

Enhanced System Awareness as Basis for Resilience of Autonomous Vehicles

Florian Woerter, Andreas Kreutz, Aniket Salvi, Marc Dreiser, Gereon Weiss

Fraunhofer-Institute for Cognitive Systems IKS

Munich, Germany

{firstname.lastname}@iks.fraunhofer.de

Abstract—The transition to autonomous driving and increasing automation of cars requires these systems to take correct decisions in very complex situations. For this, the understanding of a vehicle system’s own capabilities and the environmental context is crucial. We introduce our approach of enhancing the system awareness of vehicles to handle changes gracefully, while optimizing the overall performance. Based on a system health management the available capabilities of the distributed vehicle system can be determined. By taking into account the environment in the form of so-called operational domains at run-time, self- and context-awareness can be established providing a situation picture to which the system can adapt. We developed a service contract based solution to trigger degradations or find optimal configurations, while not endangering safety goals. Our approach is evaluated in an intersection scenario, where we can highlight the advantages of enhanced system awareness to optimize an autonomous vehicles performance.

Index Terms—autonomous vehicle, resilience, awareness, ODD, self-adaptation

I. INTRODUCTION

On the road to intelligent mobility, cars are evolving into autonomous systems that need to be able to drive reliably and efficiently through open world scenarios [1]. Managing uncertainties in understanding different situations and adapting to changes are a major corner stone to achieve this [2]. The fundamental E/E-Architectures of upcoming vehicles are already allowing for a certain degree of adaptation, e.g., by adopting service-orientation and addressing multi-level concerns of mobility [3]. However, a main open issue is the capability of an autonomous car to be sufficiently and consistently aware of its own state and environment, so the car preserves its dependability while facing changes – and thus, become *resilient* [4]. Such a system awareness requires specific concepts which encompass the heterogeneity of a mobility system and consider the inherent uncertainties, e.g., due to sensory or algorithmic limitations.

In this paper, we introduce our approach for enhanced system awareness, which integrates with present and upcoming standards for the automotive domain. Self-awareness represents the knowledge about the capabilities of the mobility system itself, which allow to perform certain actions at the moment, also including predictions of near-future capabilities. The environmental and operational conditions at hand

are captured as context-awareness. For this, we propose our concept for enhanced system awareness integrating operational conditions, contract-based service-orchestration, and overall health monitoring. We apply our approach to an autonomous car use case at an intersection which highlights the advantages of obtaining enhanced system awareness.

The remainder of the paper is as follows. With the next Section II, we describe related work. In Section III, we introduce our approach for enhanced system awareness, with details about reflecting on the operational design domain, usage of service-contracts and integration with system health monitoring. Our approach is evaluated in Section IV by the case study of an autonomous car intersection use case and concluded with Section V.

II. RELATED WORK

The backbone of modern vehicle system is the underlying architecture. With AUTOSAR [5] a development partnership of car manufacturers and suppliers develops an internationally standardized software architecture for automotive. With latest updates, it considers self-awareness by Health Monitoring [5] with different supervision types for detecting unwanted behavior in control flow, timing or aliveness of code execution, as also required by ISO 26262. Additionally, external monitoring results can be used to estimate the health of platforms. Our approach builds on top of this, combining health information of single platforms on the system level using a System Health Monitor.

A key enabler for developing higher automated driving levels is the capability of a vehicle to correctly apprehend its environment as context-awareness. To this end, perceptual uncertainties need to be properly considered to obtain an accurate representation of the surroundings. This task is further exacerbated by high environmental variability. For this, the concept of an Operational Design Domain (ODD) [6] defines functional boundaries for an automated system. Through this, developers are enabled to make certain assumptions about the operational conditions, thus, alleviating this challenge to certain extent. Besides the general application of ODDs for the system design, the work by [7], [8] deals with specifying and monitoring ODDs at run-time. The research by [9] specifically focuses on handling faults/failures and related safety aspects. In our approach, we use knowledge about the currently active Operational Domain (OD) to infer related context.

This work was partially supported by the Bavarian Ministry of Economic Affairs, Regional Development and Energy as Fraunhofer High Performance Center Secure Intelligent Systems.

In turn, such context-awareness can be used for making dynamic adaptations of the vehicle system, resulting in an optimal performance without violating safety. For ensuring the adaption of valid configurations, contract-based approaches of software services pose promising solutions. Specifically addressing safety, with Digital Dependability Identities (DDIs) [10] an approach towards an analyzable and potentially executable model of information about the dependability of a component or system has been proposed. Such representations can be used to derive dependable system configurations on top of acquired system awareness. Our proposed approach utilizes similar contract-based representations of the available vehicle services, which allow to infer service orchestration that preserves dependability by solving constraints on the present system and environment conditions exploiting enhanced system awareness.

III. ENHANCED SYSTEM AWARENESS

As a solution towards enabling situation-aware behavior of an autonomous car's system, we introduce our approach of enhanced system awareness (cf. Figure 1) and main involved concepts in the following.

A. Overview

The approach aims at allowing to maximize a vehicle system's flexibility during run-time by making the system aware of both its context and the state of the *autonomous driving system (ADS)* itself. Such a *self- and context-awareness* allows the system to dynamically adapt to changing situations. In this case, the adaptation targets to alter the provided functionalities of the vehicle, to which we refer to as *capabilities*. A capability is the ability of the system to perform a certain task. Capabilities are implemented by software services that provide them with certain service guarantees. For example, the ability of an ADS to detect objects in its proximity can be subsumed in a capability *ObjectDetection*. Services that implement this capability might differ in the type of sensor they use, such as lidar, radar, or camera. In this example, the choice of sensor could affect the detection range that the services can guarantee.

Formalized descriptions of services, including their service guarantees and requirements, are provided by *capability contracts*. Requirements of contracts are formulated with respect to information from context- and self-awareness. Continuing the previous example, the capability contract of a service *LidarObjectDetection* could state that the service

- provides the capability *ObjectDetection* with a 50 meter detection range,
- requires a Lidar-sensor with a field of view of at least 270°, and
- requires that it does not rain to ensure that the detections are reliable.

The necessary context and local information is obtained by the system using external and local monitors, which we base on recent developments in the automotive community. External monitors for context-awareness are based on ODDs, which

describe the operating conditions in which a vehicle or vehicle service is designed to correctly operate. For self-awareness, the local monitors are implemented using internal health monitoring, developed as part of the AUTOSAR initiative, e.g., Platform Health Management and Sensor Health Information [5]. In the next sections, we detail how to obtain enhanced system awareness and present how this can be used to improve the resilience of an autonomous vehicle system.

B. Operational Domain Monitoring for Context-Awareness

As designed context, an ODD defines under which conditions an automated function is designed to operate. The ODD specification process involves compiling a list of supported/prohibited *operating conditions*, e.g., weather-conditions, geo-fenced regions, traffic participants, etc. [11]. An ODD specification is created at design-time for requirements gathering and test-scenarios generation. An ADS is validated for safety against a given ODD specification, restricting the possible contexts in the general open world of an autonomous car. Therefore, an ODD specification might consist information relevant only for the design-phase. The runtime monitorable fragments of the ODD are reflected in its Runtime Operational Domain, referred as Operational Domain (OD) throughout this paper. Within our approach, we see that it is crucial to monitor run-time operational conditions to check whether the ADS is currently (and in near future will still be) operating inside its supported OD.

An OD at run-time is a composition of monitorable *operating conditions* and contexts valid over these *operating conditions*, referred to as *OD Entities*. Since, ODD specification consists of both supported and prohibited *operating conditions*, it is therefore necessary to perform monitoring considering both polarities. In our architecture, an *Operational Domain Monitor (ODM)* component is responsible for individually detecting OD entities and their associated contexts as external monitor. The ODM provides a probabilistic estimate for being inside a given OD, whilst accounting for the perceptual uncertainties. In our approach, the ODM provides context information to the System Health Monitor, so it can react on OD changes. For instance, depending on the level of autonomy, an OD-exit trigger might be followed by a fallback *Dynamic Driving Task*, bringing vehicle to a predefined known *Minimal Risk Condition*. Moreover, the ODM can provide knowledge of the context being responsible for triggering the entry or exit of an OD. In our approach, such information is aggregated as system awareness leading to an overall health indication, so that the vehicle services can be configured to safe states according to the current OD.

C. System Health Monitoring

The system awareness is realized by the System Health Monitoring [12], which comprises safety run-time analysis according to ISO 26262 for malfunctioning of E/E systems with performance limitations of ISO/PAS 21448.

In our architecture, (cf. Figure 1) the *System Health Monitor (SHM)* is responsible for estimating the health of all capabilities

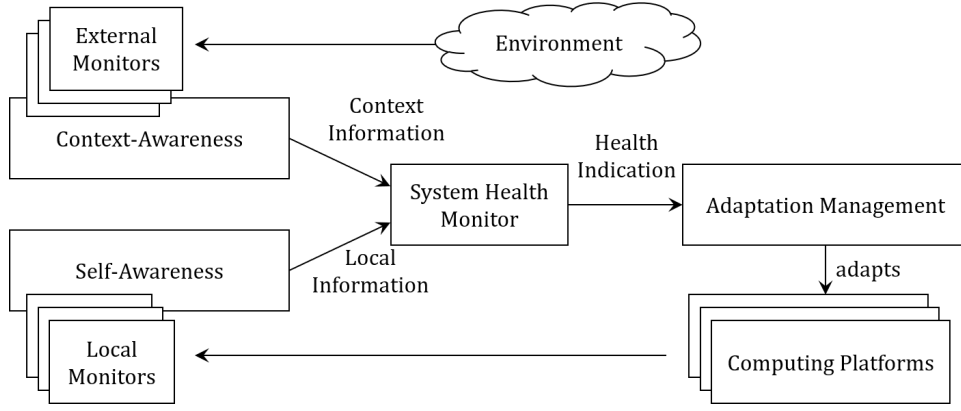


Fig. 1. Overview of the concept and involved architecture for enhanced system awareness

ties based on external context and internal monitoring results which are provided by local monitors. Depending on the current OD, different metrics for calculating the health can be applied. The SHM calculates so-called *Health Indicators (HI)*, which contain the result of the health analysis for capabilities, platforms, and domains. Autonomous driving functionalities like a Highway Pilot have dependencies to multiple capabilities deployed on different platforms. By exchanging HIs between platforms, the HI of a capability can be used to decide which capabilities currently fulfill the performance and safety requirements of the higher level functionality. The HI acts as a "Health-of-Service" attribute, which can be used for checking the fulfillment of safety contracts without deeper understanding of the function providing the capability. This is especially helpful in future Mobility System-of-Systems scenarios, where the systems are separated and the underlying configuration unknown, e.g., C2X scenarios with externally provides services.

D. Adaptation Management based on Enhanced System Awareness

In order to derive valid vehicle system configurations based on the system awareness provided as HI, we utilize service-contracts defined by a Domain Specific Language. These contracts include requirements that must be fulfilled to use the attached service as well as guarantees which the service implementation provides if all requirements are met. These requirements can be dependencies to capabilities, to system properties, or to the external context. For example, the usage of a service can be restricted to a specific OD, or a contract can state that for the using a certain object detection service, there must be a capability in the system that can deliver a video feed with adequate minimal resolution.

Besides needed requirements and guarantees, contracts may moreover describe performance attributes associated with services. This enables the Adaptation Management to not only find a configuration for which all requirements are fulfilled, but to find optimized configurations for the current situation with respect to individual metrics (e.g., speed, fuel consumption, etc.).

For this, the Adaptation Management utilizes the HIs generated by the SHM to compile a list of potentially safe and runnable services. Out of this list, a configuration of services is selected which fulfills all safety goals and requirements while aiming to optimize current performance goals, like fastest estimated time of arrival.

In addition, by linking services to ODs, our approach is able to manage capabilities proactively. If an ODM detects a change of the present OD is highly probable (e.g., entering a tunnel where limited sensor capabilities may be available), the system can prepare or even perform capability adaptation beforehand.

IV. SELF-ADAPTIVE AUTOMATED DRIVING AT INTERSECTIONS

For evaluating the benefits being added by enhanced system awareness, we apply our approach for autonomous driving at intersections. In our use case the ADS approaches a four-way intersection. Before entering the intersection, the vehicle needs to monitor its environment for approaching traffic participants that also want to cross to avoid accidents. At intersections, a different set of perception capabilities is required for effective monitoring of its environment than in other contexts, e.g., on highways. With the help of the ODM, the vehicle can determine its context and deploy the appropriate services providing the required capabilities.

In our intersection case study, we consider the capabilities *ObjectDetection* and *IntersectionClearance*. Additionally, a capability *IntersectionAssistant* is used to control the vehicle at the intersection. *IntersectionAssistant* depends on *IntersectionClearance*, which in turn depends on *ObjectDetection* to decide whether the intersection is free to enter.

ObjectDetection is implemented by multiple services making use of different sensors that provide guarantees with regards to the `detection range` at which they can reliably detect objects. Depending on the available range, a service providing *IntersectionClearance* is executed, which guarantees the `lookahead distance` of the service. Finally, this value influences the available *IntersectionAssistant* service and, thus, the maximum admissible driving speed.




Fig. 2. Automated driving at an intersection. The two cones visualize the *ObjectDetection* capabilities of the ego vehicle. The detection ranges provided by the services *LidarObjectDetection* and *CameraObjectDetection* are shown in red and blue, respectively.

Listing 1 Contract for *LongRangeIntersectionClearance*.

Contract *LongRangeIntersectionClearance*
provides *IntersectionClearance* **with**
lookahead distance = 50.0
requires *ObjectDetection* **with**
detection range ≥ 50.0 ,
health indicator ≥ 70
requires *OD* *four_way_intersection*

Which services are available for deployment depends on the HI of the services, calculated by the SHM. From these services, the Adaptation Management chooses the configuration that results in the highest possible driving speed. Thus, the ADS achieves maximum performance based on the information from self- and context-awareness.

Figure 2 shows a screenshot of a simulation in which we examined the intersection use case. The detection ranges provided by two *ObjectDetection* services, *LidarObjectDetection* and *CameraObjectDetection* are visualized. *LidarObjectDetection* provides a detection range of 50 meters. If this service is deployed, the contract of the service *LongRangeIntersectionClearance*, shown in Listing 1, is fulfilled. The vehicle can, thus, deploy this service and, consequently, also an *IntersectionAssistant* service that achieves a high driving speed.

Changes in the context or state of the vehicle system itself can lead to reduced HI values for some of the deployed services, which necessitates an adaptation to a different configuration. For example, a service health reduction for *LidarObjectDetection* could be caused by strong rains setting in, making the Lidar sensor less performant. Another cause could be an internal hardware failure in the Lidar device or a programming fault in its driver. In any of these situations, the SHM reduces the service health of *LidarObjectDetection* to such an extent that it can no longer satisfy the requirements of the service *LongRangeIntersectionClearance*. The adaptation manager therefore needs to degrade to the service *CameraObjectDetection* providing a detection range of only 25 meters. As the contract depicted in Listing 1 cannot be fulfilled by this service, further degradation is required. A less

performant *IntersectionClearance* service must be deployed, reducing the guaranteed lookahead distance and, in turn, the admissible driving speed.

This use case emphasizes that enhanced system awareness allows to choose the configuration of services that results in the highest possible driving speed for the current context and state of the ADS. Necessary degradations can be triggered automatically, leading to improved resilience of the vehicle system. If no valid configuration can be found, a pre-calculated and -configured minimal risk maneuver can be executed as safe fallback to preserve dependability.

V. CONCLUSION AND OUTLOOK

For autonomous driving the reasoning about a cars own capabilities and context will be a decisive challenge. We have introduced our approach of enhanced system awareness which includes self- and context-awareness of autonomous vehicles and integrates with present and upcoming automotive standards([5]). In the initial case study, we simulated the adaptation of an autonomous car's service configuration at an intersection based on the specific operational domain and available capabilities. By this, we could highlight how autonomous driving systems can benefit from enhanced system awareness.

In future, we plan to extend our approach, with respect to use cases, operational conditions, and capabilities, in order to further study and evaluate the impact of enhanced system awareness on the overall performance of autonomous driving.

REFERENCES

- [1] P. Koopman and M. Wagner, "Autonomous Vehicle Safety: An Interdisciplinary Challenge," *IEEE Intelligent Transportation Systems Magazine*, vol. 9, no. 1, 2017.
- [2] K. Czarnecki and R. Salay, "Towards a Framework to Manage Perceptual Uncertainty for Safe Automated Driving," *Computer Safety, Reliability, and Security (SAFECOMP)*, 2018.
- [3] A. Magnusson, L. Laine, and J. Lindberg, "Rethink EE architecture in automotive to facilitate automation, connectivity, and electro mobility," in *ICSE'18: Software Engineering in Practice*. ACM, May 2018.
- [4] J.-C. Laprie, "From Dependability to Resilience," in *38th IEEE/IFIP International Conference on Dependable Systems and Networks*, 2008.
- [5] AUTOSAR, "Specification of Health Monitoring," 2021, AUTOSAR R21-11 0.
- [6] P. Koopman and F. Fratrick, "How Many Operational Design Domains, Objects, and Events?" *SafeAI@AAAI*, 2019.
- [7] M. Horwick and K.-H. Siedersberger, "Strategy and architecture of a safety concept for fully automatic and autonomous driving assistance systems," in *2010 IEEE Intelligent Vehicles Symposium*, Jun. 2010.
- [8] D. Wittmann, C. Wang, and M. Lienkamp, "Definition and identification of system boundaries of highly automated driving," *7. Tagung Fahrerassistenz*, 2015.
- [9] I. Colwell, "Runtime Restriction of the Operational Design Domain: A Safety Concept for Automated Vehicles," Master's thesis, University of Waterloo, 2018.
- [10] D. Schneider et al., "WAP: Digital Dependability Identities," *2015 IEEE 26th International Symposium on Software Reliability Engineering (ISSRE)*, Nov. 2015.
- [11] The British Standards Institution BSI PAS 1883, "Operational Design Domain (ODD) taxonomy for an automated driving system (ADS) – Specification," 2020.
- [12] F. Dollinger, R. Asmus, and M. Dreiser, "System health indicators in mixed criticality e/e systems in automated driving context," in *Software Architecture. ECSA 2020.*, 2020.