Leveraging External Hazard Data to Safeguard Automated Driving System

Kamrul Hasan, Ali Shakeri, and Bernd Westphal
Institute of Systems Engineering for Future Mobility
German Aerospace Center (DLR)
Oldenburg, Germany
{kamrul.hasan, ali.shakeri, bernd.westphal}@dlr.de

***Abstract*—Automated driving systems (ADS) are constrained by the limits of onboard perception, especially when visual obstructions create hazards. Using external data, such as V2X hazard messages, can mitigate these shortcomings but may introduce additional hazards due to their unreliability. We propose an Early Hazard Processing System (EHPS) that informs an extended ADS about external hazards, enabling it to perform proactive risk mitigation. Extended ADS has a cautious state that proactively reduces speed and applies comfort-bounded deceleration when EHPS reports potential hazards. All other safety-critical actions depend on internal sensor confirmation and remain under the control of ADS fallback mechanisms. Our proposed solution is formalised with a three-state model that is shown to be as safe as baseline ADS, supported by safety argumentation. Finally, our simulation-based evaluation using the Automated Lane Keeping System shows how early awareness provides more reaction time and more comfortable braking, even when external information is unreliable. Our results show that even unreliable external hazard data improves ADS safety and comfort without compromising baseline safety guarantees.**

***Index Terms*—Automated Driving Systems (ADS); Vehicle-to-Everything (V2X); Collective Perception; Software Defined Vehicle (SDV); External hazard data.**

I. Introduction

The continued advancement of Automated Driving System (ADS) holds the promise of significantly improving traffic safety and efficiency. A key challenge to realising this promise, however, lies in the inherent limitations of onboard perception systems. Even state-of-the-art sensors remain constrained by line of sight, finite range, and field of view restrictions, leading to functional insufficiencies as described in ISO 21448 (Safety of the Intended Functionality, SOTIF) [1]. These insufficiencies can increase risk in common driving scenarios such as blind curves, hill crests, and occluded intersections, where an ADS may only detect a hazard when collision avoidance requires abrupt emergency manoeuvres [2].

Vehicle-to-Everything (V2X) communication, and in particular the exchange of Cooperative Awareness Messages (CAM) [3] and Collective Perception Messages (CPM) [4] as standardised by ETSI , offers a promising way to increase awareness of ADS. In addition, data ecosystems such as GAIA-X and the software-defined vehicle (SDV) paradigm enable ADS to receive and utilise such external data. By receiving external hazard information from other vehicles or infrastructure that are part of a broader data ecosystem, an ADS can be made aware of hazards much earlier than onboard sensors can detect them. Nevertheless, this external information is potentially unreliable, as it may suffer from uncertainty, latency, incompleteness, or even temporary unavailability [5], [6]. As a result, external hazard data have not been widely integrated into safety-critical ADS functions, despite their potential to mitigate functional insufficiencies.

In this study, we address how external hazard data, despite their potential unreliability, can improve ADS behaviour. Specifically, we argue that even when not fully reliable, such information remains safety-relevant if it is used to guide preparatory behaviour rather than to trigger safety-critical manoeuvres directly. We propose the Early Hazard Process System (EHPS), which aims to process external hazard information and send it to ADS. Additionally, we propose an extension to ADS that utilises EHPS information to enter a Cautious State, a preparatory state in which the system adopts more conservative driving behaviour until hazards are confirmed by onboard sensors. All decisive safety actions, including Transition Demand (TD), Minimum Risk Maneuver (MRM), and Emergency Maneuver (EM), remain exclusively the responsibility of baseline ADS.

Our contribution is the introduction of the EHPS and the extension to ADS, which we formalise through a three-state model. Building on this foundation, we define cautious state properties that ensure safe operation, particularly in the event of false alarms. To demonstrate the practical benefits of this approach, we present a simulation-based evaluation on Automated Lane Keeping System (ALKS) as the system under test. Our evaluation shows that EHPS significantly improves Time to Collision (TTC) for extended ALKS, resulting in a more comfortable driving compared to a baseline ALKS.

The content of the paper is structured as follows. In Section II, we review related work on cooperative perception and external data reliability. Section III starts with the assumptions, scope, and enabling technologies, and then presents our core contribution, including its state machine model and system properties. This is followed by a safety argument comparing baseline and EHPS enabled systems in Section IV. Section V describes the scenario-based simulation and presents quantitative results. Finally, Section VI concludes the paper and outlines future work.

II. Related Work

Research on cooperative and collective perception (CP) in V2X communication shows that sharing sensed objects among vehicles and infrastructure can extend the perception horizon of ADS. Huang *et al.* provide a comprehensive survey of CP pipelines, including agent selection, alignment, and fusion strategies, and compare early, intermediate, and late fusion methods [7]. The survey emphasises the role of standardised message formats such as CAM [3] and CPM [4], while highlighting delay, packet loss, and uncertainty as key limitations.

The reliability and quality of shared information are active research topics. Allig *et al.* propose a Dempster–Shafer-based method for estimating the trustworthiness of entities within CP [8]. Their results show that trust metrics help filter inconsistent or uncertain objects before fusion, which underlines that external data should not be treated as ground truth but can serve as advisory input. Thunberg *et al.* analyse unreliable V2X communication and introduce the notion of safety time for emergency braking [5], providing formal bounds on the Age of Information (AoI) that must be respected to guarantee that cooperative safety manoeuvres remain feasible. This motivates bounded use of external hazard data.

The S2V2X framework positions V2X communication as a sensor subsystem within the functional safety and SOTIF processes of automated driving [9]. It identifies safety requirements, validation needs, and standardisation aspects necessary to ensure that V2X data can be considered in type approval and certification. While it offers an important architectural foundation, it does not define concrete runtime logic for handling potentially unreliable external data.

A Study on the broader data ecosystem of SDV examines how decision-making can be distributed across a vehicle–edge–cloud continuum. Peltonen *et al.* propose a data-centric, context-aware framework that employs open-source components such as Eclipse Kuksa[1] and in-vehicle data brokers, emphasizing metadata-driven context fusion, governance of heterogeneous data flows, and integration of sources ranging from environmental sensing to driver feedback [10]. The architecture orchestrates diverse internal and external data streams, with real-time safety decisions remaining onboard and less time-critical tasks delegated to the edge or cloud.

Existing research has largely focused on extending the perception horizon and improving object-level detection and tracking by sharing sensor data among vehicles and infrastructure. Trust- and quality-aware approaches introduce metrics or evidential reasoning (e.g., Dempster–Shafer based methods) to filter unreliable objects before fusion. Likewise, V2X-centric frameworks such as S2V2X and SDV data-space architectures position V2X as a sensor subsystem or data source, but do not define concrete runtime logic for how potentially unreliable external hazard data should influence driving behaviour while remaining compliant with existing ADS safety regulations. In contrast, we introduce a dedicated preparatory state in the ADS behaviour model that is driven by external hazard data but keeps all decisive safety actions tied to internal perception and existing fallback mechanisms. Our approach provides a formal safety argument, showing that the extended ADS is at least as safe as the baseline ADS, even when external data are unreliable.

[1]Eclipse Kuksa: https://www.eclipse.org/kuksa/

III. An Extended ADS Model for External Hazard Integration

A. System Prerequisites and Scope

Our proposed solution relies on several assumptions. First, ADS are assumed to be equipped with the capability to receive external hazard information through V2X communication. Such data are considered authentic, in the sense that they originate from trusted sources, but may still be uncertain, delayed, or incomplete. In this study, we focus on uncertainty (see [11] for a formal definition of uncertainty used in this study) as the main source of unreliability in the hazard data. This uncertainty can show itself in two different ways. Uncertainty may exist in measurements that report the existence of the hazard, or in the measured distance to the hazard [4].

The fallback behaviours, such as TD, MRM, and EM, are assumed to be implemented within the ADS. Our solution does not aim to modify these safety-critical functions. Instead, it provides an additional preparatory capability for the ADS. Our proposed solution builds on several technologies and regulatory frameworks. For concreteness, ALKS serves as our specific system under test, regulated by UNECE R157 [12] which specifies minimum forward detection ranges, transition conditions, and fallback requirements.

B. Context

Several data space initiatives, such as Gaia-X [13], aim to enable federated and secure data exchange among users. In the automotive domain, the Gaia-X 4 AGEDA project seeks to develop an in-vehicle software architecture capable of hosting applications that interact with cloud-based services. These applications are designed not only to enhance safety but also to provide added value for stakeholders in the automotive ecosystem. One class of such applications focuses on processing traffic data within smart city infrastructures.

In this study, we consider a Gaia-X compliant V2X infrastructure that publishes traffic data to all interested road users that are also Gaia-X compliant. Such V2X traffic data can be used in various ways. Our focus is on a subsystem, the EHPS, which subscribes to V2X data, processes them, generates hazard information, and transmits it to the ADS. This hazard information enables the ADS to anticipate and react to hazards earlier than on-board sensors can detect.

As shown in Figure 1, the V2X infrastructure broadcasts relevant traffic information for a specific region. This traffic information is represented by the TrafficData interface, which is intentionally left underspecified for two main reasons. First, the specific data content depends on the region (e.g., highway or urban), and second, a detailed specification of this interface is not essential for the objectives of this study. Communication can be realised through different frameworks.

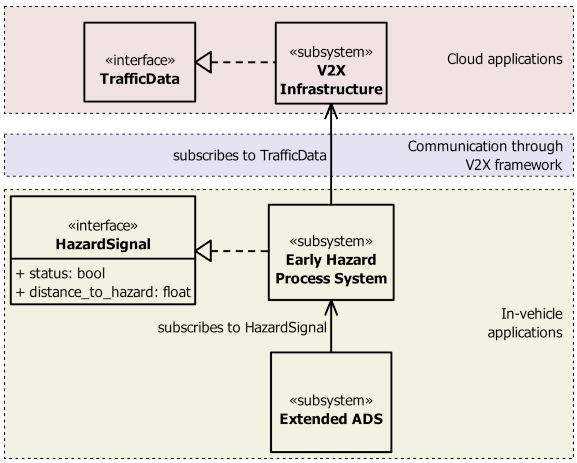

Fig. 1. Schematic view of the subsystems considered in this study and their interfaces.

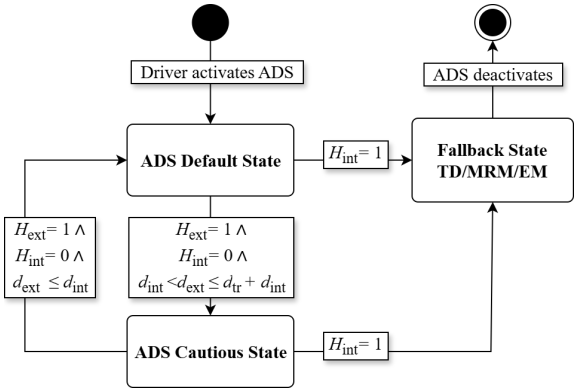

Fig. 2. The state machine model of extended ADS with three states and transitions between states.

Here, the AGEDA is used as an example of such frameworks. Inside the vehicle, the EHPS application subscribes to relevant information from the V2X infrastructure and processes it to derive hazard information in the vicinity of the ADS. This hazard information is represented by the HazardSignal interface, which includes a boolean variable status indicating the presence or absence of a hazard, and a floating-point variable denoting the distance from the ADS to the hazard. The HazardSignal is assumed to be inherently unreliable, meaning there is uncertainty regarding both the existence of the hazard and the distance estimation. Section IV discusses how this unreliability is addressed in safety argumentation.

The EHPS transmits hazard information to the ADS constantly. The ADS uses the HazardSignal to adapt its behaviour in response to potential hazards ahead. In this regard, the ADS needs an extension, in order to achieve that goal. We propose a minimal extension represented as a three-state model formalised by a state machine (see Figure 2). Before detailing the properties of this model, it should be emphasised that the simplicity of the model is intentional. This simple model enables analysis of the safety implications of the proposed extension while preserving the overall system functionality.

C. The Three-State Behavioral Model

After activation of ADS, it is assumed to operate in a Default State, whose properties are aligned with those of Dynamic Driving Task (DDT) as specified in UNECE R157 [12]. In addition, the ADS has a Fallback State, which, depending on the level of automation, may correspond to Transition Demand, Minimal Risk Manoeuvre, or an Emergency Manoeuvre. We extend the ADS by introducing a new Cautious State. The Cautious State is a preparatory state which the ADS transitions into based on external hazard information. The Cautious State enforces conservative longitudinal and lateral limits. This state never commands decisive safety actions and those actions remain tied to onboard sensor confirmation and ADS safety functionalities.

In order to describe the three-state model we define the parameters of the model. Let v be the ego speed, and d_{int} be the minimum distance that internal sensors of ADS require to detect a hazard ahead and d_{tr} be the distance that ADS travels during transition from Default state to Cautious state. This transition involves decelerating the vehicle from maximum operational speed in default state, $v_{\max}$, to a reduced cautious speed, v_c in cautious state. The distance covered during this transition, is given by $d_{\text{tr}} = (v_{\max}^2 - v_c^2)/2a_{\text{tr}}$, where a_{tr} is the moderate deceleration rate used during the transition. The transition to Cautious state takes place if at the time of receiving external hazard data, the distance to hazard reported by EHPS, d_{ext}, follows the following relation:

$$d_{\text{int}} < d_{\text{ext}} \leq d_{\text{tr}} + d_{\text{int}}. \tag{1}$$

Let $H_{\text{ext}} \in \{0, 1\}$ be a boolean variable indicating whether a hazard from an external source exists (1) or not (0). Similarly, let $H_{\text{int}} \in \{0, 1\}$ denote whether the hazard is detected by internal sensors (1) or not (0). The transition guards are specified as follows:
Default $\rightarrow$ *Cautious* (early awareness):

$$H_{\text{ext}} = 1 \ \wedge \ H_{\text{int}} = 0 \ \wedge \ d_{\text{int}} < d_{\text{ext}} \leq d_{\text{tr}} + d_{\text{int}}. \tag{2}$$

Cautious $\rightarrow$ *Default* (no hazard detected by on-board sensors):

$$H_{\text{ext}} = 1 \ \wedge \ H_{\text{int}} = 0 \ \wedge \ d_{\text{ext}} \leq d_{\text{int}}, \tag{3}$$

Any $\rightarrow$ *Fallback* (on-board sensors detect the hazard):

$$H_{\text{int}} = 1. \tag{4}$$

This structure ensures that external information influences only preparatory behaviour, while safety-critical decisions remain with the certified onboard perception and control logic.

D. Operational Behaviour of the Cautious State

The properties of Default and Fallback states are mentioned in UNECE Regulation No. 157 (ALKS) [12]. In the following we present the properties of cautious state. In this state the system operates similar to Default state while giving up some of the functionalities and avoids unnecessary manoeuvres.

- The maximum operational speed in cautious state is less than default state.
- The lateral manoeuvres are suppressed unless required by the collision avoidance or fallback mechanisms (e.g. emergency manoeuvre).
- The deceleration during transition to cautious state, a_{tr}, remains within comfort deceleration limit [14].

These preparatory measures are intentionally bounded and non-invasive. They improve comfort while maintaining the baseline safety in cases of external hazard indication. We describe the safety argumentation in the next section.

IV. Safety Argumentation

The analysis in this section compares the safety performance of the baseline ADS with an EHPS-enabled ADS. The baseline ADS relies solely on its internal sensors, which is able to detect a hazard ahead from distance d_{int}. The EHPS enhances the baseline system by sending external hazard data. From the external data source, a distance estimate to the hazard ahead will be derived that is denoted with d_{ext}, which includes an uncertainty margin d_ε. After receiving, the external hazard data, the ADS transitions into a Cautious state according to the state machine logic that is presented in section III.

We start our analysis by assuming that external hazard data is reliable and there is no uncertainty in the existence of hazard or its location. To quantify and compare the safety of the two systems, we use two indicators, the TTC and the required deceleration to stop before the hazard. We analyse a scenario where a hazard is first detected by the internal sensors at ADS detection range. According to UNECE regulation R157 for a maximum speed of $v_{\max} = 130\,\mathrm{km/h}$, forward detection range must be at least $150\,\mathrm{m}$. We therefore assume that the hazard detection distance of baseline ADS, d_{int} is according to the UNECE regulation R157.

The TTC represents the time remaining before a collision would occur if the vehicle were to maintain its current speed. For the baseline ADS, which is traveling at $v_{\max}$ upon detection, the TTC can be calculated as $\mathrm{TTC}_B = {}^{d_{\mathrm{int}}}/_{v_{\max}}$. Also, for the EHPS-enabled ADS, which has already reduced speed to v_c by the time of internal detection, the TTC is $\mathrm{TTC}_{\mathrm{EHPS}} = {}^{d_{\mathrm{int}}}/_{v_c}$.

Since $v_c < v_{\max}$, it is clear that $\mathrm{TTC}_{\mathrm{EHPS}} > \mathrm{TTC}_B$. This indicates that the EHPS enables the ADS to have more time to react and execute a fallback manoeuvre.

While TTC indicates available time before collision, the required deceleration provides a more direct indication of the physical feasibility of stopping. This measurement defines the minimum constant deceleration needed to stop safely before the hazard's location. A lower value signifies a larger safety margin. The deceleration that a vehicle driving with a speed v, requires to fully stop over a distance d can be calculated as $a_{\mathrm{req}} = {}^{v^2}/_{2d}$. For the baseline ADS traveling at $v_{\max}$, the required deceleration is $a_{\mathrm{B}} = {}^{v_{\max}^2}/_{2d_{\mathrm{int}}}$. For the EHPS-enabled ADS traveling at v_c, the required deceleration is $a_{\mathrm{EHPS}} = {}^{v_c^2}/_{2d_{\mathrm{int}}}$. Since $v_c < v_{\max}$, one can show that $a_{\mathrm{EHPS}} < a_{\mathrm{B}}$. As a result, the EHPS-enabled ADS requires a significantly lower, more manageable braking to avoid a collision.

We now consider other cases where uncertainty exists in external hazard data. This uncertainty could show itself in the reported location of the hazard or its existence. So the following cases could appear:

Case 1: Hazard exists and is reported by EHPS without uncertainty.
The system transitions to the Cautious state. As mentioned earlier, this enhances safety by increasing the TTC and reducing the required deceleration compared to the baseline.

Case 2: Hazard does not exist but is reported by EHPS.
The ADS slows down, confirms the absence of a hazard with its internal sensors, and then transition back to Default state. This scenario prioritizes safety over performance, introducing no collision risk.

Case 3: Hazard exists but is not reported by EHPS.
In this case, the ADS operates identically to the baseline system, detecting the hazard at d_{int} while at speed $v_{\max}$. Its safety performance is therefore exactly that of the baseline.

Case 4: Hazard is closer to ADS than the reported location by EHPS.
In this case, since the hazard appears earlier than expected, the ADS does not have enough time to completely transition to cautious state. However, the speed of the vehicle is less than $v_{\max}$ at the time of detecting the hazard with internal sensors. Therefore, this case is still an improvement compared to baseline.

Case 5: Hazard is farther to ADS than the reported location by EHPS.
In this case, since the hazard is farther than expected, the vehicle starts transitioning back to default state. This is not different from baseline ADS.

This case split analysis shows that in all cases, the EHPS-enabled ADS performs at least as safely as the baseline system. In most cases, the EHPS-enabled ADS even improves safety. Whether the external hazard data is fully reliable (Case 1), the hazard is closer than reported (Case 4), or the hazard is farther than reported (Case 5), the ADS is in a more prepared state. If an existing hazard is not reported by EHPS (Case 3), the system's performance remains identical to the baseline, as the ADS relies on its internal sensors and therefore does not create a less-safe state.

On the other hand, Case 2 is the only one where the ADS's behaviour deviates from the baseline in the absence of a true hazard. While this is an "unnecessary" manoeuvre that impacts performance and traffic flow, it does not create a rear-end collision risk, as the deceleration during the transition to a cautious state is within the bounds of comfortable deceleration.

V. Experimental Evaluation: ALKS with EHPS

Building on the conceptual model as discussed in Section III and the analytical safety argumentation presented in

Fig. 3. Highway scenario in CARLA: ALKS approaches a hazard and performs a fallback manoeuvre.

Section IV, we now assess whether the preparatory Cautious State yields measurable benefits under realistic driving conditions. We therefore conduct a scenario-based evaluation in a simulation environment, comparing a regulatory compliant baseline ALKS against an EHPS enabled ALKS in different highway scenarios. The evaluation uses the CARLA as simulation environment [15] and reproducible scenarios created with ScenarioRunner.

A. Evaluation Scenario

The scenario recreates a highway driving situation where a stationary hazard blocks the ALKS lane of travel as depicted in Figure 3. The ALKS performs a fallback manoeuvre and stops before the hazard. The hazard is positioned at a predefined distance along the ego's nominal route, with the environment initialized to ensure repeatable hazard exposure across multiple runs. Each simulation tick captures vehicle state (position, velocity, acceleration), inter-actor distances, and control commands. From this information, we derive the performance metrics, TTC and average required deceleration rate for offline analysis.

B. Driving Agent Configurations

The simulation is evaluated for two distinct driving behaviours, each implemented as a unique agent configuration within the CARLA framework. CARLA agents [16] are automated driving control modules implemented within the CARLA simulator. ScenarioRunner loads the agent specified at runtime and controls the vehicle based on the scenario. We use agents to compare different driving styles under exactly the same conditions

- **Default agent configuration:** The agent employs a conventional, regulation-compliant driving policy representing baseline ALKS behaviour. It operates at a nominal speed of $90\,\mathrm{km/h}$, and activates braking when a hazard enters the on-board perception field. For the evaluation, we selected this speed to be comparable to existing SAE Level 3 ADS [17], such as Mercedes-Benz DRIVE PILOT [18]. Upon detection, braking follows the regulatory deceleration requirements without external data intervention.
- **Cautious agent configuration:** The agent extends the Baseline ALKS through integration of the EHPS. Upon receiving the external hazard data, the agent triggers the Cautious State transition, adopting more conservative driving behaviour representing Case 1 mentioned in IV. In this state, operational speed reduces to $60\,\mathrm{km/h}$ following minimum speed limit in the highway, and employs gentler deceleration trajectories. Additionally, it can also handle scenarios similar to Case 4 where the EHPS receives unreliable hazard data. In these cases, the agent enters the Cautious State in response to external hazard data.

TABLE I
SUMMARY OF TTC AND AVERAGE REQUIRED DECELERATION

Behaviour [$run(n) = 10$]	**TTC** [s]±**std**	**Required Deceleration** [$\mathrm{m/s^2}$]±**std**
Default	3.102 ±0.021	3.306 ±0.032
Cautious	4.497 ±0.018	1.267 ±0.001
Cautious (Unreliable data)	3.730 ±0.001	1.957 ±0.009

C. Simulation Results

Across the different behaviour configurations with $n = 10$ simulation runs, the experiments show consistent improvement. Compared to Default, the Cautious configuration increases TTC by 45%, from $3.102\,\mathrm{s}$ to $4.497\,\mathrm{s}$ and lowers the average required deceleration by 61.7%, from $3.306\,\mathrm{m/s^2}$ to $1.267\,\mathrm{m/s^2}$ (see Table I). When the external data are unreliable (Cautious, Unreliable; $d_\varepsilon = 50\,\mathrm{m}$, hazard closer than reported), TTC still improves by 20.3% and the average required deceleration decreases by 40.8% relative to Default. The agent behaviours show minimal absolute standard deviations, as the specificity of the scenarios constrains variability in performance metrics.

Figure 4 shows required deceleration versus distance to the obstacle. The comfort boundary is at $3.0\,\mathrm{m/s^2}$ and the hard-brake boundary at $6.0\,\mathrm{m/s^2}$. The Default configuration stays above the comfort level for most of the approach, therefore a regulation compliant but its late response drives harder braking. Both Cautious configurations shift the curve downward. Required deceleration remains below the comfort boundary throughout. The reliable data case stays at or below about $2.0\,\mathrm{m/s^2}$ whereas, the unreliable data case stays below about $2.9\,\mathrm{m/s^2}$.

The improvements come from adopting the Cautious State prior to internal detection. Although the detection distance is essentially the same across runs, at approximately $74\,\mathrm{m}$ to $75\,\mathrm{m}$, the speed at detection is lower when EHPS is enabled in ALKS: $86.5\,\mathrm{km/h}$ in the Default configuration, $59.6\,\mathrm{km/h}$ in the Cautious configuration, and $71.5\,\mathrm{km/h}$ in the Cautious–Unreliable hazard data configuration. Retaining the same distance while reducing speed increases TTC and decreases the required deceleration a_{req}, as expected from Section IV. The main trade-off is a temporary reduction in ALKS cruising speed: $60\,\mathrm{km/h}$ to $72\,\mathrm{km/h}$ in the Cautious State and $90\,\mathrm{km/h}$ in Default. This reduction is limited by the Cautious State rules. Once internal sensors confirm there is no hazard in the detection range the system returns to Default. Our scenario isolates a single-lane, stationary hazard to study

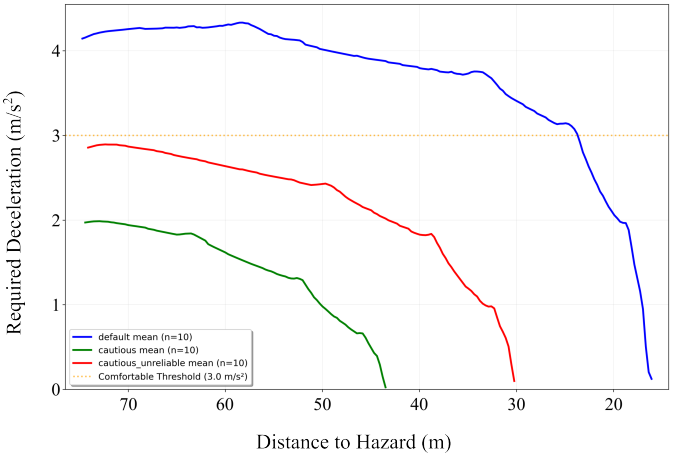

Fig. 4. Comparison of required deceleration profiles for Default and Cautious ALKS configurations as the vehicle approaches a hazard. The Cautious modes remain within comfortable deceleration limits.

longitudinal effects, while multi-lane traffic and interactions with dynamic actors are left for future work.

VI. Conclusion

This paper presents a solution to mitigate the limitations of onboard perception in hazardous scenarios. We introduced the EHPS together with a formal three-state ADS model, which uses potentially unreliable external data to trigger a preparatory Cautious State. This solution preserves the safety guarantees of the baseline system while enabling earlier, noninvasive risk mitigation. Our safety argumentation shows that, if EHPS is enabled, the extended ADS reduces speed before internal sensors detect a hazard, TTC increases, and less deceleration is required to stop safely before the hazard. A simulation-based evaluation on ALKS as the system under test, shows that our proposed solution increases TTC compared to baseline by 45% and reduces the average required deceleration by 61.7%. Even when the external hazard data are unreliable, EHPS still improves TTC by 20.3% and lowers the required deceleration by 40.8%.

Future research can extend the identified sources of unreliability to include additional factors such as latency, dropout, network or system-level disturbances, and incorporate more realistic V2X unreliability model. Moreover, the scope of the EHPS can be extended beyond advisory level to also have control over safety-critical functions. Achieving this will require the development of a more comprehensive and detailed safety case template.

VII. Acknowledgment

This work was supported by the GAIA-X 4 AGEDA project, funded by the German Federal Ministry of Economic Affairs and Climate Action (BMWK) under the funding code 19S22004N.

References

[1] *ISO 21448:2022 – Road vehicles – Safety of the intended functionality*, International Organization for Standardization Technical Specification, 2022.

[2] A. Cassel, C. Bergenhem, O. M. Christensen, H.-M. Heyn, S. Leadersson-Olsson, M. Majdandzic, P. Sun, A. Thorsén, and J. Trygvesson, "On perception safety requirements and multi sensor systems for automated driving systems," *SAE International Journal of Advances and Current Practices in Mobility*, vol. 2, no. 6, pp. 3035–3043, 2020.

[3] European Telecommunications Standards Institute, "Etsi tr 102 638 v1.1.1 – intelligent transport systems (its); vehicular communications; basic set of applications; definitions," 2009, technical Report.

[4] ——, "Etsi tr 103 562 v2.1.1 – intelligent transport systems (its); vehicular communications; basic set of applications; analysis of the collective perception service (cps)," 2019, technical Report.

[5] J. Thunberg, M. Segata, and A. Vinel, "Unreliable v2x communication in cooperative driving: Safety times for emergency braking," *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 11, pp. 6922–6932, 2021.

[6] U. Bauder, Q.-T. Le, J. Josefiok, and S. Schönfeld, "Data accuracy in vehicle-to-x cooperative awareness messages: An experimental study for the first commercial deployment of c-its in europe," in *2020 IEEE 91st Vehicular Technology Conference (VTC2020-Spring)*. IEEE, 2020, pp. 1–7.

[7] T. Huang, J. Song, F. Wang, W. Ding, M. Tomizuka, and W. Zhan, "V2x cooperative perception for autonomous driving: Recent advances and challenges," *IEEE Transactions on Intelligent Vehicles*, vol. 8, no. 7, pp. 3925–3945, 2023.

[8] C. Allig, C. Gündogan, and U. Rückert, "Trustworthiness estimation of entities within collective perception," in *2023 IEEE Vehicular Networking Conference (VNC)*. IEEE, 2023, pp. 221–224.

[9] M. Nolte, G. Bagschik, J. Tölke, and M. Schaper, "S2v2x – a new v2x framework to support and enable the design of automated vehicles with a sufficient level of safety," in *2023 IEEE International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2023, pp. 1821–1827.

[10] M. Peltonen, N. Pohlmann, S. Lenz, and A. Wolf, "Towards data-centric and context-aware decision making in software-defined vehicles," in *2025 IEEE International Conference on Connected Vehicles (ICCVE)*. IEEE, 2025, pp. 112–119.

[11] I. BIPM, I. IFCC, I. ISO, and O. IUPAP, "Evaluation of measurement data—guide to the expression of uncertainty in measurement, jcgm 100: 2008 gum 1995 with minor corrections," *Joint Committee for Guides in Metrology*, vol. 98, 2008.

[12] United Nations Economic Commission for Europe, "Unece regulation no. 157 – uniform provisions concerning the approval of vehicles with regard to automated lane keeping systems (alks)," 2021, regulation.

[13] Gaia-X European Association for Data and Cloud, "Gaia-X Architecture Document - 24.04 Release," April 2024, accessed: 30-Jan-2025. [Online]. Available: https://docs.gaia-x.eu/technical-committee/architecture-document/24.04/

[14] S. Carlowitz, P. Rossner, R. Madigan, H.-J. Bieg, C. Marberger, P. Alt, H. Otto, M. Schulz, A. Schultz, E. Kenar, A. C. Bullinger, and N. Merat, "User evaluation of comfortable deceleration profiles for highly automated driving: Findings from a test track study," *Transportation Research Part F: Traffic Psychology and Behaviour*, vol. 105, pp. 206–221, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1369847824001244

[15] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "CARLA: An open urban driving simulator," in *Proceedings of the 1st Annual Conference on Robot Learning*, 2017, pp. 1–16.

[16] "Agent evaluation – scenariorunner documentation," https://scenario-runner.readthedocs.io/en/latest/agent_evaluation/, ScenarioRunner Project, 2025, accessed: 2025-11-05.

[17] O.-R. A. D. O. Committee, *Taxonomy and definitions for terms related to driving automation systems for on-road motor vehicles*. SAE international, 2021.

[18] M.-B. G. AG, "Mercedes-benz increases top speed of its level 3 automated driving system to 95 km/h," [Online], 2024, accessed: Jan. 29, 2026. [Online]. Available: https://group.mercedes-benz.com/innovations/product-innovation/autonomous-driving/drive-pilot-95-kmh.html