Cooperation and Assistance for Autonomous Driving under Realistic Conditions

Mikkel Skov Maarssø* , Giovanni Lucente † , Julian Schindler ‡

Abstract: Although automated driving systems (ADS) increasingly operate in public traffic, cooperation between automated vehicles and humans must not be forgotten in light of improving safety and efficiency. However, deploying such cooperative functions is challenging in real traffic situations. This paper presents developments from the German National project Managed Automated Driving (MAD) Urban using the open-source Eclipse-Automated Driving Open Research (ADORe®) stack based on ROS2. ADORe® enables coordinated decision making through two complementary modes: Single-Agent Autonomous Driving (SAAD) and Multi-Agent Autonomous Driving (MAAD). While ADORe SAAD provides robust onboard automated driving, capable of Vehicle-to-Everything (V2X) cooperation and remote operations, ADORe MAAD is the version designed to run on the infrastructure, allowing to jointly predict and plan trajectories for multiple vehicles, by using infrastructure sensing and Infrastructure-to-Everything (I2X) communication. Together, these components form a cooperative driving framework designed to function reliably under realistic conditions, which have been deployed in real traffic situations.

Keywords: Autonomous Driving, Cooperative Driving, Open-Source ADS, Remote Operations, Intelligent Infrastructure

1 Introduction

With the continuous expansion of Automated Driving Systems (ADS) introduced in public traffic, research on ways to increase safety and efficiency has remained a priority, although currently outperformed by deployment and sovereignty demands [1]. Cooperation between connected and automated vehicles (CAVs) and humans is seen as an avenue for acquiring safety and efficiency. However, the practical aspects of actually deploying cooperative solutions in real-world traffic pose a large number of challenges. Challenges for enabling cooperation between CAV systems themselves and between CAV systems and humans under realistic conditions include unreliable data connection, complexity in multi-agent planning and prediction, and reliability of available infrastructure. This paper aims to introduce the work done throughout the MAD-Urban research project to address these problems in DLRs' ADORe® software stack [2] [3].

*Mikkel Skov Maarssø at German Aerospace Center (DLR) (e-mail: mikkel.maarssoe@dlr.de).
†Giovanni Lucente at German Aerospace Center (DLR) (e-mail: giovanni.lucente@dlr.de).
‡Julian Schindler at German Aerospace Center (DLR) (e-mail: julian.schindler@dlr.de).

2 Related Work

Several surveys provide a comprehensive overview of cooperative intersection systems that integrate V2X communication with infrastructure support for perception, planning, and control [4, 5, 6]. The majority of existing approaches are primarily evaluated in simulation and span a broad spectrum of concepts, including reservation-based intersection management, centralized optimization, and learning-based multi-agent planning for connected and automated vehicles, typically under idealized assumptions on sensing and communication [4]. In terms of the SAE J3216 taxonomy [7], many of these methods are conceptually aligned with higher cooperation classes (e.g., agreement-seeking and prescriptive cooperation), although they are rarely demonstrated beyond simulation environments.

In contrast, real-world deployments of cooperative intersection systems remain limited. Current implementations predominantly focus on information and intent sharing with standardized messages such as Cooperative Awareness (CAM), Decentralised Environmental Notification (DENM), Signal Phase and Timing (SPATEM), mapping (MAPEM), or collective perception (CPM) messages, which largely corresponds to Classes A and B in the SAE J3216 framework.

Klimke and Mertens et al. [8] present one of the few real-world demonstrations in mixed traffic, where an edge computing unit maintains a dynamic environment model and computes cooperative maneuvers for connected and automated vehicles at an unsignalized intersection. However, in this setup, the infrastructure only provides time-window constraints at conflict zones, while trajectory generation and execution remain fully on-board the vehicles.

To the best of our knowledge, no real-world system has yet demonstrated infrastructure-based prescriptive cooperation as defined by Class D at complex urban intersections, where the infrastructure performs joint multi-vehicle trajectory planning for connected and automated vehicles.

3 Methodology

Multi-agent trajectory planning is performed using ADORe® . The same multi-agent framework is employed in both single-vehicle and multi-vehicle control cases. In the single-vehicle setting, the planned trajectories of surrounding vehicles are based on realistic predictions. In the multi-vehicle control configuration, where the centralized infrastructure (Edge Processing Unit, EPU) controls a subset of the vehicles in the scenario, the generated trajectories serve a dual role: they represent predicted trajectories for human-driven vehicles and planned trajectories for automated vehicles, respectively.

The proposed approach must address two key challenges in multi-agent trajectory planning:

- **Heterogeneity:** Mixed-traffic environments require the planner to account for the behavior of human-driven vehicles and other non-controllable agents, incorporating their predicted influence into the planning process for controllable vehicles.
- **Scalability:** The presence of a large number of interacting agents in dense traffic scenarios leads to significant computational complexity, causing many standard

multi-agent planning methods to scale poorly and become computationally prohibitive.

To mitigate these challenges, a lightweight multi-agent trajectory planning scheme based on forward integration of a control law is employed. At each timestep, control actions are computed for both controllable and non-controllable agents using the current system state and local interaction information. The system state is then propagated forward through time via numerical integration, and this procedure is iteratively repeated over the prediction horizon.

$$\begin{cases} \delta_t = k_{\psi_t} \cdot (\psi_{ref} - \psi_t) + k_{lat} \cdot e_{lat} \\ a_t = k_{v_t} \cdot (v_{IDM} - v_t) + k_{curve} \cdot |\theta_{ref}| \end{cases} \tag{1}$$

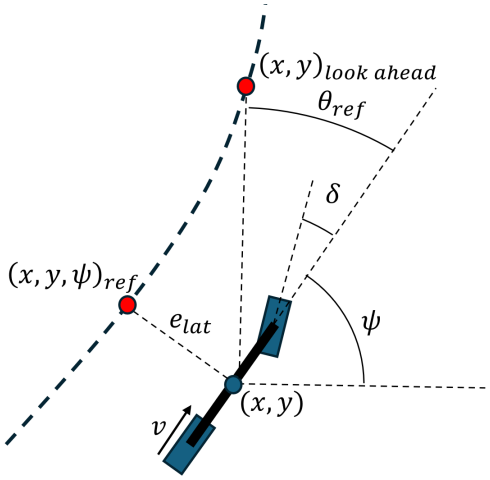

Figure 1: Model representation used for planning [9].

The controller formulation is provided in Eq. 1, while the associated state and control variables are illustrated in Fig. 1. In this framework, v_{IDM} denotes the reference velocity computed using the Intelligent Driver Model (IDM), which captures longitudinal interactions with the nearest preceding vehicles. Pedestrian motion is modeled under a constant-speed, constant-heading assumption. The resulting output consists of a coupled set of trajectories, where controllable vehicles are assigned prescriptive plans, while non-controllable agents are represented through predictive trajectory estimates.

From a computational perspective, the primary bottleneck arises from the nearest-neighbor search used to model inter-agent interactions. In the current implementation, this is performed via a pairwise comparison strategy with quadratic complexity, $\mathcal{O}(N^2)$. Although this could be reduced to $\mathcal{O}(N \log N)$ using spatial data structures such as k-d trees, the simplicity and low cost of the control law ensure that the overall system remains suitable for real-time operation even without such optimizations.

In the single-vehicle control case, the resulting ego-vehicle trajectory, accounting for interactions with surrounding vehicles through their predicted motions, is forwarded to a low-level MPC-based trajectory tracker for execution.

In the multi-agent control case, the planned trajectories for controlled vehicles are encoded into so called EPU-to-Vehicle messages (E2V), which have been specifically designed

for this purpose during the project, based on early drafts of the MCM, and transmitted via the Roadside Unit (RSU). The vehicles have then their own low-level trajectory tracker to follow the reference trajectory from the infrastructure.

Within the infrastructure-controlled section, a controllable vehicle prioritizes the infrastructure-generated trajectory over its locally computed plan, provided that the message passes local safety consistency checks. In parallel, the onboard ADAS maintains a fallback trajectory to ensure safety in the event of communication loss. Upon exiting the infrastructure-controlled section, control authority is smoothly transferred back to the vehicles' internal automation system.

4 ADORe® Overview

ADORe® is an open-source Autonomous Driving Stack (ADS) initiated by the German Aerospace Center (DLR). ADORe® separates itself from similar ADS stacks by focusing on system-level design to enable realistic coordinated decision making. The ADORe® stack is modular and includes decision making, control, planning, and visualization. There are two different *flavors* of ADORe® , SAAD and MAAD.

ADORe® SAAD is designed to be deployed on a single vehicle, enabling ADS functionalities, assuming the availability of an upper stream sensor stack providing obstacles and a situation representation. ADORe® SAAD also allows to interface with any external service, provided through V2X (like e.g. traffic light state reception), backend or edge communication (like e.g. fleet management or remote operation).

ADORe® MAAD is designed to be deployed on infrastructure level. This version of the software is able to predict the behavior of human driven vehicles (HDVs), to manage the interactions with multiple traffic participants and traffic lights, within a defined area of influence. ADORe MAAD's capabilities increase when communicating with one or more CAV systems equipped with ADORe SAAD or a similar AD stack capable of dealing with the related V2X information inside of its area of influence, as in that case the planning of multiple connected cooperative vehicles can be done centrally on the local infrastructure [2].

4.1 Single-Agent Autonomous Driving

A CAV equipped with ADORe® SAAD is able to drive at up to SAE level 3, although mimicking level 4 behavior for research. ADORe® SAAD is a ROS2-based software, requiring the presence of a sensor stack in ROS2 and a high-definition map. The information needed to drive consists of the vehicle's position and orientation and the list of detected objects including their positions, orientations, linear and angular velocities of each agent in the vicinity. ADORe® SAAD also receives vehicle state data from the host vehicle. It provides a Lichtblick based GUI for the test drivers and passengers, allowing to provide goal position to the automation. Lastly, V2X communication topics are received by ADORe® SAAD, but are not mandatory for its automated driving capability. The basic V2X communication topics include handling of CAMs, SPaTEMs, CPMs as well as different Vehicle-to-Vehicle (V2V) cooperation messages such as early drafts of Maneuver Coordination (MCM). Besides, ADORe® SAAD also handles suggested trajectories created by ADORe® MAAD and is integrated into a remote operation service.

The ADORe® SAAD outputs can be summarized into vehicle commands, i.e. steering angle and acceleration, which are sent to the low-level controller. Vehicle signals are further messages used to trigger other vehicle functions, like indicators. Lastly, SAAD outputs V2X information (e.g. CAM), as well as state information to the remote operation system as described in sec. 4.2. ADORe® SAAD also sends out information about its capabilities and executed commands when driving in an area where ADORe® MAAD is available 4.3 [2].

4.2 Remote Operations

Even most sophisticated autonomous driving systems on SAE Level 4 will be restricted to specific Operational Design Domains (ODD). An example of such restriction is the CAV driving along a one-lane road with a solid center line (meaning the opposing road cannot be entered) while encountering a road obstacle blocking any further progress along the road. While a human driver would most likely ignore the solid line if the traffic situation allows it, such a context-based decision would possibly be problematic, unsafe or not allowed for CAVs in real traffic. In such cases, a remote operator (RO) is linked to the vehicle, allowing to remotely assist the vehicle to solve the situation. A CAV using ADORe® SAAD is constantly providing current state information to the remote operation system. In case of a detected violation of the ODD, it is requesting assistance by the RO. The RO can see all CAV positions of the fleet in real time, while a CAV that encounters a problematic situation is highlighted. Through cooperation (e.g. providing sensor inputs, establishing video connections, etc.), the CAV and the RO can find a solution to the problem by providing the ADS the required support and trusted interpretation of the data so that it can execute the command when ready.

In practice, the procedure is as follows: First, the CAV detects and provides the issue to the RO. The RO can then get an overview of the CAVs situation, e.g. by establishing video transmissions from the CAV. In case of the blockage, the RO can define a path the car should drive, in the form of waypoints, to overcome the situation. These waypoints are sent to the CAV, converted into a drivable trajectory in the ADS, sent back to the RO and visualized. If the RO is satisfied with the trajectory provided, it can be either approved, which will start the execution of the generated trajectory on the CAV, or rejected. This setup allows for human decision-making to be part of the automated driving behavior in edge and corner cases, and for a single remote operator to oversee multiple CAVs simultaneously.

4.3 Multi-Agent Autonomous Driving

ADORe Multi-Agent Autonomous Driving (MAAD) extends ADORe® to support joint planning, prediction, and control of multiple CAVs operating within mixed traffic, made possible by infrastructure support and V2X communication technologies.

The framework predicts the behavior of non-controllable agents while optimizing trajectories for controllable vehicles. ADORe® adopts a non-hierarchical logic that models CAVs and HDVs symmetrically. Trajectory planning and behavior prediction are executed in parallel, yielding an interactive multi-agent system in which CAVs and other traffic participants mutually influence each other. The adopted control logic, based on

forward integration of a control law, allows to manage real world traffic scenarios, with the computational complexity scaling quadratically with the number of traffic participants.

5 Implementation and Demonstration

ADORe® has been successfully deployed within the project MAD-Urban [3] [9], which integrates collective perception from infrastructure and vehicle sensors with infrastructure-level multi-vehicle trajectory planning for the control of automated vehicles in mixed traffic. To the best of our knowledge, this constitutes the first real-world implementation of closed-loop, infrastructure-based trajectory prescription on public urban roads.

In the project, different CAVs such as the DLR ViewCar2 [10] were driving around the signalized intersection of Tostmannplatz in Braunschweig, Germany [11]. While the ViewCar2 is equipped with ADORe® SAAD, the intersection is equipped with ADORe MAAD. The vehicles were broadcasting extended V2X CAMs describing the state, intent, and willingness to follow external guidance. The intersection is equipped with several different sensors to provide full coverage of the intersecting area. The sensors were connected to a perception stack, serving ADORe® MAAD.

When a CAV reaches the area covered by the infrastructure-based sensors, the infrastructure fuses these CAMs with its own sensor data to maintain a unified dynamic environment model and predict short-term behaviors. Based on this fused view, ADORe® MAAD computes cooperative, conflict-free reference trajectories to enhance safety and traffic smoothness. These trajectories are sent back to CAVs via project designed E2V V2X messages, which CAVs validate through internal safety checks before execution. The E2V includes administrative information such as the polygon in which the service is offered, as well as individual maneuver, trajectory and motion commands. If the received commands are valid, the CAV performs a transition of control to the external system and executes them. Depending on the set level of allowed external influence, the CAV may decide whether to follow the general maneuver, trajectory or motion command, or to just treat the commands as advisories. This provides backward compatibility and ensures benefiting also on lower levels of cooperation. Once ADORe SAAD predicted that it would leave the area of influence within four seconds, it would transition back to vehicle-based control.

Excluding development activities, the majority of demonstration activities for MAD-Urban with the described behavior took place on 17th of September 2025 at Tostmannplatz. This event had the CAVs drive the area observed in Fig. 2, where the CAVs' start positions are highlighted in green, their routes in yellow, and the infrastructure-support area is in blue. During the event, multiple groups participated as civil passengers in the DLR Viewcar2. Each demonstration with passengers lasted a few minutes, depending on the traffic light, with the CAV driving automated from each marked start position till the end of the scenario, and with the safety-driver then manually performing U-turns to return the CAV to the next scenario's starting position due to the road's layout.

During the demonstration activities, the CAVs successfully drove automated around Tostmannplatz, and transitioned to infrastructure-based cooperative driving when inside the infrastructure-observed area. During each scenario, no passenger mentioned or indicated noticing a transition of control to and from MAAD-driving.

As part of the project, the Remote Operation implementation was also successfully

Figure 2: Traffic light at Tostmannplatz and demonstration area of MAD-Urban, each green marker indicates a possible start position of the CAVs, and the route/scenario they drive is marked yellow with a letter to separate them. The area marked blue is what could be observed by the infrastructure.

explored in real traffic situations. The CAV was driving along Fig. 2 route C, and when reaching the end it entered a problematic area for remote operations. The ADORe SAADs' request for assistance was triggered, the CAV in collaboration with a Remote Operator, was successfully capable of leaving the problematic area.

6 Limitations and Outlook

In this paper, we presented a real-world V2X cooperative framework built upon the open-source automated driving software ADORe® . The underlying multi-agent trajectory planning framework enables ADORe® to operate both as an ADS on a single automated vehicle and as an infrastructure-based planner capable of jointly predicting and generating trajectories for multiple vehicles using infrastructure sensing and infrastructure-to-vehicle communication.

The successful implementation and deployment at a real urban intersection on public roads demonstrates the flexibility and potential of this approach. Despite these promising results, several limitations remain and give rise to new research challenges:

- **Pedestrian and cyclist modeling.** In the current implementation, the behavior of pedestrians and cyclists is approximated using a constant-velocity, constant-heading model. However, these road users, particularly pedestrians, exhibit highly stochastic and context-dependent motion, making realistic prediction significantly more difficult. Future work should incorporate more expressive behavioral models to better capture their dynamics.

- **Scalability and computational complexity.** Although the planner can handle a substantial number of agents, its computational complexity remains $\mathcal{O}(N \log N)$, which may become problematic for very large or highly complex intersections where

both agent count and computational constraints are significant. Emerging approaches based on generative AI offer a promising alternative: by predicting future top-view scene frames directly in the pixel domain, trajectory prediction becomes independent of the number of agents. However, this introduces new challenges, such as conditioning the generative process on control inputs, preventing artifacts or hallucinations, and ensuring generalization from dataset-specific training to diverse intersection geometries. These represent valuable directions for future research.

References

[1] "ERTRAC position paper on Road Transport Research in Horizon Europe follow-up," https://www.ertrac.org/wp-content/uploads/2025/12/ERTRAC-Position_Paper-Decembre-2025-Road-Transport-Research.pdf, accessed: 2025.

[2] M. S. Maarssoe, S. Konthala, M. Mizdrak, G. Lucente, M. Nichting, T. Lobig, and A. Koerner, "ADORe: Unified Modular Framework for Vehicle and Infrastructure-Based System Level Automation," in *Proceedings of the 11th International Conference on Vehicle Technology and Intelligent Transport Systems (VEHITS)*, INSTICC. SciTePress, 2025, pp. 571–581.

[3] "MAD Urban - Managed Automated Driving," https://www.dlr.de/en/fk/research-and-transfer/projects/global-projects/mad-urban, Braunschweig, Germany, DLR.

[4] E. Yazdani Bejarbaneh, H. Du, and F. Naghdy, "Exploring shared perception and control in cooperative vehicle-intersection systems: A review," *IEEE Transactions on Intelligent Transportation Systems*, vol. 25, no. 11, pp. 15 247–15 272, 2024.

[5] Z. Yao, Y. Zhao, H. Jiang, and Y. Jiang, "A review of cooperative intersection: From design to management," *IEEE Transactions on Intelligent Vehicles*, pp. 1–23, 2025.

[6] C. Cress, Z. Bing, and A. C. Knoll, "Intelligent transportation systems using roadside infrastructure: A literature survey," *IEEE Transactions on Intelligent Transportation Systems*, vol. 25, no. 7, pp. 6309–6327, 2024.

[7] "J3216_202107: Taxonomy and Definitions for Terms Related to Cooperative Driving Automation for On-Road Motor Vehicles - SAE International," Dec. 2024, [Online; accessed 13. Dec. 2024]. [Online]. Available: https://www.sae.org/standards/content/j3216_202107

[8] M. Klimke, M. B. Mertens, B. Voelz, and M. Buchholz, "Real-world evaluation of two cooperative intersection management approaches," 2025. [Online]. Available: https://arxiv.org/abs/2403.16478

[9] M. Mizdrak, G. Lucente, S. Konthala, M. S. Maarssoe, C. Boeker, J. Burger, M. Heinrich, A. Lee, J. Doll, S. Ochs, T. Fleck, and J. Schindler, "Managed Automated Driving (MAD): Real-World Demonstration of Infrastructure Controlled Vehicle Automation," in *IEEE Intelligent Vehicles Symposium (IV)*, 2026, accepted for publication.

[10] "ViewCar2," https://www.dlr.de/en/research-and-transfer/research-infrastructure/view-car, Braunschweig, Germany, DLR.

[11] "Traffic research for greater safety and smoother traffic flow," https://www.dlr.de/de/aktuelles/nachrichten/2025/verkehrsforschung-fuer-mehr-sicherheit-und-besser-fliessenden-verkehr, Braunschweig, Germany, DLR.