shows that fine-tuning a lightweight model can recover, and even exceed, the effectiveness of larger, computation-heavier and more accurate models.

The confidence score [96] measures how sure the model is that a detection belongs to a class. Predictions below a threshold are filtered. Higher mean confidence at fixed threshold indicates more reliable detections. With the proposed solution, the average confidence goes from 0.59 to 0.72 in the highway case and from 0.54 to 0.58 in the rural case.

Runtime performance was measured on the device. The proposed model yields higher *Frames per Second (FPS)* processed and shorter inference times than the base-

lines (Table 3.2). In the highway scenario, FPS increases from 4.46 (*YOLOv11s*) to 43.3; in the rural scenario, from 6.57 to 38.05. Inference time drops from 224 ms to 27.9 ms (highway) and from 152.2 ms to 32.1 ms (rural). Notably, the fine-tuned model also outperforms *YOLOv11n-NCNN*, which is already optimized for efficiency.

These gains translate into lower on-device resource usage: the measured current decreases from 0.902 A to 0.742 A in highway and from 0.857 A to 0.703 A in rural.

Metrics	**v11s**		**v11n**		**v11s-NCNN**		**v11n-NCNN**		**v11n-NCNN-FT**	
	H	**R**	**H**	**R**	**H**	**R**	**H**	**R**	**H**	**R**
mAP50-95^{train}	0.464	0.441	0.363	0.358	0.464	0.441	0.363	0.358	0.649	0.581
Confidence score	0.59	0.54	0.55	0.51	0.59	0.54	0.55	0.51	0.72	0.58
FPS	4.46	6.57	10.3	14.6	16.05	17.14	35.9	35.8	43.3	38.05
Inference time [ms]	224	152.2	97.28	68.48	70.4	61.9	34.6	36.13	27.9	32.1
Current usage [A]	0.902	0.857	0.878	0.831	0.844	0.840	0.771	0.771	0.742	0.703

Table 3.2: Comparison of accuracy and efficiency across models and scenarios (H = Highway, R = Rural) [88].

3.3.2 Driver state monitoring with lightweight personalization

Motivations

Beyond environment perception, monitoring the driver state can provide additional safety value in CCAM scenarios, especially when human supervision or takeover readiness is relevant. As modern (and future) vehicles build increasingly upon on-board automation functions, driver's attention in controlling vehicle kinematics and trajectory tends to decrease, which can induce a progressive relaxation and reduced vigilance during the driving. Indeed, according to the SAE J3016 driving automation taxonomy (Fig. 3.3 [3]), a human driver (or, more precisely, a fallback-ready user) is still expected to resume control up to SAE Level 3 when the automated driving system requests an intervention, e.g., due to system limitations, unexpected events, or sensors failure. In this context, excessive relaxation can lead to critical situations such as drowsiness, distraction, or delayed reactions to re-take the control of the vehicle.

In practice, driver monitoring pipelines must remain lightweight, deployable on embedded hardware, and capable of detecting hazardous driver conditions in real-time. At the same time, different drivers' characteristics (e.g., facial traits, habitual expressions, eyewear, illumination) can significantly affect performance of a generic

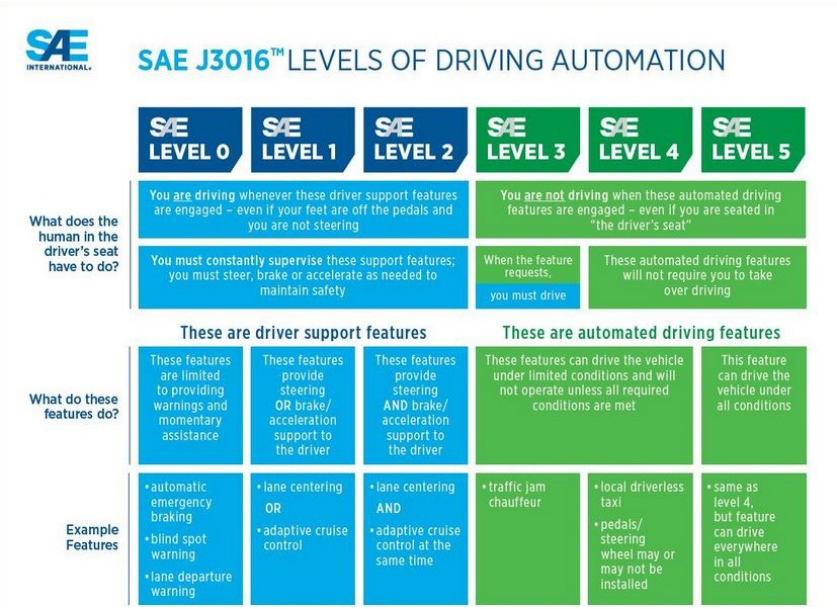

Fig. 3.3: SAE Automation levels [3].

model. This motivates lightweight personalization strategies, where a legacy model is adapted to a specific driver. This is possible using fine-tuning techniques which help to improve the performance of the detection system.

Proposal and methodology

We considered a vision-based drowsiness detection pipeline and evaluated a lightweight personalization strategy based on driver-specific fine-tuning. The task is formulated as a binary classification problem (*active* vs. *fatigued*). Fine tuning is applied on driver-specific images starting from three pre-trained CNN models (ResNet50 [97], InceptionV3 [98], and MobileNetV2 [99]), producing an updated model that can be deployed on the on-board perception module.

Results

Fine-tuning increases classification accuracy by about three percentage points (e.g., from $\approx 93\%$ to $\approx 96\%$ for MobileNetV2), while keeping model sizes in the range 14–98 MB and tuning times on the order of 8–16 minutes depending on the allocated CPU share (Tab. 3.3). These results suggest that driver-specific personalization can be performed prior to the trip, either at the edge or on an external compute node, and then deployed onboard, providing measurable gains with limited overhead.

ML model	Size [MB]	Accuracy before FT [%]	Accuracy after FT [%]
ResNet50	98	91	94
InceptionV3	92	92	95
MobileNetV2	14	93	96

Table 3.3: Model size and accuracy before and after fine-tuning (FT).

3.4 Data collection and dissemination: improvements and validation

The integrated platform was designed to enable the collection of heterogeneous data streams on board and disseminate them towards nearby road users, through short-range communications, and/or towards remote applications, via a cellular interface. For this purpose, in the following, details are provided about how different components have been properly interfaced and protocols configured for the aforementioned purpose.

Moreover, validation results conducted with real experiments on the road are also discussed, when considering short-range communications and 5G connectivity in a stand-alone manner, as well as their interplay for improved data dissemination.

3.4.1 Short-range communications

Experimental set-up

The platform enables the broadcasting of ETSI CAMs and DENMs over ITS-G5, carrying kinematics and perception-driven alerts, respectively.

Tests focused on the exchange of CAMs were conducted with two equipped private vehicles. Each car was fitted with an OBU, with the GNSS antenna mounted on the rooftop. The vehicles followed an urban path near the campus, passing multiple roundabouts, with speeds typically below 50 km/h for most of the trip (Fig. 3.4).

Each OBU was configured to generate CAMs according to ETSI rules [19]. Such rules are intended to reduce the number of messages exchanged over the wireless channel and to trigger their transmission only when changes occur in the mobility pattern of the vehicle. In particular, the trigger conditions are the following:

- the absolute difference between the current heading of the originating vehicle and the heading included in the CAM previously transmitted by the same vehicle exceeds 4°;
- the distance between the current position of the vehicle and the position included in the CAM previously transmitted by the same vehicle exceeds 4 m;

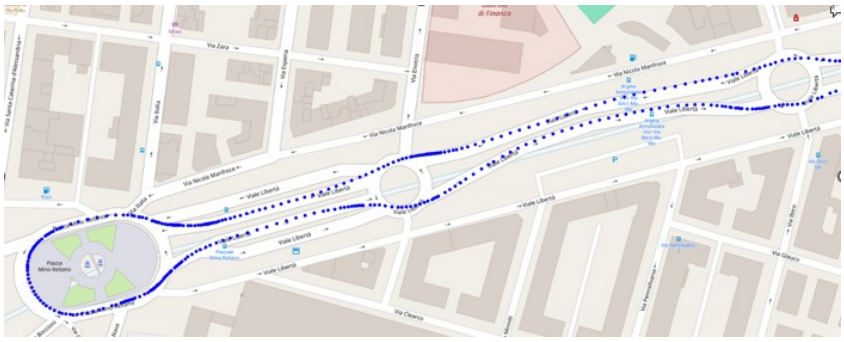

Fig. 3.4: Urban path followed by the test vehicles and position of the transmitting vehicle when CAMs are generated.

- the absolute difference between the current speed of the vehicle and the speed included in the CAM previously transmitted by the vehicle exceeds 0.5 m/s;
- time-based trigger: a CAM is also generated when the time elapsed since the previous CAM reaches the maximum CAM generation interval (i.e., at least one CAM shall be generated within the maximum interval even if none of the above dynamics thresholds is met).

In addition, CAM generation is subject to interval constraints: the time between two consecutive CAMs is bounded between a minimum and a maximum value (commonly $T_{\text{GenCamMin}}$ = 100 ms and $T_{\text{GenCamMax}}$ = 1000 ms), and the dynamics-based triggers are evaluated only if at least the minimum generation interval has elapsed since the last transmitted CAM. In ITS-G5 deployments, this minimum interval can be further constrained by DCC (i.e., an additional minimum interval may apply), effectively postponing CAM generation under channel congestion conditions [19].

The position retrieved from GNSS was inserted in the CAM fields, as visible from the logged messages (Fig. 3.5).

[2025-10-01 02:24:52.130] [cabs] camTransmission - {"header":{"protocolVersion":2,"messageID":2,"stationID":33},"cam":
{"generationDeltaTime":11771,"camParameters":{"basicContainer":{"stationType":5,"referencePosition":
{"latitude":381211326,"longitude":156654494,"positionConfidenceEllipse":
{"semiMajorConfidence":4095,"semiMinorConfidence":4095,"semiMajorOrientation":3601},"altitude":
{"altitudeValue":9297,"altitudeConfidence":"unavailable"}}},"highFrequencyContainer":{"basicVehicleContainerHighFrequency":{"heading":
{"headingValue":3457,"headingConfidence":127},"speed":{"speedValue":5,"speedConfidence":127},"driveDirection":"unavailable","vehicleLength":
{"vehicleLengthValue":1023,"vehicleLengthConfidenceIndication":"unavailable"},"vehicleWidth":62,"longitudinalAcceleration":
{"longitudinalAccelerationValue":161,"longitudinalAccelerationConfidence":102},"curvature":
{"curvatureValue":1023,"curvatureConfidence":"unavailable"},"curvatureCalculationMode":"unavailable","yawRate":
{"yawRateValue":32767,"yawRateConfidence":"unavailable"},"accelerationControl":"00"}},"customContainer":{"generationTimestamp":6863336310911}}}}
[2025-10-01 02:24:53.154] [cabs] camTransmission - {"header":{"protocolVersion":2,"messageID":2,"stationID":33},"cam":
{"generationDeltaTime":11771,"camParameters":{"basicContainer":{"stationType":5,"referencePosition":
{"latitude":381211326,"longitude":156654494,"positionConfidenceEllipse":

Fig. 3.5: Example of a log of a generated CAM with the relevant fields.

Results

Fig. 3.4 shows the positions of the vehicle when CAMs are transmitted, blue markers.

It can be observed that they become denser near roundabouts. This is due to the heading-based trigger: a CAM is generated whenever the absolute difference between the current heading and the heading in the last transmitted CAM exceeds 4°. Fig. 3.6 shows the speed profile sampled from the CAMs, while Fig. 3.7 reports the IGG, which represents the time between two consecutive CAMs transmitted from the same vehicle.

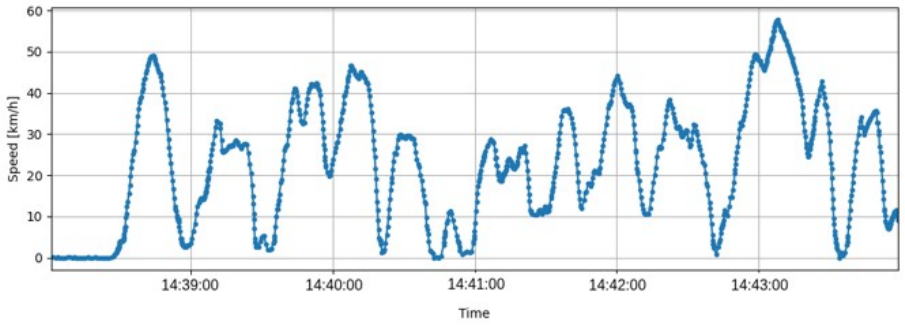

Fig. 3.6: Speed profile of the transmitting vehicle, with speed values sampled from CAMs.

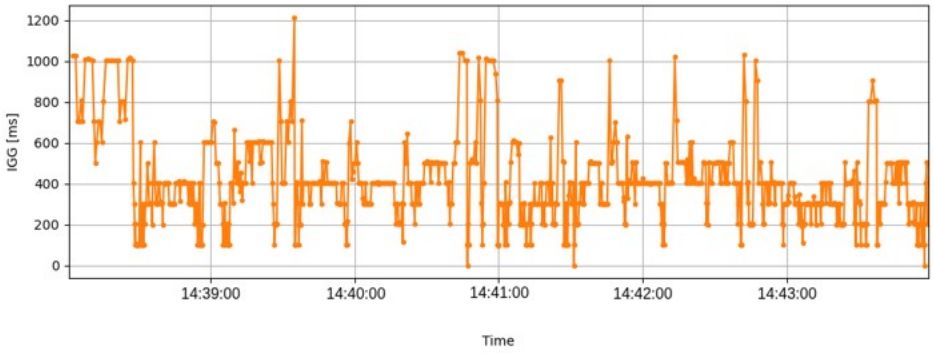

Fig. 3.7: IGG of transmitted CAMs.

By comparing the two plots, it emerges that the IGG decreases whenever there is a sudden variation in speed, consistent with the speed-based trigger condition (difference greater than 0.5 m/s between the current speed and the last CAM). In straight segments with almost constant speed the IGG grows, leading to sparser CAMs and thus, reduced channel load.

Additional interoperability tests were conducted in the campus premises through the exchange of CAMs between an equipped vehicle and an RSU, again over ITS-G5, as shown in Fig.3.8. In this testbed, we used the proposed platform, a commercial Cohda Wireless OBU and, a Movyon Electronics RSU. In all cases, CAMs were correctly received and decoded across heterogeneous devices, confirming cross-vendor interoperability.

Fig. 3.8: Testbed in the campus premises: OBU mounted on board (left) and RSU covering the area where the equipped vehicle (the one within the red circle) moves (right).

Finally, we complemented the functional tests with a packet-level validation by capturing over-the-air traffic with Wireshark [100]. While CAMs validate periodic awareness dissemination, we also inspected event-driven messages DENMs to verify that messages are correctly encoded and dissected. Fig. 3.9 reports an excerpt of the captured trace, where DENMs are broadcast repeatedly with a constant frame size (139 bytes in our capture), and shows the decoded header and payload fields for a representative DENM (e.g., protocolVersion=2, messageId=denm, stationId, eventPosition, validityDuration, and eventType). These packet-level checks strengthen the evidence of standards compliance and interoperability, and provide the basis for the multi-protocol forwarding workflows discussed next.

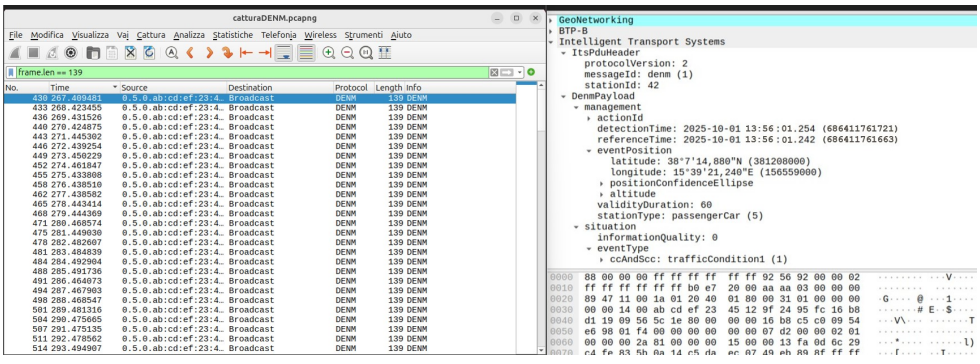

Fig. 3.9: Packet-level validation: Wireshark trace excerpt showing repeated DENM broadcast frames (left) and Decoded DENM fields (right).

3.4.2 5G-assisted communication scenarios

The dissemination of safety-critical vehicular messages is typically based on short-range V2V and V2I interactions.

Thanks to onboard sensors and cameras, vehicles can detect hazardous events and send DENMs accordingly. According to ETSI specifications, a DENM is an event-driven message intended to notify road entities about time-bounded hazardous conditions and traffic-relevant events, such as traffic congestion and abnormal traffic conditions, accidents, road works, obstacles on the road, or adverse weather conditions. Differently from CAMs, which primarily convey periodic kinematics, DENMs include an *actionId* to uniquely identify and manage the event lifecycle, a *management* container carrying event timestamps and geolocation, and a *situation* container describing the event class through standardized *causeCode*/*subCauseCode* fields and a qualitative confidence indicator (*informationQuality*). Fig. 3.9 shows an example decoded from our trace, where the message reports a traffic condition notification (eventType *trafficCondition*), together with the corresponding header and management fields. Typical DENM triggering events include, among others, traffic accidents, stationary vehicles, road works, obstacles on the road, adverse weather conditions, and other abnormal traffic or hazard situations [101].

However, in some circumstances, information carried in DENMs may need to reach remote entities beyond those in the immediate vicinity of the vehicles for proper event handling, e.g., the road operator. Moreover, in the near future, not all vehicles are expected to be V2V-capable.

The V2N2X paradigm, promoted by the 5GAA [75], provides a valuable means for fast, reliable, and efficient data exchange among key stakeholders. Despite its potential, to the best of our knowledge, only a few works have focused on this paradigm, mainly through theoretical studies [102], [83].

The study reported in this thesis aims to complement the state-of-the-art by exploring the capability of 5G-based V2N2X connectivity to timely disseminate traffic information through real-world experimental testbeds.

A set of different scenarios have been developed and assessed, as reported in Fig. 3.12, Fig. 3.13, Fig. 3.14.

Hybrid scenario. The first scenario (Fig. 3.12) involves a vehicle equipped with an ITS-G5 capable OBU that detects an event, triggering the transmission of a DENM. The device exposes RESTful APIs to configure transmission parameters and trigger messages for DENMs. These messages are then broadcast and received by nearby entities, including other OBUs and an RSU.

In our implementation, the RSU is connected through Ethernet to a 5G gateway, realized through a commercial router (Teltonika RUTX50). The router operates as

a User Equipment (UE) toward the private 5G-SA network and hosts a SIM provisioned for experimental access, enabling authenticated attachment to the network. It runs RutOS [103], which is based on the OpenWrt operating system [104].

We extended the legacy version of the operating system with open-source modules to enable data sharing and processing via the MQTT protocol. Indeed, the RSU captures DENMs from the nearby vehicle transmitting over ITS-G5, extracts the main event fields, and forwards them as a JSON-formatted MQTT PUBLISH message.

Specifically, the RSU decodes the received DENM and extracts a compact set of event descriptors, including the event position (latitude and longitude), the event timestamp (reference time of detection for the event), the event classification using the *causeCode* and *subCauseCode* and the event identifier (filled with originated station ID and the message sequence number). When explicitly using *causeCode*/*subCauseCode*, we follow the ETSI assignment. For instance, ETSI standard [101] assigns causeCode=2 to Accident with subCauseCode=8 for assistance requested (e-call); causeCode=3 to Roadworks with subCauseCode=4 for short-term stationary roadWorks; and causeCode=94 to Stationary vehicle with subCauseCode=2 for vehicle breakdown.

Fig. 3.10 also reports the API and the interface used to inject a DENM with controlled repetition parameters (e.g., repetition interval), which was useful to validate both short-range dissemination and the subsequent 5G/MQTT forwarding pipeline.

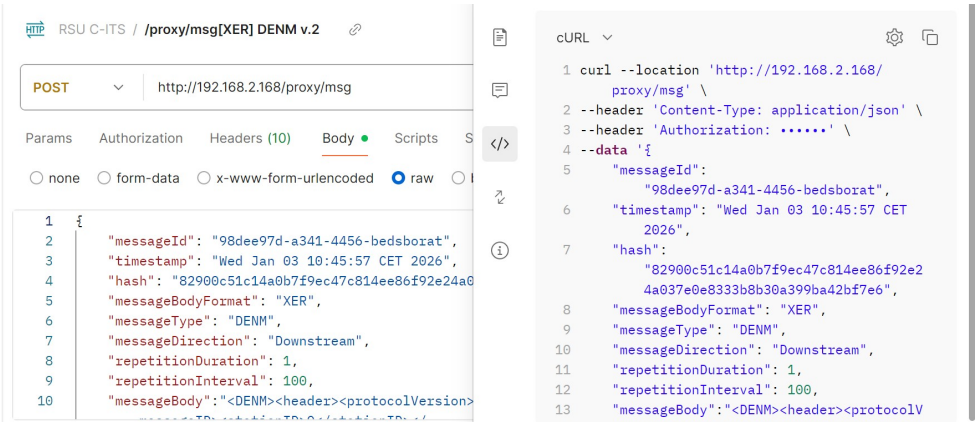

Fig. 3.10: Message injection workflow for DENM using the dedicated API.

The extracted information is then forwarded as a JSON formatted MQTT PUBLISH message. An example of the resulting MQTT payload is: *"stationId":42, "eventType":"trafficCondition","causeCode":26, "subCauseCode":8,"validityDuration":60, "eventPosition":"lat":381208000, "lon":156559000,"informationQuality":0.*

A private 5G-SA network is deployed on campus (Fig. 3.11) and operates on licensed spectrum in band n77 (centered around 3.99 GHz), made available for experimental use. The deployment implements a complete end-to-end 5G system, including both the Radio Access Network (RAN) and the 5G Core. On the access side, the RAN is provided by a gNodeB (gNB) offering radio coverage to the experimental area. On the core side, the 5G Core includes the main control-plane and user-plane functions required for registration and session establishment, such as the Access and Mobility Management Function (AMF), the Session Management Function (SMF), the User Plane Function (UPF). Authentication and subscription management are supported through the Authentication Server Function (AUSF) and the Unified Data Management (UDM).

Access to the private 5G-SA network is provided through commercial 5G routers that act as UEs. Each router hosts a SIM provisioned for the private network, which provides the subscription credentials used during network registration. In addition, the terminal identity International Mobile Equipment Identity (IMEI) can be registered and recognized by the network to explicitly authorize experimental devices and prevent unauthorized attachment.

An edge server, co-located with the gNB and the core network elements (Fig. 3.11), runs a containerized instance of Mosquitto [105] acting as an MQTT broker. This edge placement supports low-latency data dissemination workflows and enables local processing of road-related events close to the radio access.

3.4.3 5G-based long-range communication scenarios

In the first scenario, as previously described, the subscriber is assumed to be deployed on the edge server, subscribing to specific road-related topics.

Subscriber at the edge. In the second scenario (Fig. 3.13), the vehicle is not V2V-capable. Instead, it is equipped with a 5G gateway, enabling direct communication via 5G. It publishes a DENM-like message to notify remote entities of a dangerous condition detected on the road. The broker at the edge forwards the PUBLISH message to the co-located subscriber.

The subscriber could be, for example, a road operator, a traffic management application, an emergency dispatch center, an analytics service at the network edge, or a non-V2X-capable vehicle receiving warnings through 5G connectivity.

Subscriber on the road. The last scenario involves the dissemination of DENM-like messages among non-V2V-capable vehicles via 5G, as shown in Fig. 3.14. One vehicle acts as the publisher, while the other serves as the subscriber.

Fig. 3.11: 5G deployment in the campus: gNB, core network and edge server (left) and antenna (right).

Results

After proving the implementation viability of the conceived platform and validating the integration of its hardware and software components, we also aim to provide numerical results to showcase the performance of the different data dissemination solutions.

Hybrid scenario. We derive the end-to-end dissemination delay in the first scenario, the hybrid one. To this purpose, three delay contributions have been measured:

- the *time-to-802.11p*, i.e., the delay from data availability at the OBU to the completion of the DENM transmission over ITS-G5;
- the *time-before-5G*, i.e., the time to collect data at the RSU and generate the MQTT PUBLISH message;
- the one-way *5G transmission time* from the gateway to the MQTT subscriber at the edge, which is assumed to be idle.

The results reveal average values for the three different delay contributions of about 8.4 ms, 2.4 ms and 11.5 ms (Tab. 3.4), respectively, confirming that the multi-interface, multi-protocol integration introduces only a few milliseconds of additional processing delay. In the specific deployment, the private 5G network itself contributes a small and stable one-way latency.

5G-assisted scenarios. In the following, results are reported which have been derived for the scenarios where 5G connectivity is only available. Two types of subscribers are considered: *(i)* a subscriber at the 5G edge (co-located with the broker) and *(ii)* a subscriber hosted on another 5G-capable vehicle under the same gNB

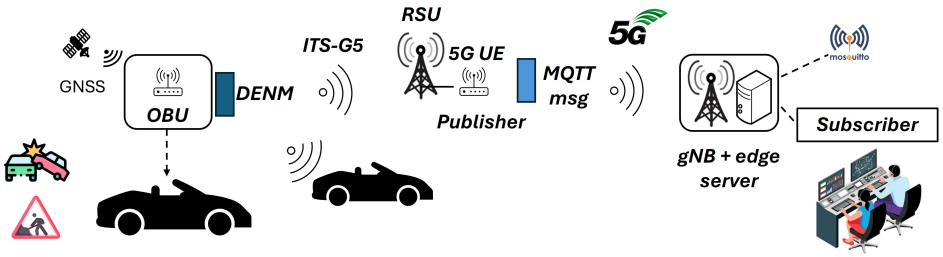

Fig. 3.12: ETSI DENM dissemination to road entities in proximity through ITS-G5 and to remote entities through 5G.

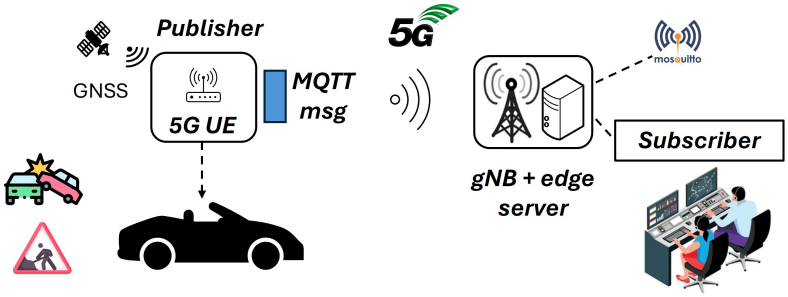

Fig. 3.13: DENM-like message dissemination to a subscriber at the edge through 5G connectivity.

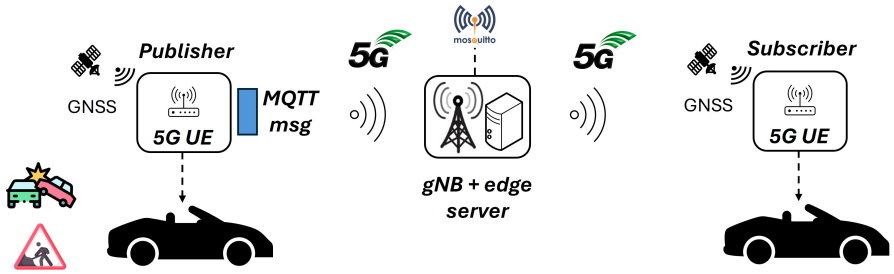

Fig. 3.14: DENM-like message dissemination among vehicles through 5G connectivity.

Time-to-802.11p	Time-before-5G	5G transmission time
8.4 ms	2.4 ms	11.5 ms

Table 3.4: Delay values in the hybrid scenario.

coverage. For each case, two broker conditions were analysed: an *idle* broker (only serving the target flow) and a *loaded* broker (ten additional clients continuously publishing messages of the same size).

The measured end-to-end delay, Fig. 3.15, from event detection at the source vehicle to notification at the subscriber, remained below 100 ms in all scenarios,

with the lowest values obtained for the edge subscriber with idle broker, and higher delays when the broker was loaded and/or the subscriber was on-board another vehicle. These results indicate that, even when accounting for broker load and an additional radio hop, 5G-based dissemination can satisfy the latency requirements of time-critical warning services, and offer a viable path to reach non-V2V-capable vehicles and remote stakeholders beyond the local ITS-G5 coverage.

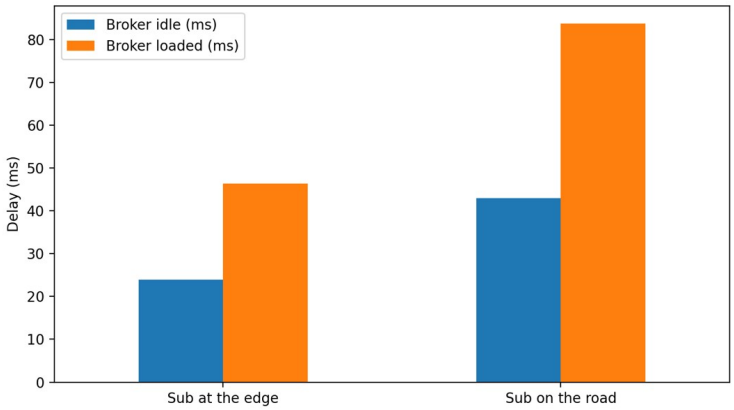

Fig. 3.15: Dissemination delay of the DENM-like message [4].

3.5 Main findings

The chapter presented an integrated sensing, perception, and communication platform for CCAM, along with a PoC evaluation of its data collection and dissemination capabilities.

The different modules are not designed as stand-alone components but are synergistically combined to enhance overall capabilities.

From a perception viewpoint, the platform demonstrates that context-aware fine-tuning of lightweight models can substantially improve accuracy and efficiency on embedded hardware. The proposed YOLOv11n NCNN Fine-Tuning (FT) configuration outperforms the heavier YOLOv11s baseline in terms of both mAP and runtime metrics. Frame rate increases by up to an order of magnitude, inference time is reduced to just a few tens of milliseconds, and current draw decreases, enabling continuous operation on a Raspberry Pi 5 without thermal throttling.

Dynamic selection of scenario-specific models (highway vs. rural) further improves robustness for safety-critical classes while keeping computational costs under control.

In addition, driver monitoring benefits from driver-specific fine-tuning, yielding measurable accuracy gains with a compact model size and minimal fine-tuning time.

On the data collection side, the sensing module shows that low-cost, modular hardware (ESP32-based external unit plus CAN-based vehicle telemetry) can be combined into a coherent stream of geo-referenced measurements.

Regarding short-range communications, trials confirm that ETSI CAM generation rules adapt the message frequency to vehicle dynamics: CAM IGG decreases around roundabouts and during speed changes, while it stretches on straight segments with nearly constant speed, naturally limiting channel load. The experiments, which also included DENMs, further reveal interoperability among vehicular devices from different vendors.

The analysis extends beyond short-range message dissemination, considering the viability of using 5G to replace V2V communications when not available, and to broaden the scope of message dissemination to entities not in close proximity. The measured delays for DENM-like message dissemination remain within the bounds required for time-critical warning services, staying below 100 ms under all the considered settings.

Besides the position of the subscriber, whether at the edge or onboard a vehicle, the analysis also reveals that the load on the MQTT broker significantly affects the results, highlighting the importance of its proper selection.

These findings support the use of the proposed platform as a basis for future large-scale trials and the integration of advanced cooperative applications, such as environmental monitoring, cooperative perception by also considering the exchange of CPMs, and network-assisted driver support.

4

Context-aware Bandwidth Adaptation for Video Stream of Teleoperated Fleets of Vehicles

This chapter investigates how to improve the network efficiency of ToD in fleet scenarios where multiple tele-operated vehicles share portions of their route and compete for the same uplink resources. The proposed solution adapts the quality of multi-camera video streams based on locally available context, namely inter-vehicle proximity and on-board detection of surrounding objects, and is integrated with an edge-assisted platoon management workflow that enables vehicles to dynamically form a platoon and safely transition follower vehicles between teleoperation and cooperative longitudinal control. The approach is evaluated in Simulation of Urban MObility (SUMO) under different traffic load conditions and compared against fixed-bandwidth streaming baselines.

4.1 Introduction

ToD is among the major use cases in the CCAM ecosystem, as it can complement fully autonomous driving and support vehicles facing challenging, dangerous, or unexpected traffic circumstances. In ToD, a remote operator supervises or directly controls the vehicle over a continuous bidirectional communication loop [106]. The scalability of ToD deployments is primarily constrained by the uplink footprint of multi-camera streaming and by the need to preserve low and stable end-to-end latency to guarantee safe maneuvering and proper situational awareness to the human-in-the-loop remotely controlling the vehicle [106, 107]. Inefficiencies in resource allocation translate directly into scalability limits at radio access and system level. For this reason, ToD is adopted here as a representative high-demand service to design and validate context-aware mechanisms that reduce network resource usage while preserving operator situational awareness in the most critical conditions.

In this chapter, we consider a fleet of ToD vehicles tele-operated through V2N connectivity and sharing parts of the trip (e.g., logistics operations). Our goal is to reduce the aggregate uplink demand without compromising the remote operator's situational awareness in critical conditions. To this end, we propose CABA, a

context-aware bandwidth adaptation method that jointly exploits (i) the proximity of adjacent ToD vehicles and (ii) the local detection of nearby objects. The approach is integrated with an edge-assisted platoon management function that enables dynamic platoon formation: the lead vehicle remains under teleoperation, while follower vehicles can switch their longitudinal control to a cooperative mode governed by CACC, supported by V2V message exchanges.

The main original contributions of this chapter are:

- We propose CABA, a context-aware multi-camera streaming scheme that leverages inter-vehicle proximity and on-board object detection to reduce uplink bandwidth requirements while preserving the remote operator's understanding of the scene.
- We couple CABA with an edge-assisted platoon management mechanism that enables dynamic platoon formation for V2V-equipped ToD vehicles and supports safe transitions between teleoperation and CACC.
- We evaluate the proposal using SUMO [108], comparing CABA against legacy fixed-bandwidth streaming under different traffic load settings.

The remainder of this chapter is organized as follows. Section 4.2 introduces the ToD service model and summarizes the main communication requirements that motivate bandwidth-efficient designs in the literature. Section 4.3 describes the proposed framework, detailing the context-aware bandwidth adaptation logic and the edge-assisted platoon management and control mechanisms. Section 4.4 presents the simulation methodology and discusses the main performance results. Section 4.5 reports a preliminary PoC implementation over a real 5G-SA testbed, aimed at validating the feasibility of an end-to-end teleoperation pipeline and at providing practical bandwidth measurements. Finally, Section 4.6 concludes the chapter and outlines future research directions.

4.2 Background and motivations

4.2.1 ToD: an overview

ToD realizes a remote-driving control loop in which a human operator, located in a control center, supervises and/or directly controls a vehicle through bidirectional communications. The vehicle continuously provides the operator with a time-critical representation of the surrounding environment, typically via multiple camera streams, when available, and auxiliary telemetry, while the operator sends control commands to actuate steering, throttle, and braking. As a result, ToD exhibits

a strongly asymmetric traffic profile: a sustained, bandwidth-intensive uplink dominated by multi-camera video, and a comparatively lightweight yet safety-critical downlink for command delivery. Fig. 4.1 illustrates this interaction.

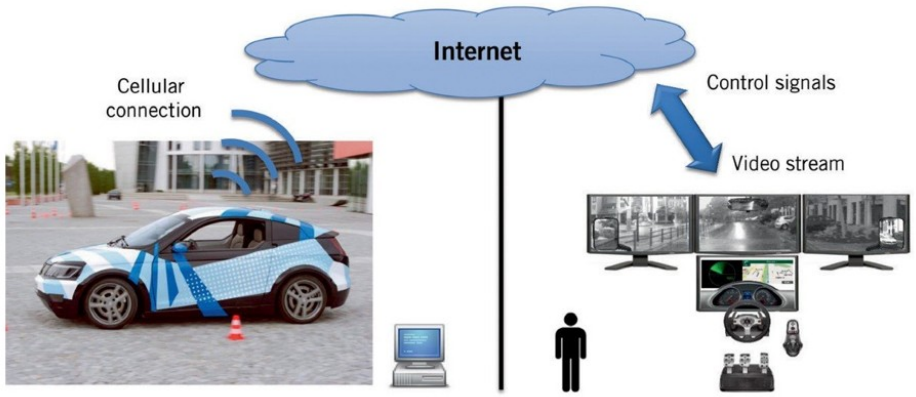

Fig. 4.1: Conceptual view of ToD: uplink video streaming from the vehicle to the remote operator and downlink control signals from the operator to the vehicle.

The 5GAA technical report on ToD system requirements and architecture [109], provides indicative service-level values for Direct Control operation, assuming a reference configuration where four cameras are streamed toward the remote operator. The resulting data-rate and reliability targets are summarized in Table 4.1. While the downlink command stream is modest in bitrate, it requires extremely high reliability; conversely, the uplink is dominated by multi-camera video and represents the main driver of spectrum and radio resource consumption.

In the 5GAA terminology, *Direct Control* refers to a teleoperation mode in which the remote operator continuously actuates steering, throttle, and braking in a closed-loop manner, relying on real-time visual feedback and low round-trip latency. Conversely, *Indirect Control* refers to a mode where the remote operator provides higher-level commands (e.g., target speed, maneuvers, or waypoints), while the vehicle executes such commands through on-board control functions. As a result, Indirect Control can tolerate looser latency budgets, whereas Direct Control requires tighter and more stable end-to-end delay to preserve controllability and safe maneuvering [109].

In addition to throughput and reliability, ToD is critically constrained by end-to-end latency and jitter. The 5GAA report discusses that latency budgets should be interpreted end-to-end and that the uplink contribution may be higher than the downlink one, since uplink includes video processing components such as encoding/decoding and rendering [109]. For Direct Control ToD (vehicle speed up to 50 km/h), an overall service-level round-trip latency on the order of 120 ms is reported, with an indicative split of 100 ms for the uplink path (vehicle to operator) and 20 ms for the

Table 4.1: Indicative communication requirements for ToD [109].

Direction	**Traffic component**	**Indicative rate**	**Notes / QoS target**
Uplink	Multi-camera video streaming	~32 Mbps	Progressive HD video, baseline assumption with H.264 compression and four cameras (one per side). Target reliability: ~99%.
Uplink	Optional interpreted object information	Up to ~4 Mbps	Additional sensor/perception data (e.g., interpreted objects) that can complement video. Target reliability: ~99%.
Uplink	Audio (optional)	~96 kbps	Environmental audio and/or conversational audio (may be embedded in the video stream).
Uplink	Vehicle information (telemetry/status)	~0.2 Mbps	Low-rate kinematic/status data supporting operator awareness and service monitoring.
Downlink	Control commands	Up to ~400 kbps	Up to 1000 bytes/message, typically sent every 20 ms (up to 50 messages/s). Very high reliability target: ~99.999%.

downlink path (operator to vehicle). For Indirect Control ToD, a looser round-trip budget on the order of 300 ms is reported, reflecting a different control paradigm and operational envelope [109].

These requirements motivate the adoption of efficient video compression and streaming configurations. In practice, ToD streaming must balance three coupled dimensions: *(i)* bandwidth consumption, driven by resolution, frame rate, codec choice and rate-control; *(ii)* processing delay, which may increase with more complex compression settings; and *(iii)* perceptual quality, which must preserve decision-relevant cues for safe maneuvering. This coupling is particularly challenging under variable radio conditions, where rate adaptation may be required to maintain latency stability and to avoid buffering-induced control impairment.

Experiments with ToD have been preliminarily performed during the research period at Universidad Politécnica de Madrid (UPM), using an operator driving station representative of a remote driving controller center, in order to understand the requirements discussed above.

Fig. 4.2 shows the driving station and two representative visual outcomes obtained under different uplink conditions. The upper half illustrates a high-quality video configuration, enabled by sufficient bandwidth, which preserves scene details and supports timely perception. The lower half illustrates an extremely bandwidth-limited condition, where strong compression and packet losses lead to severe visual artifacts and loss of semantic detail. This hands-on experience motivated the conducted study aimed to reduce the uplink bandwidth demands of ToD.

Fig. 4.2: ToD operator workstation used during the research stay at UPM, including the driving interface and the visual feedback supporting remote vehicle control.

4.2.2 State of the Art

Video coding and streaming for ToD has attracted considerable interest in recent years, motivated by the stringent requirements imposed by real-time remote control over mobile networks. In ToD, multiple high definition camera views must be delivered in uplink with high throughput and low, stable latency, in order to provide the remote operator with a sufficiently rich and timely representation of the surrounding environment. Meeting such requirements is particularly challenging in wireless

systems, where the experienced QoS can vary significantly due to interference, coverage conditions, and network congestion. In addition, the uplink is typically more constrained than the downlink, and transmission power at the vehicle side is limited compared to infrastructure nodes. For these reasons, the communication footprint of multi-camera streaming is widely recognized as one of the main obstacles to the scalability of ToD deployments.

A first and widely investigated research direction aims at reducing the bitrate of the uplink video stream through content-aware encoding. A representative approach is to identify regions of interest in each frame and encode them at higher quality, while compressing or blurring less relevant regions to save bandwidth [110]. Building upon the same principle, risk-aware strategies have been proposed in the [111], where the encoding decisions are guided by safety-related metrics, e.g., by exploiting RSS-inspired safe distances to prioritize objects that are more likely to become dangerous. Other works explored hybrid pipelines that combine object recognition with lightweight image processing to preserve decision-critical cues and reduce bitrate while controlling end-to-end latency [112, 113, 114]. When multiple camera views are available, the authors in [115], focus on view selection and prioritization. In this case, the bitrate is adapted by assigning different priorities to different camera views depending on the driving context, such as the current maneuver and the vehicle trajectory, so that the most relevant view(s) receive higher quality. A complementary solution address multi-camera adaptation by exploiting estimates of the available uplink capacity, and dynamically scaling the resolution or bitrate of each stream in order to remain within the capacity budget [116]. Finally, robustness to wireless variability has also been investigated through multi-link streaming, where the system adapts the bitrate and, when needed, switches between alternative access links to maintain low-latency playback [117].

In parallel to bitrate reduction mechanisms, other contributions emphasize the role of network-side support to keep the ToD service within its target QoS. For example, the authors in [118] present a system which dynamically requests or reserves radio resources to maintain the required streaming performance, focusing on resource provisioning rather than on video source adaptation. A further direction, more architectural in nature, proposes leveraging local perception and scene reconstruction to reduce the amount of raw data that must be transmitted. In particular, Digital Twin (DT) based solutions can reconstruct the environment using on-board sensing and computation, allowing interpreted information (e.g., object lists or reconstructed scenes) to be delivered instead of raw sensor streams, thereby alleviating network load [119].

While the above literature provides valuable building blocks, it mainly addresses the optimization of a single tele-operated vehicle, implicitly assuming that each vehicle can be treated in isolation. This assumption becomes limiting in fleet-oriented scenarios, where multiple ToD vehicles may simultaneously operate in the same geographical area and compete for the same uplink resources. Moreover, existing studies that consider fleets and tele-operated platooning primarily focus on feasibility, control, and operational aspects [120, 121, 122, 123], and do not explicitly investigate how uplink streaming demand can be reduced by exploiting cooperative opportunities arising from vehicle proximity and shared context.

This chapter is motivated by two observations. First, not all sensory data have the same relevance for the remote operator at all times: depending on the driving context, some camera views may carry limited safety-relevant information. Second, when fleet vehicles travel in proximity and can form a platoon, their sensing and control tasks become partially coupled: on-board detection can complement remote perception, and cooperative longitudinal control can reduce the need for continuous high-quality streaming from follower vehicles, provided that transitions back to teleoperation remain safe. These considerations motivate the proposed framework, which combines context-aware bandwidth adaptation with edge-assisted platoon formation and cooperative control, aiming at reducing uplink demand while preserving the operator's situational awareness in the most critical conditions.

4.3 Proposed framework

4.3.1 Reference scenario and main assumptions

Our study focuses on a fleet of ToD vehicles traveling in the same direction and sharing portions of their routes, as common for logistic operations. We propose a framework aimed at optimizing the uplink bandwidth requirements of the fleet by exploiting a context that jointly encompasses *(i)* the proximity of the vehicles and *(ii)* the detection capabilities of their on-board sensors.

The following assumptions hold for each ToD vehicle:

- It is equipped with four on-board cameras, each covering a 90° field of view [124], thereby providing full 360° perception. This configuration partitions the surrounding area into four sectors: front, rear, left, and right, similarly to [125].
- It is equipped with GNSS for absolute positioning and ultra-sonic sensors that detect the presence of objects in proximity and other obstacles.
- It can interact with the remote driver through long-range V2N connectivity (uplink video and downlink commands) and can exchange data with vehicles in proximity through short-range V2V communications.

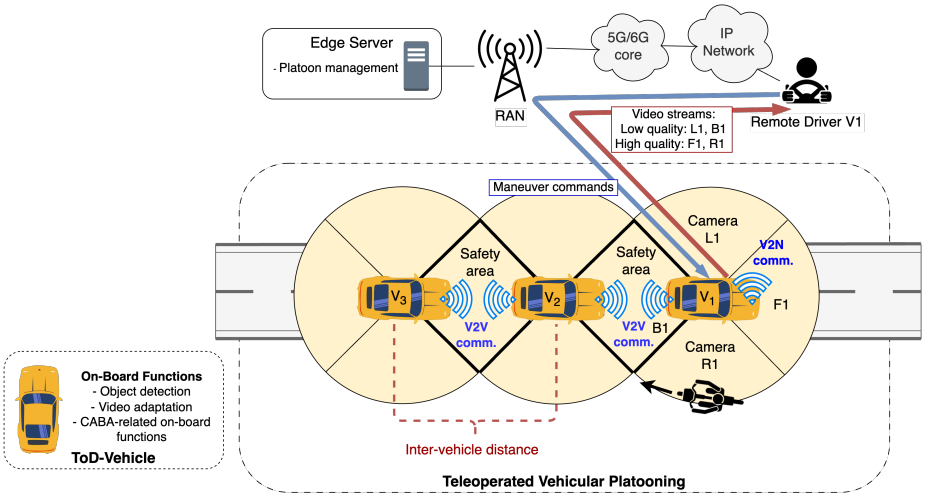

Fig. 4.3: Overview of the reference scenario. In the figure, only the video streams from the lead vehicle (V_1) and its corresponding command reception are shown for clarity. However, the follower vehicles are also teleoperated and have similar uplink and downlink traffic flows according to the CABA proposal. The figure illustrates an example platoon of three teleoperated vehicles, V_1, V_2, and V_3, with the platoon size being a configurable system parameter.

4.3.2 Our proposal

Since the ToD vehicles may share part of their trip and are equipped with V2V capabilities, they can dynamically form a platoon, as illustrated in Fig. 4.3. While the lead vehicle continues to be remotely controlled through V2N, follower vehicles switch their longitudinal control from teleoperation to a cooperative mode governed by CACC.

Besides the common advantages of platooning (fuel efficiency and improved road capacity) [126], the proposal yields two additional benefits:

- By switching from teleoperation to CACC, the workload of the remote operators of follower vehicles can be reduced.
- Once the platoon is established, follower vehicles can reduce the quality of the transmission of selected video streams and release uplink radio resources accordingly. The saved bandwidth can be exploited, for example, to increase the number of ToD vehicles supported within the resources allocated to a remote driving network slice [127].

Follower vehicles do not completely suppress video streaming to their remote operators. In conventional platooning services, where follower vehicles may require human supervision, drivers remain on-board and can continuously observe the environment, being ready to take control at any time. In contrast, in teleoperated pla-

toons the operators are located in a remote control center and must rely on transmitted sensor data to perceive the environment whenever they need to assume control, e.g., in the event of platoon break-up.

For this reason, streaming remains essential to safely support transitions back to teleoperation and can also be exploited, once the platoon is established, to maintain a minimum level of situational awareness for the remote operators of follower vehicles. This motivates the context-aware bandwidth adaptation scheme described next.

4.3.3 Bandwidth adaptation in the platoon

In a legacy scenario, each camera stream is transmitted at high quality to ensure comprehensive remote situational awareness. Our proposal dynamically adjusts video quality (and thus bandwidth) according to locally available context, particularly the detection of surrounding objects, and is activated when vehicles are sufficiently close to form and maintain a platoon.

The adaptation uses two per-camera bitrate levels:

- B_{HQ} for high-quality streaming,
- B_{LQ} for low-quality streaming.

Each vehicle continuously monitors the sector covered by each camera. When at least one object is detected within a sector (via ultra-sonic sensing), the corresponding camera stream is upgraded to B_{HQ}. Otherwise, it is downscaled to B_{LQ}. For example, in Fig. 4.3 a bicycle enters the right sector of V_1; the stream from camera R_1 is then transmitted at B_{HQ}. Conversely, if no objects are present in the left sector, camera L_1 can be transmitted at B_{LQ}.

When vehicles are in platoon formation, the presence of the adjacent platoon member is a known and continuously tracked condition. Therefore, the front camera of a platoon member following the leader and the rear camera of the leader are kept at B_{LQ} if no external objects (i.e., other than the adjacent platoon vehicle) are detected in the corresponding sectors. In other words, these streams are promoted to B_{HQ} only when additional objects appear in front of the follower or behind the leader, since the leader and follower mutual presence does not convey new situational information to the remote operators.

For safety reasons, the front camera of the leader is always transmitted at high quality.

The front camera of the leader is always transmitted at high quality. The proposed sector-based adaptation can be combined with Region of Interest (ROI)-based

encoding approaches [110, 111] to further reduce bandwidth, which is left for future work.

4.3.4 Platoon management

Initially, fleet vehicles are tele-operated. Besides streaming video and receiving commands, each vehicle periodically transmits its position (e.g., via GNSS) to an edge-based platoon management application. The latter *(i)* tracks fleet vehicles and their reciprocal distance, *(ii)* decides when the platoon can be formed and when it should be broken, and *(iii)* instructs vehicles to join/leave the platoon and assigns roles as Platoon Leader (PL) or Platoon member (PM).

4.3.5 Platoon formation

The platoon management application triggers platoon formation whenever the distance between two ToD vehicles V_i and V_j, denoted $d(V_i, V_j)$, satisfies:

$$d(V_i, V_j) \leq d_{max}. \tag{4.1}$$

The distance d_{max} is derived as the maximum spacing that still allows the definition of a *safe area* between two adjacent vehicles, namely an inter-vehicle region that cannot be entered by external objects without being detected by at least one of the on-board sensors. This property relies on the overlap between the rear sensing coverage of the leading vehicle and the front sensing coverage of the following vehicle when they travel in close proximity.

The straight-road configuration in Fig. 4.4 represents the baseline case in which the two vehicles are aligned and the overlap of the rear and front sensing coverages is symmetric. For completeness, we also consider the more general case in which the vehicles are not perfectly aligned (e.g., due to road curvature, local maneuvers, or transient tracking errors). In this case, the geometry of the overlap region changes and the maximum admissible spacing ensuring that the safe area remains closed must be derived accordingly. The corresponding derivation and the resulting expression for the maximum inter-vehicle distance under heading misalignment are provided in Appendix 6.

In the baseline case of aligned vehicles moving along a straight road (Fig. 4.4), the rear coverage area of the leading vehicle and the front coverage area of the following vehicle overlap symmetrically. The safe area, defined by the polygon AC_1C_2B, forms a square that can be decomposed into two congruent isosceles right triangles. By construction, $\overline{AC_1} = \overline{AC_2} = r$ (with r the sensing range of the ultra-sonic sensors) and $\angle A = 90°$. Hence, the maximum inter-vehicle distance (the diagonal of the square) is:

$$d_{max} = \overline{C_1C_2} = r\sqrt{2}. \tag{4.2}$$

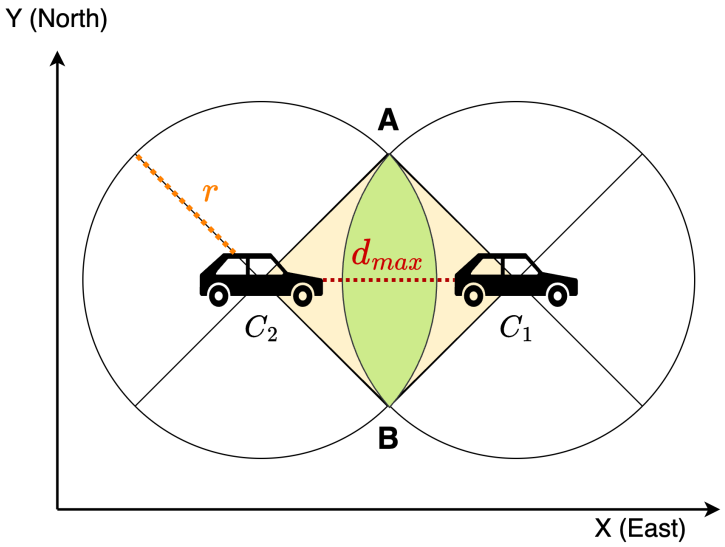

Fig. 4.4: Maximum inter-vehicle distance.

4.3.6 Platoon control

After receiving instructions from the edge application, follower vehicles switch from teleoperation to CACC. One key parameter is the spacing policy that determines the intra-platoon gap d_{gap} and enables computation of a target speed. The gap can be constant [128, 129] or determined via a constant-time headway policy [130], e.g., $d_{gap} = h \cdot v$, where v is the vehicle speed and h the headway time.

Upon platoon formation, vehicles exchange Platoon Control Messages (PCMs) through V2V communications, similarly to [131, 132], containing the information required by the CACC controller (e.g., speed and acceleration of the leader and/or predecessor, depending on the enforced scheme), while relative distance can be obtained through local sensing (e.g., radar).

When condition in Eq. 4.1 is no longer satisfied, the platoon management application detects that the platoon broke up and control of the non-leader vehicle is passed back to the remote operator. Consequently, radio resources must be reallocated to enable high-quality video streaming as needed.

4.3.7 Analysis of the platoon formation triggering condition

The identified distance for triggering platoon formation, d_{max}, should verify:

$$d_{max} \geq d_{gap}. \tag{4.3}$$

Fig. 4.5 shows that for a sensing range r in the order of 8 m [133], d_{max} is larger than the desired intra-platoon gap under a headway-based policy with $h = 0.5$ s [129] up to relatively high speed values. The condition is always satisfied, whatever

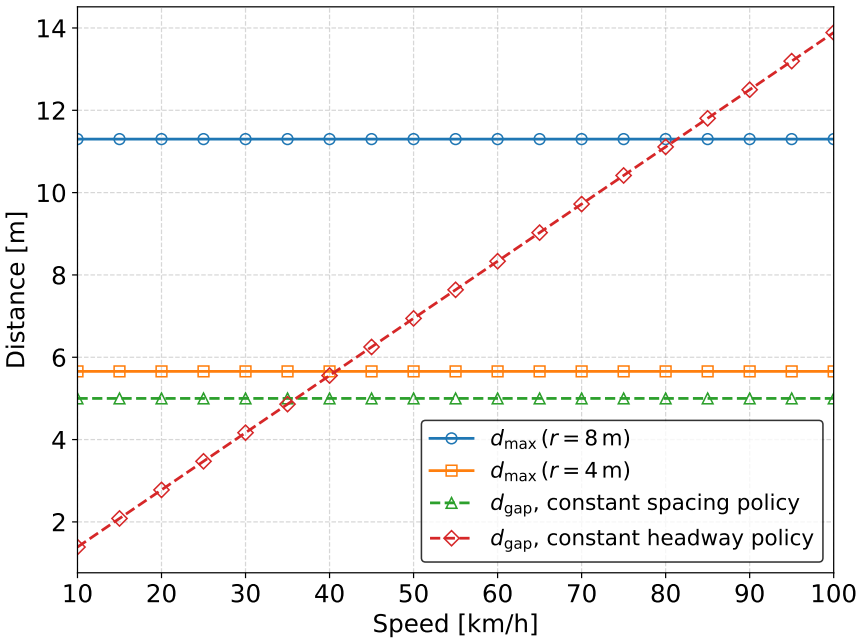

Fig. 4.5: Tele-operated platooning: feasibility regions.

the speed and sensing range, for a constant spacing policy with d_{gap} = 5 m. Conversely, for smaller sensing ranges (e.g., r = 4 m [134]) the feasible region becomes more restrictive for headway-based spacing.

Overall, these values confirm that for typical speed constraints for tele-operated vehicles, in the order of 50 km/h [124], the proposed triggering condition can be compatible with platooning policies even when the sensing range is limited.

4.4 Performance evaluation

4.4.1 Tools and main settings

The effectiveness of the proposal has been evaluated using the SUMO simulator [108], extended with custom code to replicate the on-board camera sectors and sensing devices and to implement the bandwidth adaptation logic.

The target platoon moves in an urban scenario along a straight road segment of approximately 1.1 km, with two lanes per direction. Other vehicles enter the scenario with a varying rate λ [veh/s], resembling different traffic density levels and different objects to be detected by the cameras of the tele-operated vehicles.

Table 4.2 summarizes the main simulation settings.

4.4.2 Compared schemes

The performance of the following schemes has been compared:

Parameter	Value
Platoon size, N	2
Video bitrate (B_{LQ}/B_{HQ})	1.125 Mbps [135] / 8 Mbps [124]
Object arrival rate, λ	[0.1, 0.33] veh/s
Headway time, h	0.5 s [129]
Sensing range, r	8 m

Table 4.2: Main simulation settings

- *Legacy*: all members of the tele-operated platoon stream all camera feeds at high quality, with bitrate B_{HQ}.
- *CABA*: for each vehicle, a given camera streams at high quality (B_{HQ}) if at least one object is within the sector covered by that camera; otherwise it streams at low quality (B_{LQ}). The leader front camera is always streamed at B_{HQ}.

4.4.3 Metrics

The following metrics have been derived.

Instantaneous platoon aggregate data rate

The parameter $R_{\text{agg}}(t)$ is defined as the instantaneous aggregate uplink bandwidth required by the tele-operated platoon. At time t, it is the sum of the bitrates of the four camera streams from every vehicle:

$$R_{\text{agg}}(t) = \sum_{i=1}^{N} \sum_{k=1}^{4} R_{i,k}(t), \tag{4.4}$$

where $R_{i,k}(t)$ denotes the instantaneous bitrate of camera k on vehicle i at time t, with $R_{i,k}(t) \in \{B_{LQ}, B_{HQ}\}$.

Average aggregate platoon data rate

For a trip of duration T, the time-averaged aggregate data rate is:

$$R_{avg} = \frac{1}{T} \int_0^T R_{\text{agg}}(t)\, dt. \tag{4.5}$$

4.4.4 Results

Instantaneous data rate

Fig. 4.6 reports the platoon aggregate and individual data rate for different λ settings. It can be observed that both metrics are significantly reduced for the proposed

CABA scheme compared to the legacy approach. In particular, as λ increases, objects are detected more frequently across sectors and CABA keeps a larger number of cameras at high quality, which reduces the achievable gain while preserving a clear reduction compared to the baseline.

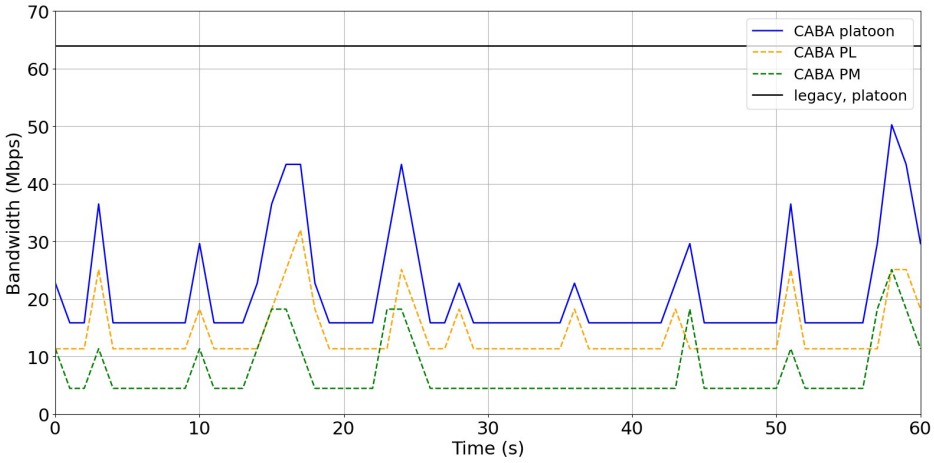

fig:instBlow-early

(a) λ=0.1 veh/s.

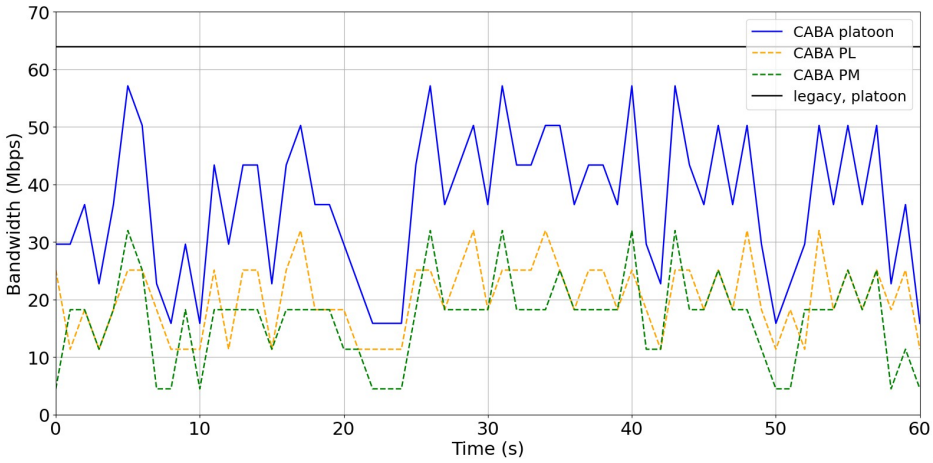

fig:instBhigh-early

(b) λ=0.33 veh/s.

Fig. 4.6: CABA vs. legacy solution: instantaneous data rate values required for different λ settings.

Average data rate

Table 4.3 reports the average data rate for the overall platoon as well as for the PL and the PM, when varying λ. Consistently with the instantaneous results, R_{avg} increases with λ. Gains are larger for the PM, because the PL always streams the front camera at high quality.

Veh/s	PL		PM		Aggregate platoon	
	[Mbps]	[%]	[Mbps]	[%]	[Mbps]	[%]
0.10	13.96	56.37	7.31	77.16	21.28	66.75
0.33	20.27	36.66	16.56	48.25	36.83	42.45

Table 4.3: Average data rate and gain metrics for CABA under different vehicle insertion rate.

Veh/s	Legacy	CABA	Teleoperated platooning
0.10	64 Mbps	21.28 Mbps	32 Mbps
0.33	64 Mbps	36.83 Mbps	32 Mbps

Table 4.4: Legacy vs. CABA vs. teleoperated platooning: aggregate platoon average data rate under different vehicle insertion rate.

CABA vs. traditional teleoperated platooning

For completeness, we compare the overall bandwidth requirements of CABA against a conventional teleoperated platooning scheme [123], where the PL streams the four cameras to the remote operator, while the follower relies on CACC. As shown in Table 4.4, for the considered two-vehicle platoon, CABA yields aggregate average rates that are comparable to, or even lower than (for low λ), the benchmark. This is relevant because CABA maintains a continuous (though potentially downscaled) visual channel also for the follower, which is beneficial for safe transitions if the platoon breaks.

4.5 Preliminary PoC implementation

Complementary to the results obtained in the simulator presented in the previous section, we developed a PoC to validate the feasibility of an end-to-end teleoperation service over a real 5G network (Fig. 4.7). The PoC replicates the theoretical framework presented in the chapter by re-creating, on real devices and hardware, the same logical roles considered in the referenced scenario:

- a terrestrial robotic platform acting as a ToD vehicle;
- a driving station hosting the operator generating control inputs;
- a 5G-SA network for V2N connectivity;
- an edge server hosting virtualized functions.

In particular, the 5G-SA connectivity used in this PoC relies on the same private campus network deployment described in Chapter 3 (Section 3.4.2, Fig. 3.11)- This

Fig. 4.7: PoC teleoperation setup: robot and driving station used over the 5G-SA testbed.

includes operation on the n77 band, a complete RAN+5G Core network architecture, and access through 5G routers acting as UEs with SIM-based subscription and device identification for authorization.

In this PoC, the remote operator station is hosted at the network edge as showed in Fig. 4.8, and the edge server is co-located within the 5G-SA infrastructure, enabling low-latency communication of the teleoperation service and local handling of the associated data flows.

4.5.1 Hardware components

The communication pattern follows the ToD service model introduced in Section 4.2.1: uplink traffic delivers video and telemetry streams to the operator, while downlink traffic conveys time-critical control commands. On the robot side, connectivity to the 5G network is provided through a commercial 5G router operating as a UE (Chapter 3, Section 3.4.2), connected via Ethernet to the on-board computing unit, implemented on an NVIDIA Jetson Nano. Fig. 4.9 reports the 5G-SA private network dashboard with the on-board UE connected.

Fig. 4.8: Remote driving station used over the 5G-SA testbed.

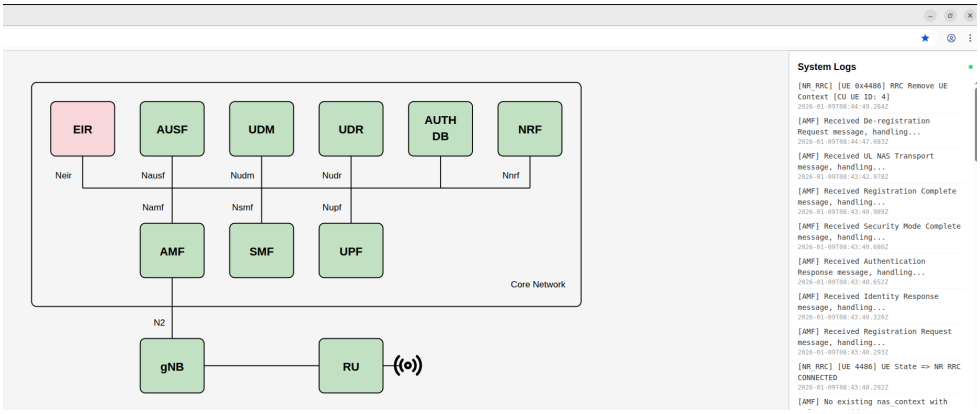

Fig. 4.9: Dashboard of the 5G-SA private network.

The robot is equipped with a multi-modal sensor suite, including a LiDAR and a front-facing camera supporting the operator view. To align the PoC with the positioning-dependent workflow assumed by CABA, a GNSS receiver is integrated following the same hardware integration pipeline adopted for the SU discussed in Chapter 3 (Section 3.2.1).

4.5.2 Software components

The Jetson Nano runs the robot software stack on top of a Linux operating system and leverages Robot Operating System (ROS)v2 as the middleware layer for integrating sensing, control, and networking functionalities. In particular, ROSv2 structures the application into distributed nodes that exchange data via *topics* (publish/subscribe) and *service/actions* (request/response), thereby enabling a modular architecture.

The GNSS node acquires the positioning measurements and publish them on legacy ROSv2 topics, enabling the trajectory logging through GPS coordinates and periodic position reporting toward periodic updating of topics.

In addition, the same positioning updates are forwarded to the edge server, where the platoon management application collects them to monitor the robot trajectory and to support coordination decisions consistent with the edge-assisted workflow assumed by CABA (e.g., tracking relative distances and triggering platoon-related actions when required).

This addition provides a step forward to realizing PoC for a real-world test of the CABA solution, aligned with the prototyping pipeline adopted throughout the thesis.

4.5.3 Video streaming configuration

A visualization pipeline is implemented to provide the operator with situational awareness based on the front camera stream. On the operator side, the video is displayed through RViz. ROSv2 messages are transported across the vehicle and operator domains using Zenoh [136], installed on both sides to facilitate communication between the mobile gateway and the wired operator connection, similarly to recent work in the literature [137]. Zenoh is an open-source, data-centric communication protocol and middleware designed to support efficient data distribution across heterogeneous and geographically distributed systems, with location-transparent abstractions for high-performance publish/subscribe and distributed queries. In robotics and edge-enabled deployments, Zenoh is commonly used to complement ROSv2 communications by providing an alternative data plane that can span LAN/WAN segments and can be integrated through bridging mechanisms that relay selected ROSv2 entities (topics/services/actions) over Zenoh. In our setup, it is used to bridge the ROSv2 environments by relaying the teleoperation-relevant *topics* across the 5G V2N segment.

4.5.4 Early testbeds

Different video streaming settings have been considered to assess perceived quality and bandwidth demands over the private 5G-SA infrastructure.

As a baseline, we first evaluated the *uncompressed* stream produced by the *main camera* and visualized in RViz. In this configuration, video stream is published on the raw topic */camera/color/image_raw*.

To reduce the uplink load, the camera stream is therefore compressed on-board before being sent over the 5G link. Compression is implemented through a dedicated ROSv2 node based on the *image_transport* package (executable *republish*, node name *image_compressor*). The node subscribes to the raw stream on */camera/color/image_raw* and republishes a compressed version on */camera/color/image_raw/compressed*. The compression level is controlled through the *jpeg_quality* parameter, configured in the 70 to 80 range (set to 75 in our setup). This configuration reduces the bandwidth demand by roughly 60% with respect to the uncompressed baseline, while preserving a video quality that was deemed adequate for early teleoperation trials.

4.5.5 Discussion and outlook

The developed PoC complements the simulation-based evaluation by providing an end-to-end experimental baseline over a real 5G-SA infrastructure. In particular, it validates the feasibility of the teleoperation loop under realistic transport conditions and demonstrates the practical integration of the robot software stack (ROSv2), the edge-hosted functions, and alternative video delivery strategies (including a Zenoh-based transport). Preliminary bandwidth measurements obtained with on-board compression confirm that the end-to-end pipeline can operate with sustained uplink rates that are consistent with the order of magnitude assumed by the CABA design; however, the current prototype still relies on lightweight, frame-based compression and does not provide fine grained control of the generated bitrate.

Accordingly, matching the lower bandwidth operating points explicitly considered in the CABA proposal (1.125 and 8 Mbps) would require more efficient and bitrate-controllable video compression schemes, so as to reproduce the target profiles while maintaining sufficient visual quality for teleoperation.

At the same time, while the proposed setup is suitable to validate the framework assumptions in a controlled environment and to complement the simulation evidence with practical measurements, a full experimental validation of CABA and of the associated edge-assisted coordination logic would require extensive and systematic field trials involving at least two robots to resemble a platoon scenario. Such trials would involve broader operational scenarios, longer measurement campaigns,

larger datasets. Nevertheless, the presented PoC constitutes a solid foundation for future work aimed at large-scale real-world validation.

4.6 Main findings and future works

This chapter addressed network efficiency for fleets of tele-operated vehicles by leveraging cooperative opportunities that arise when multiple ToD vehicles share portions of their route. To this end, we proposed CABA, a context-aware multi-camera bandwidth adaptation scheme that exploits local object detection and inter-vehicle proximity to reduce the uplink demand of video streaming while preserving operator situational awareness in safety-relevant conditions.

The proposal was integrated with an edge-assisted platoon management workflow that enables dynamic platoon formation and supports safe transitions between teleoperation and cooperative longitudinal control through CACC.

Simulation results obtained in SUMO showed that CABA can substantially reduce the aggregate uplink data rate compared to legacy fixed high-quality streaming, with gains that depend on the external traffic density and resulting rate of relevant objects detected across camera sectors.

In addition, the preliminary PoC implementation over a real 5G-SA testbed provided an end-to-end experimental baseline, validating the feasibility of the teleoperation pipeline.

Future work will primarily investigate how to further reduce the uplink footprint of CABA while preserving decision-critical visual cues, for instance by integrating finer-grained adaptation techniques such as ROI-based and risk-aware compression. In parallel, a tighter cross-layer coupling with network-side resource management will be considered to improve latency stability under fluctuating radio conditions and competing traffic.

Finally, while the preliminary PoC provides a practical baseline, a comprehensive assessment of the end-to-end trade-offs and of the deployability of CABA in real fleet operations will require extensive field trials beyond the scope of this thesis.

5

Conclusions and future works

This thesis investigated how sensing, perception, and communication can be jointly engineered to make CCAM services more scalable and dependable under real-world constraints. A distinctive aspect of the work is its strongly practical orientation: beyond proposing design principles and algorithms, each contribution was validated through concrete PoC implementations and experimental activities, aimed at assessing feasibility, integration complexity, and performance under realistic constraints. Across the different chapters, the common methodological thread is the explicit coupling between *(i)* what the road user/vehicle needs to perceive and communicate for safety and controllability, and *(ii)* what the network and embedded devices can sustain in terms of latency, throughput, and computational budget.

5.1 Summary of contributions and main outcomes

5.1.1 VRU awareness for a more comprehensive CCAM ecosystem

Chapter 2 addressed VRU awareness in cooperative ecosystems by focusing on a specific and increasingly relevant category: e-bikes. The analysis highlighted that message generation strategies designed around conventional vehicles may not fully capture the mobility patterns and stability dynamics of light VRUs. The proposed refinements showed that additional trigger logic can increase the timeliness and representativeness of VAMs, while keeping the overhead limited and maintaining compliance with practical channel-load constraints. Beyond the specific case study, the resulting takeaway is that CCAM safety services benefit from domain-specific modeling of mobility and user behavior, especially when heterogeneous road users share the same cooperative space.

5.1.2 A modular sensing, perception, and communication platform for CCAM

Chapter 3 presented an integrated platform that combines low-cost sensing, embedded inference, and hybrid connectivity to support cooperative applications and data

dissemination. The main result is that the platform can achieve useful accuracy–efficiency trade-offs on constrained hardware by leveraging context-aware fine-tuning of lightweight models, while remaining compatible with realistic end-to-end communication workflows (short-range dissemination and 5G-assisted connectivity). The proof-of-concept evaluation showed that: *(i)* context-aware specialization improves embedded perception without requiring heavier ML models, *(ii)* personalization is beneficial for driver monitoring when human factors remain part of the control loop, and *(iii)* data dissemination performance depends not only on radio access but also on application-layer design choices (e.g., broker load and placement), which must be treated as first-class engineering parameters.

5.1.3 Context-aware bandwidth adaptation for video stream of teleoperated fleets of vehicles

Chapter 4 focused on ToD as a CCAM use case characterized by stringent uplink demands and tight latency constraints. Starting from the observation that multi-camera streaming can dominate the uplink budget of teleoperated vehicles, the chapter introduced a context-aware bandwidth adaptation scheme (CABA) that reduces the communication footprint by exploiting two forms of redundancy:

- **Redundancy in the scene:** video streams are more valuable when relevant objects are present in the corresponding sector, whereas empty sectors can be safely downscaled without compromising situational awareness.
- **Redundancy across vehicles in proximity:** when vehicles are sufficiently close to form a platoon, partial overlap in the monitored space enables more aggressive adaptation, while still guaranteeing that the critical inter-vehicle region remains observed.

The resulting system architecture (Fig. 4.3) couples per-camera adaptation with an edge-assisted platoon management mechanism that can orchestrate the transition from teleoperation to cooperative longitudinal control, leveraging CACC and V2V communications.

In the considered simulation settings, CABA reduced the average uplink bitrate substantially with respect to a fixed high-quality baseline, with savings that remain significant even under higher object arrival rates. Importantly, the results suggest that adaptive streaming can bring the aggregate demand close to, and in some cases below, what would be achieved by more rigid teleoperated-platooning baselines, while retaining a continuous (although downscaled) visual channel for safer transitions in case of platoon breakdown.

Complementary to the simulation-based evaluation, a preliminary PoC was developed over a real 5G-SA testbed to validate the feasibility of an end-to-end teleoperation pipeline. The PoC demonstrated the integration of a robot software stack based on ROSv2, edge-hosted functions, and alternative video delivery strategies, including Zenoh-based transport. The preliminary measurements confirmed that the uplink demand is strongly influenced by the adopted video delivery configuration, and highlighted the need for suitable compression schemes to reproduce, in practice, the operating points assumed by the CABA design.

5.2 Implications for CCAM design

Three broader implications emerge from the thesis:

- **Context should drive resource usage.** Across cooperative messaging, embedded perception, and teleoperation streaming, the results support the same design rule: allocate communication and computation resources primarily where/when context indicates potential safety relevance.
- **Hybridization is not optional.** Practical CCAM systems require hybrid designs that combine local sensing and inference with cooperative data exchange and, when available, network assistance.
- **System performance is end-to-end.** For both cooperative awareness and teleoperation, the final quality experienced by users (drivers, operators, road users) depends on cross-layer interactions: sensor assumptions, inference latency, message/video generation logic, connectivity choice, and application-layer orchestration.

5.3 Future research directions

While the results support the viability of the proposed approaches, there is room for improvements. Building on the thesis outcomes, the following directions appear both technically relevant and practically impactful.

5.3.1 Teleoperation: from bitrate savings to end-to-end guarantees

A first priority is to extend CABA from a demand-reduction mechanism to a full end-to-end strategy that jointly optimizes encoding, scheduling, and slice-level resource allocation. This includes:

- integrating CABA with finer-grained adaptation mechanisms such as ROI-based or risk-aware encoding while controlling added processing delay;

- coupling the adaptation logic with network-side predictors and multi-link strategies to mitigate radio fluctuations;
- quantifying the impact on teleoperation performance through human-in-the-loop studies that relate streaming configurations to task success rates and reaction times, rather than to bit rate alone. On the experimental side, reproducing the target operating points considered by CABA requires bitrate-controllable compression and rate-control strategies that can sustain low-bitrate profiles while preserving decision-critical visual cues.

A subject matter of future work could also concern the investigation of integrated prototypes combining the context-aware perception capabilities presented in Chapter 3 with the context-aware streaming logic of ToD services discussed in Chapter 4, by combining vision-based detections with ultra-sonic sensing to obtain more robust CABA triggers, and extending the driver monitoring pipeline to the remote ToD operator to quantify workload and relate the operator state to network impairments and adaptive streaming decisions.

5.3.2 Platooning under richer geometry and mixed traffic

The safe-area formulation for the ToD platooning should be generalized and validated for non-trivial vehicle kinematics, including non-zero heading differences (already identified as a system-level requirement in Chapter 4), curved roads, and lane changes. A complementary step is to study mixed settings where teleoperated vehicles share the road with partially automated vehicles and VRUs, to understand when platoon formation remains beneficial and when it becomes fragile.

5.3.3 Security and privacy extensions

The prototypes and experiments of the thesis highlight that scaling CCAM services is not only a performance problem but also a security and privacy one, especially when VRU devices and driver monitoring are involved. Future deployments should therefore incorporate privacy-by-design principles, including data minimization (collect only what is required by the target service), on-device processing whenever feasible, and strict retention policies.

For VRUs, identifiers and traces should be protected through pseudonymization and controlled re-identification procedures when required by regulation.

For driver/operator monitoring, sensitive visual data should be handled locally or under strong access control with explicit consent, for instance, shared with the paired digital counterpart through secured interfaces.

From a security perspective, end-to-end protection is needed across heterogeneous links (V2V and cellular), including authentication of messages, integrity protection, and secure provisioning of devices and software updates. Interoperability testing should explicitly include security mechanisms, since lightweight prototypes may otherwise overestimate performance by omitting cryptographic and key-management overheads.

5.3.4 From small-scale platform assessment to large-scale trials

The platform introduced in Chapter 3 provides a practical basis for larger experiments. Future work should run extended trials combining short-range V2V communications and 5G-assisted data realistic settings in terms of number of equipped vehicles and mobility patterns in more crowded road contexts, including security constraints and interoperability testing.

Experiments conducted under large-scale settings would also provide the chance to collect datasets that could be shared for reproducibility purposes.

5.4 Closing remarks

Overall, the thesis supports the position that scalable CCAM services will be enabled by **context-driven** system designs, where sensing, perception, and communication are co-optimized instead of being engineered in isolation. Within this perspective, the proposed contributions provide both *(i)* concrete mechanisms, from refined message triggering to embedded model specialization and context-aware teleoperation streaming, and *(ii)* a unifying architectural rationale that can guide future research and deployments in heterogeneous, resource-constrained mobility environments.

6

Appendix: Maximum Inter-vehicle Distance for Maintaining Platoon Formation

6.1 Problem formulation

This appendix derives the maximum inter-vehicle distance that guarantees the platoon remains cohesive and safe, i.e., that the **safe area** between each pair of adjacent vehicles remains continuously monitored by onboard sensors.

We consider two consecutive vehicles in the platoon: the leader, centered at C_1, and its follower, centered at C_2. The relevant sensing coverage for monitoring the inter-vehicle region is modeled, for each vehicle, as a circular region of radius r (with r being the sensing radius of ultra-sonic sensors) centered at the vehicle position (Figs. 6.1–6.4). For analytical convenience, range is represented through circles of radius r, while the safe-area construction considers only the rear/front directional portions that are relevant to monitoring the inter-vehicle region.

Let θ_1 and θ_2 denote the leader and follower headings, respectively, defined with respect to the global reference axes shown in Fig. 6.2 (Y axis oriented North and X axis oriented East). We define the relative heading mismatch as:

$$\Delta\theta \triangleq \theta_1 - \theta_2. \tag{6.1}$$

The subsequent derivation depends on the magnitude of the mismatch only, since changing the sign of $\Delta\theta$ produces a mirrored geometry (the construction is symmetric and the pivot point for the auxiliary rotation switches between A and B). Therefore, without loss of generality, the analysis can be restricted to $\Delta\theta \geq 0$.

Moreover, the notion of *safe area* adopted in this work is tied to the directional sensing assumed by the system architecture: the inter-vehicle region must be monitored by the *rear-side* sensing of the leader and the *front-side* sensing of the follower. Under the reference configuration, this directional coverage corresponds to a $90°$ sector for each relevant side (consistent with the $90°$ field of view of each camera sector). To guarantee that the inter-vehicle region remains fully contained in both the leader rear sector and the follower front sector, the relative yaw mismatch must not exceed such half-aperture; otherwise, one of the two directional coverages would no

longer include part of the inter-vehicle region, and the safe area could not be guaranteed to remain closed and continuously monitored. Hence, the feasible heading mismatch domain is conservatively bounded as

$$0 \le \Delta\theta \le \frac{\pi}{4}, \tag{6.2}$$

or, equivalently, $|\Delta\theta| \le \pi/4$ when symmetry is made explicit.

Formally, for any admissible relative heading mismatch, the goal is to determine the largest inter-vehicle spacing that still guarantees that the safe area remains closed and continuously monitored. Specifically, for

$$\Delta\theta \in \left[-\frac{\pi}{4}, \frac{\pi}{4}\right], \tag{6.3}$$

let $d(\Delta\theta)$ denote the inter-vehicle distance between the vehicle centers, i.e., $d(\Delta\theta) = \overline{C_1 C_2}$. For a given $\Delta\theta$ and a given spacing d, the sensing geometry determines (when it exists) a unique pair of circle-intersection points (A, B) and, consequently, a unique safe-area polygon with vertices (A, C_1, C_2, B). The problem is therefore to characterize the maximum admissible distance

$$d_{\max}(\Delta\theta) \triangleq \sup\{\, d \ge 0 \mid \text{the safe area associated with } (\Delta\theta, d) \text{ is closed and fully monitored} \,\}, \tag{6.4}$$

so that feasibility can be stated as

$$d(\Delta\theta) \le d_{\max}(\Delta\theta), \qquad \Delta\theta \in \left[-\frac{\pi}{4}, \frac{\pi}{4}\right]. \tag{6.5}$$

The two coverage circles intersect at two points, denoted by A and B in the original (non-rotated) configuration (Fig. 6.2). The safe area is the region that must remain jointly monitored by the rear-side sensing of the leader and the front-side sensing of the follower; at the feasibility boundary this region is enclosed by a quadrilateral whose vertices are determined by the two vehicle centers and by the circle-intersection geometry.

In the considered teleoperated/platooning setting, heading mismatches are small. Moreover, the geometric construction adopted in this appendix is defined over the limiting mismatch of ±45° shown in the reference figures, which corresponds to the extreme configuration supported by the sensing geometry under study. Therefore, the analysis is restricted to:

$$|\Delta\theta| \le \frac{\pi}{4}. \tag{6.6}$$

For a given $\Delta\theta$, let $d(\Delta\theta) = \overline{C_1 C_2}$ denote the inter-vehicle distance. We define the maximum admissible spacing as

$$d_{\max}(\Delta\theta) \triangleq \max\{\, d \mid \text{the safe area remains closed and fully monitored} \,\}, \qquad |\Delta\theta| \le \frac{\pi}{4}. \tag{6.7}$$

Thus, the platoon configuration is feasible whenever:

$$d(\Delta\theta) \le d_{\max}(\Delta\theta), \qquad |\Delta\theta| \le \frac{\pi}{4}. \tag{6.8}$$

6.1.1 Same-heading configuration ($\Delta\theta = 0$)

We first consider the baseline case in which the two vehicles have identical headings (Fig. 6.1). In this case the overlap geometry is symmetric and the maximum admissible spacing corresponds to the boundary configuration in which the safe area is just closed. The distance between C_1 and C_2 at this boundary is:

$$d_{\max}(0) = r\sqrt{2}. \tag{6.9}$$

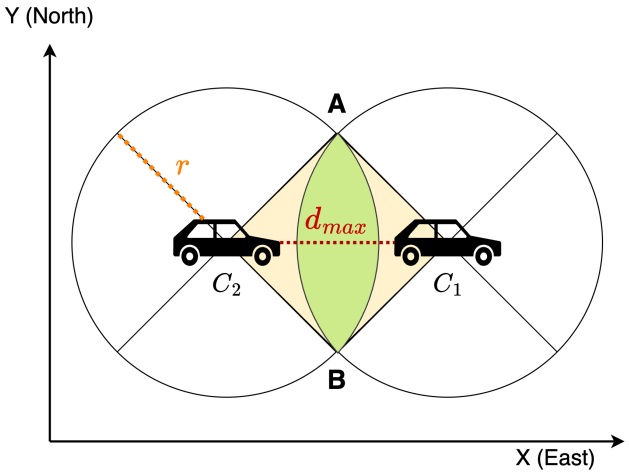

Fig. 6.1: Same-heading configuration ($\Delta\theta = 0$). The leader (C_1) and the follower (C_2) have the same heading. Each circle models the sensing region of radius r. The feasibility boundary corresponds to the maximum admissible inter-vehicle spacing $d_{\max} = \overline{C_1 C_2}$ for which the safe-area overlap remains closed and continuously monitored.

6.1.2 Nonzero heading mismatch ($\Delta\theta \neq 0$)

We now consider the general case in which the follower heading differs from the leader heading, i.e., $\Delta\theta \neq 0$ (Fig. 6.2). The two sensing circles still intersect at two points, A and B, but the symmetry of the same-heading case is lost.

To simplify the trigonometric derivation without altering the geometry, we apply a rigid rotation (change of reference frame) of the entire construction about one of the intersection points (Fig. 6.3). Specifically, for $\Delta\theta < 0$ the rotation is performed about A, while for $\Delta\theta > 0$ it is performed about B. Under this rigid rotation, all distances and angles are preserved; however, the second intersection point is mapped

to a new point, denoted by P (i.e., P is the rotated image of B, or vice versa depending on the sign of $\Delta\theta$). In the rotated frame, the feasibility-boundary quadrilateral is conveniently described by the vertices A, C_1, C_2, and P.

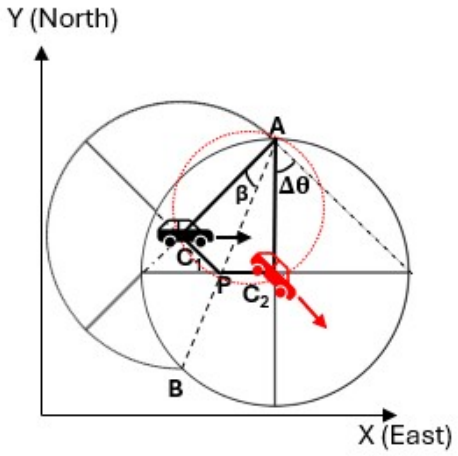

Fig. 6.2: Nonzero heading mismatch ($\Delta\theta \neq 0$). The sensing circles (radius r) centered at C_1 (leader) and C_2 (follower) intersect at points A and B. The relative heading mismatch $\Delta\theta = \theta_1 - \theta_2$ is shown at the intersection vertex. The angle at A is split by its bisector, yielding the half-angle β used in the analytical derivation.

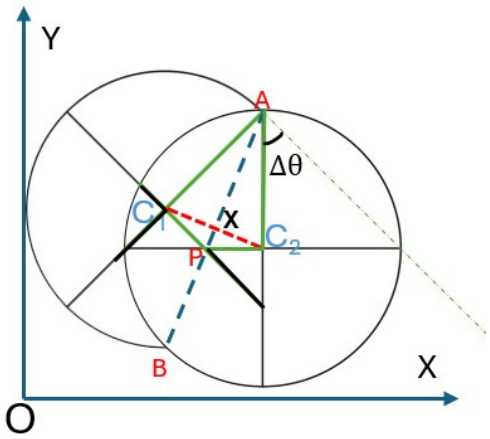

Fig. 6.3: Rigid rotation (change of reference frame) used to simplify the derivation. The construction is rotated about A for $\Delta\theta < 0$ and, analogously, about B for $\Delta\theta > 0$. Under this rotation, the second intersection point is mapped to P. Distances and angles are preserved. Point X denotes the intersection of the deltoid diagonals in the rotated frame.

Characterization of the feasibility-boundary polygon and angle β.

At the feasibility boundary (i.e., when $d = d_{\max}(\Delta\theta)$), the safe-area quadrilateral $A\,C_1\,P\,C_2$ forms a kite (deltoid). This follows from the equal radii of the sensing circles ($\overline{AC_1} = \overline{AC_2} = r$) and from the symmetry induced by the angle bisector at A in the rotated frame, which yields equal adjacent sides $\overline{PC_1} = \overline{PC_2}$.

The internal angle at A is reduced by the heading mismatch. In the same-heading case, this angle equals $\pi/2$; for $\Delta\theta \neq 0$, it becomes:

$$\angle C_2AC_1 = \frac{\pi}{2} - |\Delta\theta|. \tag{6.10}$$

Since the diagonal AP is the bisector of $\angle C_2AC_1$, the half-angle β used in the right-triangle decomposition is:

$$\beta = \frac{\frac{\pi}{2} - |\Delta\theta|}{2}. \tag{6.11}$$

Analytical derivation of $d_{\max}(\Delta\theta)$.

Consider the two triangles $\triangle AC_1P$ and $\triangle AC_2P$ in Fig. 6.4. By construction, $\overline{AC_1} = \overline{AC_2} = r$, side $\overline{AP}$ is common, and $\angle C_1AP = \angle C_2AP = \beta$. Hence, $\triangle AC_1P \cong \triangle AC_2P$, which implies $\overline{PC_1} = \overline{PC_2}$ (deltoid property).

In $\triangle AC_2P$, one can express the side $\overline{C_2P}$ as:

$$\overline{C_2P} = r\tan(\beta), \tag{6.12}$$

which is useful to characterize the deltoid sides (note that $\overline{C_2P}$ is *not* the inter-vehicle distance).

To obtain the inter-vehicle spacing, let X be the intersection of the deltoid diagonals (Fig. 6.4). In a kite (deltoid), one diagonal bisects the other; therefore X is the midpoint of segment $\overline{C_1C_2}$ and

$$\overline{C_2X} = \frac{\overline{C_1C_2}}{2} = \frac{d_{\max}(\Delta\theta)}{2}. \tag{6.13}$$

In the right triangle $\triangle AC_2X$, the hypotenuse is $\overline{AC_2} = r$ and the angle at A equals β, hence:

$$\overline{C_2X} = r\sin(\beta). \tag{6.14}$$

Combining (6.13) and (6.14) yields the closed-form maximum admissible spacing:

$$d_{\max}(\Delta\theta) = 2r\sin(\beta) = 2r\sin\left(\frac{\frac{\pi}{2} - |\Delta\theta|}{2}\right), \qquad |\Delta\theta| \leq \frac{\pi}{4}. \tag{6.15}$$

Equation (6.15) correctly reduces to (6.9) for $\Delta\theta = 0$, since $\sin(\pi/4) = \sqrt{2}/2$.

Boundary case ($|\Delta\theta| = \pi/4$).

For the extreme mismatch considered in this work, $|\Delta\theta| = \pi/4$, Eq. (6.15) gives:

$$d_{\max}\left(\frac{\pi}{4}\right) = 2r\sin\left(\frac{\pi}{8}\right). \tag{6.16}$$

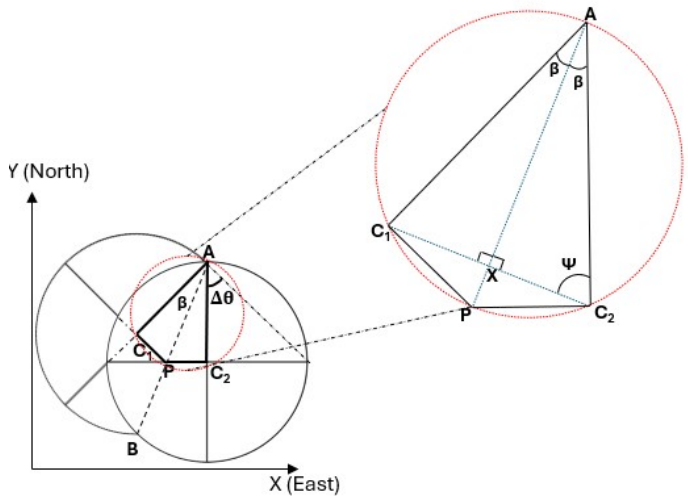

Fig. 6.4: Geometric quantities used in the derivation of $d_{\max}(\Delta\theta)$. Left: original configuration and rotated reference frame (with B mapped to P). Right: zoom on the deltoid $A\,C_1\,P\,C_2$ at the feasibility boundary. The diagonals intersect at X. The half-angle at A is $\beta = \left(\frac{\pi}{2} - |\Delta\theta|\right)/2$, yielding $d_{\max}(\Delta\theta) = 2\,\overline{C_2X} = 2r\sin(\beta)$.

6.1.3 Implications for spacing policies

Let v denote the vehicle speed, h the headway time, and r the sensing radius. In a platoon, the desired intra-platoon gap can be either set as a constant spacing [128, 129] or determined by a constant-headway policy [130]:

$$d_{gap} = \begin{cases} h\,v, & \text{constant headway,} \\ d_0, & \text{constant spacing.} \end{cases} \tag{6.17}$$

The ToD-assisted platoon can be maintained provided that:

$$d_{\max}(\Delta\theta) \geq d_{gap}, \tag{6.18}$$

where $d_{\max}(\Delta\theta)$ is given by Eq. (6.15). If $\Delta\theta$ is not estimated/controlled, a conservative sufficient condition is obtained by using the worst-case bound:

$$d_{gap} \leq \min_{|\Delta\theta| \leq \pi/4} d_{\max}(\Delta\theta) = 2r\sin\left(\frac{\pi}{8}\right). \tag{6.19}$$

References

1. C. Campolo, A. Molinaro, B. Pizzimenti, M. Schrapel, A. Vinel, and D. M. Zappalà, "Understanding vulnerable road users awareness messages generation dynamics: The case of e-bikes," in *2024 IEEE Vehicular Networking Conference (VNC)*, pp. 336–342, 2024.
2. M. Schrapel, T. Graf, D. M. Zappalà, C. Campolo, A. Molinaro, and A. Vinel, "When the e-bike takes over: Speed precision and perception of cruise control for cyclists," in *Proceedings of the 16th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*, AutomotiveUI '24, (New York, NY, USA), p. 272–282, Association for Computing Machinery, 2024.
3. J. Shuttleworth, "Sae standards news: J3016 automated-driving graphic update." `https://www.sae.org/news/2019/01/sae-updates-j3016-automated-driving-graphic`, 2019. Accessed: 2026-01-02.
4. D. Gioffrè, B. Pizzimenti, H. Rashid, D. M. Zappalà, C. Campolo, A. Molinaro, and G. Ruggeri, "Traffic condition dissemination through 5g-based vehicle-to-network-to-everything communications," in *2025 IEEE Conference on Standards for Communications and Networking (CSCN)*, pp. 1–1, 2025.
5. W. H. Organization, "Road traffic injuries." `https://www.who.int/news-room/fact-sheets/detail/road-traffic-injuries`, 2023. Accessed: 2026-01-02.
6. INRIX, "Global traffic scorecard." `https://inrix.com/scorecard/`, 2024. Accessed: 2026-01-02.
7. E. R&I, "CCAM Association." `https://www.ccam.eu/`, 2024. Accessed: 2026-01-02.
8. "ETSI TS 103 300-1 v2.1.1. Intelligent Transport System (ITS); Vulnerable Road Users (VRU) awareness; Part 1: Use Case definition," September 2019.
9. IEEE, "IEEE Standard for Information technology– Local and metropolitan area networks– Specific requirements– Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 6: Wireless Access in Vehicular Environments," July 2010.
10. D. Eckhoff, N. Sofra, and R. German, "A performance study of cooperative awareness in ETSI ITS G5 and IEEE WAVE," in *2013 10th Annual Conference on Wireless On-demand Network Systems and Services (WONS)*, pp. 196–200, 2013.

11. B. Y. Yacheur, T. Ahmed, and M. Mosbah, "Analysis and comparison of ieee 802.11 p and ieee 802.11 bd," in *International Workshop on Communication Technologies for Vehicles*, pp. 55–65, Springer, 2020.
12. 3rd Generation Partnership Project (3GPP), "3gpp release 14 description; tr 21.914." https://www.3gpp.org/specifications-technologies/releases/release-14, March 2017. Accessed: 2026-01-02.
13. 3rd Generation Partnership Project (3GPP), "3gpp release 15 description; tr 21.915." https://www.3gpp.org/specifications-technologies/releases/release-15, April 2019. Accessed: 2026-01-02.
14. T. 3rd Generation Partnership Project (3GPP), "3GPP Release 16 Description; TR 21.916." https://www.3gpp.org/specifications-technologies/releases/release-16, July 2021.
15. T. 3rd Generation Partnership Project (3GPP), "3GPP Release 17 Description; TR 21.917." https://www.3gpp.org/specifications-technologies/releases/release-17, June 2022.
16. T. 3rd Generation Partnership Project (3GPP), "3GPP Release 18 Description; TR 21.918." https://www.3gpp.org/specifications-technologies/releases/release-18, December 2023.
17. "ETSI TS 103 300-2 v2.1.1. Intelligent Transport System (ITS); Vulnerable Road Users (VRU) awareness; Part 2: Functional Architecture and Requirements definition," May 2020.
18. "ETSI TS 103 300-3 v2.2.1. Intelligent Transport System (ITS); Vulnerable Road Users (VRU) awareness; Part 3: Specification of VRU awareness basic service," February 2023.
19. "ETSI EN 302 637-2 v1.3.1 Intelligent Transport System (ITS); Vehicular Communications; Basic Set of Applications; Part 2: Specification of Cooperative Awareness Basic Service," September 2014.
20. A. Banks and R. Gupta, "Mqtt version 3.1. 1," *OASIS standard*, vol. 29, p. 89, 2014.
21. E. Commission, "Eu road fatalities drop by 3% in 2024, but progress remains slow." https://transport.ec.europa.eu/news-events/news/eu-road-fatalities-drop-3-2024-progress-remains-slow-2025-03-18_en, March 2025. Accessed: 2026-01-02.
22. C. Zoghlami, R. Kacimi, and R. Dhaou, "5g-enabled v2x communications for vulnerable road users safety applications: A review," *Wireless Networks*, vol. 29, no. 3, pp. 1237–1267, 2023.
23. A. Boukerche and M. Sha, "Design guidelines on deep learning–based pedestrian detection methods for supporting autonomous vehicles," vol. 54, no. 6, 2021.
24. S. El Hamdani, N. Benamar, and M. Younis, "Pedestrian support in intelligent transportation systems: Challenges, solutions and open issues," *Transportation Research Part C: Emerging Technologies*, vol. 121, p. 102856, 2020.
25. S. K. Maurya and A. Choudhary, "Deep learning based vulnerable road user detection and collision avoidance," in *2018 IEEE International Conference on Vehicular Electronics and Safety (ICVES)*, pp. 1–6, 2018.

26. H. Chen, Y. Liu, C. Hu, and X. Zhang, "Vulnerable road user trajectory prediction for autonomous driving using a data-driven integrated approach," *IEEE Transactions on Intelligent Transportation Systems*, vol. 24, no. 7, pp. 7306–7317, 2023.
27. P. Merdrignac, O. Shagdar, I. B. Jemaa, and F. Nashashibi, "Study on perception and communication systems for safety of vulnerable road users," in *2015 IEEE 18th International Conference on Intelligent Transportation Systems*, pp. 1876–1881, IEEE, 2015.
28. P. Merdrignac, O. Shagdar, and F. Nashashibi, "Fusion of perception and v2p communication systems for the safety of vulnerable road users," *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 7, pp. 1740–1751, 2017.
29. E. N. C. A. P. E. NCAP), "Euro ncap for engineers - protocols," tech. rep., European New Car Assessment Programme (Euro NCAP), Oct. 2025. Accessed: 2026-01-02.
30. M. Bachmann, M. Morold, S. Engel, J. Götz, and K. David, "Camera vs. cooperative vru collision avoidance," in *2020 IEEE 91st Vehicular Technology Conference (VTC2020-Spring)*, pp. 1–5, 2020.
31. E. N. C. A. P. E. NCAP), "Test protocol - aeb vru systems (version 2.0.4)," tech. rep., European New Car Assessment Programme (Euro NCAP), Oct. 2019. Accessed: 2026-01-02.
32. M. Morold, Q.-H. Nguyen, M. Bachmann, K. David, and F. Dressler, "Requirements on delay of vru context detection for cooperative collision avoidance," in *2020 IEEE 92nd Vehicular Technology Conference (VTC2020-Fall)*, pp. 1–5, 2020.
33. P. Wang, M. Zhou, and Z. Ding, "A vru collision warning system with kalman-filter-based positioning accuracy improvement," in *2021 IEEE International Conference on Information Communication and Software Engineering (ICICSE)*, pp. 191–198, 2021.
34. M. Morold, M. Bachmann, and K. David, "Toward context awareness for cooperative vulnerable road user collision avoidance: Incorporating related contextual information," *IEEE Vehicular Technology Magazine*, vol. 17, no. 3, pp. 75–83, 2022.
35. A. Flach, A. Q. Memon, S. L. Lau, and K. David, "Pedestrian movement recognition for radio based collision avoidance: A performance analysis," in *2011 IEEE 73rd Vehicular Technology Conference (VTC Spring)*, pp. 1–5, 2011.
36. L. Lusvarghi, C. A. Grazia, M. Klapez, M. Casoni, and M. L. Merani, "Awareness messages by vulnerable road users and vehicles: Field tests via lte-v2x," *IEEE Transactions on Intelligent Vehicles*, vol. 8, no. 10, pp. 4418–4433, 2023.
37. A. Genovese, M. Rapelli, F. Raviglione, and C. Casetti, "Etsi-compliant protection of connected vulnerable road users: Simulation and field trial," in *2024 IEEE Vehicular Networking Conference (VNC)*, pp. 343–350, IEEE, 2024.
38. N. Ostendorf, K. Garlichs, and L. C. Wolf, "Analysis of vulnerable road user awareness message generation rules for pedelecs," in *2024 IEEE Vehicular Networking Conference (VNC)*, pp. 321–328, IEEE, 2024.
39. M. Schrapel, M. C. Anfang, and A. Vinel, "Poster: Analyzing vam sampling rates for e-bikes in vanets," in *2025 IEEE Vehicular Networking Conference (VNC)*, pp. 1–2, 2025.

40. S. Lobo, A. Festag, and C. Facchi, "Enhancing the safety of vulnerable road users: Messaging protocols for v2x communication," in *2022 IEEE 96th Vehicular Technology Conference (VTC2022-Fall)*, pp. 1–7, IEEE, 2022.
41. S. Lobo, L. B. Da Silva, and C. Facchi, "To cluster or not to cluster: A vru clustering based on v2x communication," in *2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC)*, pp. 2218–2225, IEEE, 2023.
42. M. Rupp and L. Wischhof, "Evaluation of the effectiveness of vulnerable road user clustering in c-v2x systems," in *2023 IEEE International Conference on Omni-layer Intelligent Systems (COINS)*, pp. 1–5, IEEE, 2023.
43. M. Bieshaar, S. Zernetsch, M. Depping, B. Sick, and K. Doll, "Cooperative starting intention detection of cyclists based on smart devices and infrastructure," in *2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*, pp. 1–8, 2017.
44. M. Segata, R. Vijeikis, and R. L. Cigno, "Communication-based collision avoidance between vulnerable road users and cars," in *2017 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pp. 565–570, 2017.
45. P. A. Plazier, G. Weitkamp, and A. E. van den Berg, ""cycling was never so easy!" an analysis of e-bike commuters' motives, travel behaviour and experiences using gps-tracking and interviews," *Journal of transport geography*, vol. 65, pp. 25–34, 2017.
46. P. Huertas-Leyva, M. Dozza, and N. Baldanzini, "Investigating cycling kinematics and braking maneuvers in the real world: e-bikes make cyclists move faster, brake harder, and experience new conflicts," *Transportation research part F: traffic psychology and behaviour*, vol. 54, pp. 211–222, 2018.
47. K. Schleinitz, T. Petzoldt, L. Franke-Bartholdt, J. Krems, and T. Gehlert, "The german naturalistic cycling study–comparing cycling speed of riders of different e-bikes and conventional bicycles," *Safety science*, vol. 92, pp. 290–297, 2017.
48. O. A. Molina, E. Ronelöv, K. Boustedt, J. Blidkvist, and A. Vinel, "Protection of vulnerable road users using hybrid vehicular networks," in *2022 IEEE International Conference on Vehicular Electronics and Safety (ICVES)*, pp. 1–6, 2022.
49. A. Clérigo, M. Schrapel, P. Rito, S. Sargento, and A. Vinel, "Microservice-based architecture for enhancing road safety with support for low-latency services," in *NOMS 2025-2025 IEEE Network Operations and Management Symposium*, pp. 1–7, 2025.
50. D. H. De La Iglesia, J. F. De Paz, G. Villarrubia González, A. L. Barriuso, J. Bajo, and J. M. Corchado, "Increasing the intensity over time of an electric-assist bike based on the user and route: The bike becomes the gym," *Sensors*, vol. 18, no. 1, p. 220, 2018.
51. J. A. Afonso, F. J. Rodrigues, D. D. R. Pedrosa, and J. L. Afonso, "Automatic control of cycling effort using electric bicycles and mobile devices," 2015.
52. D. Meyer, M. Körber, V. Senner, and M. Tomizuka, "Regulating the heart rate of human–electric hybrid vehicle riders under energy consumption constraints using an optimal control approach," *IEEE Transactions on Control Systems Technology*, vol. 27, no. 5, pp. 2125–2138, 2018.

53. M. Corno, P. Giani, M. Tanelli, and S. M. Savaresi, "Human-in-the-loop bicycle control via active heart rate regulation," *IEEE Transactions on Control Systems Technology*, vol. 23, no. 3, pp. 1029–1040, 2014.
54. M. G. Wilson, A. M. Lane, C. J. Beedie, and A. Farooq, "Influence of accurate and inaccurate 'split-time'feedback upon 10-mile time trial cycling performance," *European journal of applied physiology*, vol. 112, no. 1, pp. 231–236, 2012.
55. R. H. Morton, "Deception by manipulating the clock calibration influences cycle ergometer endurance time in males," *Journal of Science and Medicine in Sport*, vol. 12, no. 2, pp. 332–337, 2009.
56. A. Matviienko, D. Mehmedovic, F. Müller, and M. Mühlhäuser, "" baby, you can ride my bike" exploring maneuver indications of self-driving bicycles using a tandem simulator," *Proceedings of the ACM on Human-Computer Interaction*, vol. 6, no. MHCI, pp. 1–21, 2022.
57. M.-C. H. Oczko, L. Stratmann, M. Franke, J. Heinovski, D. S. Buse, F. Klingler, and F. Dressler, "Integrating haptic signals with v2x-based safety systems for vulnerable road users," in *2020 International Conference on Computing, Networking and Communications (ICNC)*, pp. 692–697, IEEE, 2020.
58. S. A. Useche, J. Gene-Morales, F. W. Siebert, F. Alonso, and L. Montoro, ""not as safe as i believed": differences in perceived and self-reported cycling behavior between riders and non-riders," *Sustainability*, vol. 13, no. 4, p. 1614, 2021.
59. S. Boufous, J. Hatfield, and R. Grzebieta, "The impact of environmental factors on cycling speed on shared paths," *Accident Analysis & Prevention*, vol. 110, pp. 171–176, 2018.
60. J. Andres, T. Kari, J. Von Kaenel, and F. Mueller, "Co-riding with my ebike to get green lights," in *Proceedings of the 2019 on Designing Interactive Systems Conference*, pp. 1251–1263, 2019.
61. S. G. Hart and L. E. Staveland, "Development of nasa-tlx (task load index): Results of empirical and theoretical research," in *Advances in psychology*, vol. 52, pp. 139–183, Elsevier, 1988.
62. T. Franke, C. Attig, and D. Wessel, "A personal resource for technology interaction: development and validation of the affinity for technology interaction (ati) scale," *International Journal of Human–Computer Interaction*, vol. 35, no. 6, pp. 456–467, 2019.
63. P. W. Jordan, B. Thomas, I. L. McClelland, and B. Weerdmeester, *Usability evaluation in industry*. CRC press, 1996.
64. "ETSI EN 302 663 v1.3.0. Intelligent Transport System (ITS); ITS-G5 Access layer specification for Intelligent Transport Systems operating in the 5 GHz frequency band," October 2019.
65. SAE, "Taxonomy and definitions for terms related to driving automation systems for on-road motor vehicles." `https://www.sae.org/standards/content/j3016_202104/`, 2021. Accessed: 2026-01-02.
66. J. Zhao, W. Zhao, B. Deng, Z. Wang, F. Zhang, W. Zheng, W. Cao, J. Nan, Y. Lian, and A. F. Burke, "Autonomous driving system: A comprehensive survey," *Expert Systems with Applications*, vol. 242, p. 122836, 2024.

67. A. Balasubramaniam and S. Pasricha, "Object detection in autonomous vehicles: Status and open challenges," *arXiv preprint arXiv:2201.07706*, 2022.
68. A. Bazzi, A. O. Berthet, C. Campolo, B. M. Masini, A. Molinaro, and A. Zanella, "On the design of sidelink for cellular v2x: A literature review and outlook for future," *IEEE Access*, vol. 9, pp. 97953–97980, 2021.
69. H. Zhou, W. Xu, J. Chen, and W. Wang, "Evolutionary v2x technologies toward the internet of vehicles: Challenges and opportunities," *Proceedings of the IEEE*, vol. 108, no. 2, pp. 308–323, 2020.
70. S. Lu, N. Ammar, A. Ganlath, H. Wang, and W. Shi, "A comparison of end-to-end architectures for connected vehicles," in *2022 Fifth International Conference on Connected and Autonomous Driving (MetroCAD)*, pp. 72–80, IEEE, 2022.
71. O. Briante, C. Campolo, A. Iera, A. Molinaro, S. Y. Paratore, and G. Ruggeri, "Supporting augmented floating car data through smartphone-based crowd-sensing," *Vehicular Communications*, vol. 1, no. 4, pp. 181–196, 2014.
72. D. Budimir, N. Jelušić, and M. Perić, "Floating car data technology," *Pomorstvo*, vol. 33, no. 1, pp. 22–32, 2019.
73. D. A. Kountche, F. Raissi, M. R. Rakotondravelona, E. Bonetto, D. Brevi, A. Martin, O. Otaegui, and G. Velez, "Monetisation of and access to in-vehicle data and resources: The 5gmeta approach," *arXiv preprint arXiv:2208.11335*, 2022.
74. G. Naik, B. Choudhury, and J.-M. Park, "Ieee 802.11 bd & 5g nr v2x: Evolution of radio access technologies for v2x communications," *IEEE access*, vol. 7, pp. 70169–70184, 2019.
75. 5GAA, "Vehicle-to-network-to-everything (v2n2x) communications: Architecture, solution blueprint, and use case implementation examples." `https://5gaa.org/content/uploads/2024/06/v2n2x-communications-architecture-solution-blueprint-and-use-case-implementation.pdf`, 2025. Accessed: 2026-01-02.
76. 5GAA, "Marketplace." `https://5gaa.org/marketplace/`, 2025. Accessed: 2026-01-02.
77. M. Klapež, C. A. Grazia, and M. Casoni, "Experimental evaluation of ieee 802.11 p in high-speed trials for safety-related applications," *IEEE Transactions on Vehicular Technology*, vol. 70, no. 11, pp. 11538–11553, 2021.
78. M. Klapez, C. A. Grazia, and M. Casoni, "Application-level performance of ieee 802.11 p in safety-related v2x field trials," *IEEE Internet of Things Journal*, vol. 7, no. 5, pp. 3850–3860, 2020.
79. F. Raviglione, M. Malinverno, and C. Casetti, "Open source platform for ieee 802.11 p nics evaluation," in *2019 IEEE 20th International Symposium on" A World of Wireless, Mobile and Multimedia Networks"(WoWMoM)*, pp. 1–3, IEEE, 2019.
80. M. Rapelli, F. Raviglione, and C. Casetti, "Oscar: An etsi-compliant c-its stack for field-testing with embedded hardware devices," in *2024 22nd Mediterranean Communication and Computer Networking Conference (MedComNet)*, pp. 1–4, IEEE, 2024.

81. J. Zheng, T. H. Luan, Y. Zhang, R. Li, Y. Hui, L. Gao, and M. Dong, "Data synchronization in vehicular digital twin network: A game theoretic approach," *IEEE Transactions on Wireless Communications*, vol. 22, no. 11, pp. 7635–7647, 2023.
82. J. He, Z. Tang, X. Fu, S. Leng, F. Wu, K. Huang, J. Huang, J. Zhang, Y. Zhang, A. Radford, *et al.*, "Cooperative connected autonomous vehicles (cav): Research, applications and challenges," in *2019 IEEE 27th International Conference on Network Protocols (ICNP)*, pp. 1–6, IEEE, 2019.
83. B. Coll-Perales, M. Lucas-Estañ, O. Altintas, J. Gozalvez, M. Khan, and M. Sepulcre, "V2X Service Provisioning with 5G V2N2V Communications with Cross-Stakeholder Information Sharing," in *IEEE Vehicular Networking Conference (VNC)*, pp. 109–112, 2024.
84. S. Chouali, A. Boukerche, and A. Mostefaoui, "Towards a formal analysis of mqtt protocol in the context of communicating vehicles," in *Proceedings of the 15th ACM International Symposium on Mobility Management and Wireless Access*, pp. 129–136, 2017.
85. B. Ganguly and A. Chatterjee, "Mqtt protocol based extensive smart motor control for electric vehicular application," in *2020 IEEE 7th Uttar Pradesh Section International Conference on Electrical, Electronics and Computer Engineering (UPCON)*, pp. 1–5, IEEE, 2020.
86. C. Campolo, A. Molinaro, B. Pizzimenti, and D. M. Zappalà, "Poster: Multi-rat multi-protocol obus for vehicle-to-everything communications," in *2025 IEEE Vehicular Networking Conference (VNC)*, pp. 1–2, 2025.
87. C. Campolo, G. Genovese, A. Molinaro, B. Pizzimenti, G. Ruggeri, and D. M. Zappalà, "An edge-based digital twin framework for connected and autonomous vehicles: Design and evaluation," *IEEE Access*, vol. 12, pp. 46290–46303, 2024.
88. C. Campolo, A. Molinaro, B. Pizzimenti, G. Singh, and D. M. Zappalà, "Digital twin-assisted context-aware object detection for connected and automated vehicles," in *2025 IEEE Vehicular Networking Conference (VNC)*, pp. 1–4, 2025.
89. C. Campolo, N. Mammone, A. Molinaro, B. Pizzimenti, G. Singh, and D. M. Zappalà, "Poster: Digital twins for personalized and safer automated driving," in *2024 IEEE Vehicular Networking Conference (VNC)*, pp. 277–278, IEEE, 2024.
90. M. A. R. Alif, "Yolov11 for vehicle detection: Advancements, performance, and applications in intelligent transportation systems," *arXiv preprint arXiv:2410.22898*, 2024.
91. C. Consortium, "Coco - common objects in context." `https://cocodataset.org/#home`. Accessed: 2026-01-02.
92. X. Yin *et al.*, "Fine-tuning and visualization of convolutional neural networks," in *2017 12th IEEE Conference on Industrial Electronics and Applications (ICIEA)*, pp. 1310–1315.
93. W. Liu *et al.*, "Edge-assisted vehicle mobility prediction to support V2X communications," *IEEE Transactions on Vehicular Technology*, vol. 68, no. 10, pp. 10227–10238, 2019.
94. H. Ni and the NCNN contributors, "Ncnn." `https://github.com/Tencent/ncnn`. Accessed: 2026-01-02.
95. U. Inc., "Yolo performance metrics." `https://docs.ultralytics.com/guides/yolo-performance-metrics/#class-wise-metrics`. Accessed: 2026-01-02.
96. U. Inc., "Insights on model evaluation and fine-tuning." `https://docs.ultralytics.com/guides/model-evaluation-insights/`. Accessed: 2026-01-02.

97. K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.
98. C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," in *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2818–2826, 2016.
99. M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "Mobilenetv2: Inverted residuals and linear bottlenecks," in *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4510–4520, 2018.
100. W. Foundation, "Wireshark: The world's foremost network protocol analyzer." `https://www.wireshark.org/`. Accessed: 2026-01-02.
101. "ETSI EN 302 637-3 v1.3.1 Intelligent Transport System (ITS); Vehicular Communications; Basic Set of Applications; Part 3: Specification of Decentralized Environmental Notification Basic Service."
102. M. C. Lucas-Estañ, B. Coll-Perales, T. Shimizu, J. Gozalvez, T. Higuchi, S. Avedisov, O. Altintas, and M. Sepulcre, "An analytical latency model and evaluation of the capacity of 5g nr to support v2x services using v2n2v communications," *IEEE Transactions on Vehicular Technology*, vol. 72, no. 2, pp. 2293–2306, 2023.
103. T. Networks, "Rutos." `https://wiki.teltonika-networks.com/view/Teltonika-Networks_Operating_System_-_RutOS`, 2025. Accessed: 2026-01-02.
104. OpenWrt, "Technical reference." `https://openwrt.org/docs/techref/start`, 2025. Accessed: 2026-01-02.
105. E. Foundation, "Mosquitto." `https://mosquitto.org/`, 2025. Accessed: 2026-01-02.
106. O. Amador, M. Aramrattana, and A. Vinel, "A survey on remote operation of road vehicles," *IEEE Access*, vol. 10, pp. 130135–130154, 2022.
107. S. B. Kamtam, Q. Lu, F. Bouali, O. C. Haas, and S. Birrell, "Network latency in teleoperation of connected and autonomous vehicles: A review of trends, challenges, and mitigation strategies," *Sensors (Basel, Switzerland)*, vol. 24, no. 12, p. 3957, 2024.
108. P. A. Lopez, M. Behrisch, L. Bieker-Walz, J. Erdmann, Y.-P. Flötteröd, R. Hilbrich, L. Lücken, J. Rummel, P. Wagner, and E. Wießner, "Microscopic traffic simulation using sumo," in *2018 21st international conference on intelligent transportation systems (ITSC)*, pp. 2575–2582, 2018.
109. 5GAA, "Tele-operated driving (tod): System requirements analysis and architecture." `https://5gaa.org/tele-operated-driving-tod-system-requirements-analysis-and-architecture/`, 2021. Accessed: 2026-01-08.
110. S. Neumeier, V. Bajpai, M. Neumeier, C. Facchi, and J. Ott, "Data rate reduction for video streams in teleoperated driving," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 10, pp. 19145–19160, 2022.
111. D. Peled, A. Shmilovici, and O. Hadar, "Automative video compression for remote driving via safety considerations," in *IEEE Conference on Standards for Communications and Networking (CSCN)*, pp. 54–58, 2023.

112. N. Adwani, K. Silvestrini-Cordero, R. Rojas-Cessa, T. Han, and C. Wang, "A special-purpose video streaming codec for internet-based remote driving," in *IEEE International Conference on Advanced Intelligent Mechatronics (AIM)*, pp. 1109–1116, 2024.
113. M. Hussain, "Yolo-v1 to yolo-v8, the rise of yolo and its complementary nature toward digital manufacturing and industrial defect detection," *Machines*, vol. 11, no. 7, 2023.
114. J. Canny, "A computational approach to edge detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-8, no. 6, pp. 679–698, 1986.
115. M. Hofbauer, C. B. Kuhn, M. Khlifi, G. Petrovic, and E. Steinbach, "Traffic-aware multi-view video stream adaptation for teleoperated driving," in *IEEE 95th Vehicular Technology Conference: (VTC2022-Spring)*, pp. 1–7, 2022.
116. A. Schimpe, S. Hoffmann, and F. Diermeyer, "Adaptive video configuration and bitrate allocation for teleoperated vehicles," in *IEEE Intelligent Vehicles Symposium Workshops (IV Workshops)*, pp. 148–153, 2021.
117. H. Schippers, T. Gebauer, K. Heimann, and C. Wietfeld, "Rise: Multi-link proactive low-latency video streaming for teleoperation in fading channels," in *IEEE 101st Vehicular Technology Conference (VTC2025-Spring)*, pp. 1–6.
118. B. Turkovic, R. van de Vlasakker, S. Potnuru, R. S. Schwartz, L. D'Acunto, O. Held, S. Karasek, R. Kusumakar, A. Cremona, and R. Fardanian, "On-demand network qos adaptation for automated tele-operated driving," in *Proceedings of the 3rd International Workshop on Testing Distributed Internet of Things Systems*, TDIS '25, (New York, NY, USA), p. 19–22, Association for Computing Machinery, 2025.
119. P. Kremer, N. Nourani-Vatani, and S. Park, "A digital twin for teleoperation of vehicles in urban environments," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 12521–12527, IEEE, 2023.
120. S. Bhadra and S. Park, "Criticality-aware load-balanced scheduling algorithm for teleoperated vehicles controlled over vehicle-to-infrastructure network," in *IEEE 100th Vehicular Technology Conference (VTC2024-Fall)*, pp. 1–6, 2024.
121. N. Slamnik-Kriještorac, W. Vandenberghe, X. Limani, E. Oostendorp, E. de Groote, V. Maglogiannis, D. Naudts, P.-P. Schackmann, R. Kusumakar, K. Kural, *et al.*, "5G-enhanced Teleoperation in Real-Life Port Environments: Lessons Learned from the 5G-Blueprint Project," in *Joint European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit)*, pp. 973–978, 2024.
122. J. M. Marquez-Barja, S. Hadiwardoyo, B. Lannoo, W. Vandenberghe, E. Kenis, L. Deckers, M. C. Campodonico, K. dos Santos, R. Kusumakar, M. Klepper, *et al.*, "Enhanced teleoperated transport and logistics: A 5G cross-border use case," in *IEEE Joint European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit)*, pp. 229–234, 2021.
123. F. J. Jiang, M. Al-Janabi, T. Bolin, K. H. Johansson, and J. Mårtensson, "SVEA: an experimental testbed for evaluating V2X use-cases," in *IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*, pp. 3484–3489, 2022.
124. 5G Automotive Association (5GAA), "C-V2X Use Cases and Service Level Requirements, Volume II," tech. rep., 5GAA, 2023. Technical Report.

125. G. Kakkavas, K. N. Nyarko, C. Lahoud, D. Kühnert, P. Küffner, M. Gabriel, S. Ehsanfar, M. Diamanti, V. Karyotis, K. Mößner, *et al.*, "Teleoperated support for remote driving over 5G mobile communications," in *2022 IEEE International Mediterranean Conference on Communications and Networking (MeditCom)*, pp. 280–285, IEEE, 2022.
126. D. Jia, K. Lu, J. Wang, X. Zhang, and X. Shen, "A survey on platoon-based vehicular cyber-physical systems," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 1, pp. 263–284, 2015.
127. C. M. Lentisco, C. Campolo, A. Molinaro, J. Martín-Pérez, and L. Bellido, "Digital twin-assisted radio resource allocation for tele-operated driving," in *2024 20th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*.
128. R. Rajamani, H.-S. Tan, B. K. Law, and W.-B. Zhang, "Demonstration of integrated longitudinal and lateral control for the operation of automated vehicles in platoons," *IEEE Transactions on Control Systems Technology*, vol. 8, no. 4, pp. 695–708, 2000.
129. M. Segata, "Platooning in SUMO: an open source implementation," in *SUMO User Conference*, pp. 51–62, 2017.
130. J. Ploeg, B. T. Scheepers, E. Van Nunen, N. Van de Wouw, and H. Nijmeijer, "Design and experimental evaluation of cooperative adaptive cruise control," in *14th International IEEE Conference on Intelligent Transportation Systems (ITSC)*, pp. 260–265, 2011.
131. C. Hidalgo, A. Arizala, N. Iturbe-Olleta, A. Brazalez, A. Zubizarreta, E. Asua, and J. P. Rastelli, "A Cooperative Driving Framework For Platooning Using V2X Messages In Urban Environments," *IEEE Transactions on Vehicular Technology*, 2025.
132. C. M. R. Carletti, C. Casetti, J. Härri, and F. Risso, "Platoon-Local Dynamic Map: Hierarchical cooperative perception for platooning vehicles," *IEEE Transactions on Vehicular Technology*, 2025.
133. V. M. Rafael, J. Caceres, A. H. Abel, A. Montoya, J. Guevara, G. B. Jarelh, and J. Talavera, "Development of a low-cost teleoperated explorer robot (TXRob)," *International Journal of Advanced Computer Science and Applications*, vol. 13, no. 7, 2022.
134. S. Rameshselvakumar, T. Arjun, T. M. Kumar, and C. N. Kumar, "Teleoperated campus utility vehicle using SDAS based on CAN," in *2021 International Conference on Advancements in Electrical, Electronics, Communication, Computing and Automation (ICAECA)*, pp. 1–6, IEEE, 2021.
135. M. Lucas-Estañ, B. Coll-Perales, M. Khan, J. Gozalvez, S. Avedisov, O. Altintas, and M. Sepulcre, "5G Network Architecture and Configuration Choices to Support Teleoperated Driving at Scale," in *IEEE 100th Vehicular Technology Conference (VTC2024-Fall)*, pp. 1–6, 2024.
136. A. Corsaro, L. Cominardi, O. Hecart, G. Baldoni, J. E. P. Avital, J. Loudet, C. Guimares, M. Ilyin, and D. Bannov, "Zenoh: Unifying communication, storage and computation from the cloud to the microcontroller," in *2023 26th Euromicro Conference on Digital System Design (DSD)*, pp. 422–428, IEEE, 2023.
137. J. N. Hark, J. Pfeifer, J. Kouril, C. Schahn, B. Schäufele, and I. Radusch, "Extending teleoperated driving: Advanced design for versatile environments and communications," in *2025 IEEE Vehicular Networking Conference (VNC)*, pp. 1–8, IEEE, 2025.