A Strategy for Assessing Safe Use of Sensors in Autonomous Road Vehicles

Rolf Johansson[1,2], Samieh Alissa[3], Staffan Bengtsson[4], Carl Bergenhem[5], Olof Bridal[6], Anders Cassel[7], De-Jiu Chen[8], Martin Gassilewski[4], Jonas Nilsson[2], Anders Sandberg[9], Stig Ursing[3], Fredrik Warg[1], Anders Werneman[5]

[1]RISE, [2]Zenuity, [3]Semcon, [4]Volvo Car Corporation, [5]Qamcom Research & Technology, [6]Volvo Group Trucks Technology, [7]Autoliv, [8]KTH, [9]Delphi E&S

Published in:

Proceedings of 36th International Conference on Computer Safety, Reliability, and Security, SAFECOMP 2017,Trento, Italy, September 13-15, 2017

The final publication is available at Springer via http://dx.doi.org/10.1007/978-3-319-66266-4_10

A Strategy for Assessing Safe Use of Sensors in Autonomous Road Vehicles

Rolf Johansson[1,2]✉, Samieh Alissa[3], Staffan Bengtsson[4], Carl Bergenhem[5], Olof Bridal[6], Anders Cassel[7], De-Jiu Chen[8], Martin Gassilewski[4], Jonas Nilsson[2], Anders Sandberg[9], Stig Ursing[3], Fredrik Warg[1], Anders Werneman[5]

[1]RISE, [2]Zenuity, [3]Semcon, [4]Volvo Car Corporation, [5]Qamcom Research & Technology, [6]Volvo Group Trucks Technology, [7]Autoliv, [8]KTH, [9]Delphi E&S
✉rolf.johansson@zenuity.com

Abstract. When arguing safety for an autonomous road vehicle it is considered very hard to show that the sensing capability is sufficient for all possible scenarios that might occur. Already for today's manually driven road vehicles equipped with advanced driver assistance systems (ADAS), it is far from trivial how to argue that the sensor systems are sufficiently capable of enabling a safe behavior. In this paper, we argue that the transition from ADAS to automated driving systems (ADS) enables new solution patterns for the safety argumentation dependent on the sensor systems. A key factor is that the ADS itself can compensate for a lower sensor capability, by for example lowering the speed or increasing the distances. The proposed design strategy allocates safety requirements on the sensors to determine their own capability. This capability is then to be balanced by the tactical decisions of the ADS equipped road vehicle.

Keywords: — ISO 26262, automated driving systems, systematic system design faults, sensor systems, tactical decisions

1 Introduction

There is an increasing expectation for autonomous road vehicles to become available in a not too distant future, and there are many initiatives in the automotive industry aiming for such a development. One of the largest challenges is to come up with a safety case arguing that the automated drive system (ADS) feature is sufficiently safe, and has equal or better driving capability than the human driver it replaces. This is a difficult task, especially when it comes to the environment sensing capability. In order to claim that the verification and validation is sufficient, i.e. that the sensors of an ADS always have the capability implied by the task of ADS operation, all relevant scenarios, potential failure modes and corresponding hazards must be covered by relevant test cases. As the number of possible scenarios an ADS-equipped vehicle needs to handle is very large, it is hard to show that the set of test cases provide sufficient evidence that the vehicle will be able to act safely.

As a starting point when realizing an ADS feature, the functionality, limitations and system boundaries must be defined. Given this definition, the task of showing that the sensors will provide the information needed for the vehicle to behave sufficiently safe implies three challenges. The first is to identify the use scenarios, potential failure modes of the ADS feature, and carry out a hazard analysis and risk assessment. In the functional safety standard for road vehicles, ISO 26262, the result is a number of safety goals (top level safety requirements) each with an integrity level ranging from ASIL A (lowest) to ASIL D (highest), where higher integrity levels require higher confidence in the absence of failures. The second implication is the difficulty to reach the coverage of the argumentation required by the higher ASIL levels. The third is the implicit implication, begging the question if it is at all reasonable to provide the sensor capability necessary to reach the last part of the probability tail needed show coverage for the higher ASILs. One typical safety goal for an ADS might be to control the vehicle in such a way as to avoid collision with a vulnerable road user with a velocity that could cause a specific harm. In order to argue that this safety goal is satisfied, we need to show that the ADS has sufficient capability of the sensors that the vehicle can react in time to avoid a collision; under various drive scenarios, road and weather conditions. Achieving such a capability level can be both very difficult and very costly, but if we cannot argue that the ADS vehicle will fulfill this safety goal to the required ASIL level, we cannot introduce this vehicle on the market. Based on the reasoning above, a number of key questions need to be addressed for an ADS function:

- How to select a limited number of drive scenarios, carry out hazard analysis, define safety goals and be able to argue completeness.
- How to define sufficient sensing capability to sense the surroundings and drive scenarios in various conditions.
- How to define sufficient level of sensing redundancy given the limitations in technology and cost.

In this paper, we argue that these challenges are solvable. ADS-equipped vehicles need to perform tactical decisions [1,2], such as e.g. when to perform a lane change, preferred ego speed and preferred distance to vehicles in front . Therefore, we propose to allocate safety requirements on tactical decisions such that the safety goals can more easily be shown to be fulfilled. Instead of solving a statically defined sensing task, the sensors need to report what they can promise, and the tactical decisions can be adjusted by the ADS accordingly. In the example safety goal above regarding avoiding collision with vulnerable road users, the maximum safe speed of the vehicle will depend on the current ability of the sensors to detect vulnerable road users with enough confidence regarding the required ASIL. This means that we get on the one hand a framework for how to argue that the chosen sensors can have sufficiently specific test cases for showing completeness with respect to the safety requirements. On the other hand we get a pattern where we do not need to introduce unnecessarily expensive sensing solutions just to meet the very last part of the probability tail for higher ASILs.

The rest of the paper is organized as follows: Related work is discussed in Section 2, our solution is explained in detail in Section 3, and finally some conclusions and future work is provided in Section 4.

2 Related Work

Terminology and definitions in this paper aims to conform to taxonomy and definitions made in SAE J3016 [3]. This standard provides a taxonomy describing the full range of levels of driving automation in on-road motor vehicles and includes functional definitions for advanced levels of driving automation and related terms and definitions.

2.1 Standards addressing sensor performance

PAS 21448 - SOTIF[1] (Safety of the Intended Function) [4] is an ISO TC 22/SC 32/WG 8 initiative that proposes "guidance on the design, verification and validation measures applicable to avoid malfunctioning behavior in a system in the absence of faults, resulting from technological and system definition shortcomings". Hence, PAS 21448 claims that this (safety violations in a system without failure) is outside the scope of ISO 26262, and that this issue therefore requires additional guidance.

PAS 21448 proposes a process that aims to improve the nominal function specification to avoid or handle hazardous behavior due to the nominal function or technical limitations in the implementation. To ensure that the intended function is sufficiently safe PAS 21448 proposes a process to define and improve the function definition to reduce to an acceptable level the risk of:

- residual risk of the intended function, through analysis.
- unintended behavior in known situations through verification
- residual unknown situations that could cause unintended behavior, through validation of verification situations.

The updated intended function can be used as an input to the process of ensuring that functional safety is achieved using ISO26262. Although applicable to ADS, SOTIF is primarily focused on ADAS (Advanced Driver Assistance Systems) functions; that rely on environmental sensors such as camera and radar.

Another approach to addressing sensor performance is considered by IEC TC44 in the committee draft[2] "IEC 62998 CD – Safety-related sensors used for protection of person" [5], which aims to be complementary to the functional safety standards IEC 61508 [6], IEC 62061 [7] and ISO 13849 [8]. As opposed to PAS 21448, 62998 CD does not claim that safety violations due to the sensor system making hazardous deci-

[1] The SOTIF PAS is in working draft phase and is available for internal review as of Q1 2017. Hence the statements below are subject to change based on the outcome of the development process of the PAS.

[2] Since this description is based on a committee draft, the content is subject to change.

sions about the environment is outside the scope of the corresponding functional safety standard. Instead this draft provides guidance on how the safety-related sensors can achieve the integrity levels determined by the HA&RA performed in the main standard, when a sensor is used as a safety-related subsystem. The draft considers necessary performance, which static (e.g. environmental conditions or reliability) and dynamic (sensor readings and associated confidence) information must be available, and how to validate the sensors for enabling their use as part of a safety-related control system. The run-time confidence information can be used by the control system, to adapt the operation to the current capabilities of the sensors. The draft also gives guidelines on sensor fusion, and the separation of safety-related and automation-related information, where the latter is information needed for non-safety critical detection requirements.

2.2 Aerospace

Automation in the aviation world (Aviation Automation Systems (AAS)) plays a pivotal role nowadays. Its presence on board airplanes is pervasive and highly useful in improving the pilots' performance and to enhance safety. When developing safety critical system in aerospace, MIL-STD-882 [9] shall be used. This standard uses a prescriptive process that details the steps that shall be taken. The methods employed in this standard are qualitative, quantitative, or both. The development process based on MIL-STD-882 is iterative in nature. The process begins with concept design and derives an initial set of safety requirements. During design development, if any changes are made, and the modified design must be reassessed to meet safety objectives. This may create new design requirements. These in turn may necessitate further design changes. The safety assessment process ends with verification that the design meets safety requirements and regulatory standards.

In general, AAS is not designed to be responsible completely for safe operation of the aircraft. This implies that if the AAS fails, the pilot has responsibility e.g. to safely land the airplane. This is a similar relation as between the human driver and ADAS functions. The vehicle environment (ADAS and ADS), on other hand, is less cooperative than the environment in the air; meaning that Air Traffic Control (ATC) plays an important role to maintain safety, e.g. ensure the adequate separation of the airspace. Airplanes do not need to "stay on the road"; as long as they are at normal flying altitude there are no obstacles to avoid, no lane to follow, hardly any flying objects to avoid. Since an airplane operates in three dimensions it is less likely that two randomly flying objects will collide. A minor collision avoidance effort is required for the airplane, e.g. a simple radar based avoidance algorithm. The number of scenarios and tactical decisions in AAS are relatively low and the tactical decisions are the responsibility of pilot and the ATC in controlled airspace. If something happens, the human pilot usually has in the order of minutes to react. In a vehicle (ADAS and ADS), the reaction time for the driver is in the order of seconds. This leads to high requirements on the human driver (ADAS) or ADS. Further, due to the shorter time-frame any sensing may have lower precision. This is mainly a challenge for ADS since with ADAS, omission of function is generally not an issue.

In contrast to AAS and ADAS (which assume human operator(s) e.g. ATC, pilot or driver), ADS have low controllability. A challenge in ADS is to automate behavior planning on a tactical level. This has different challenges related to the complexity of real world traffic situations. Based on the above, we argue that ADS are more complex than AAS or ADAS. Table 1 gives an overview of the comparison between AAS, ADAS and ADS.

Table 1. Comparison between AAS, ADAS and ADS.

	AAS	**ADAS**	**ADS**
Environment	Cooperative	Not necessarily cooperative	Not necessarily cooperative
Sensing precision	High	Low	Low
Controllability level	High	High	Low
Scenarios number	Low	High	Low
Tactical decisions number	Low	Low	High
Failure severity	Catastrophic	Normal	Severe

2.3 ISO 26262 safety requirement refinement

The specification of requirements at different levels (Safety Goals (SG), Functional Safety Concept (FSC)/ Functional Safety Requirements (FSR), Technical Safety Concept (TSC)/ Technical Safety Requirements (TSR)) are described in the ISO 26262 standard. The definitions of SG, FSC and FSR are:

- SG is a top-level safety requirement as a result of a hazard analysis and risk assessment
- FSR is specification of implementation-independent safety behavior, or implementation-independent safety measure, including its safety-related attributes.
- FSC specification of the functional safety requirements, with associated information, their allocation to architectural elements, and their interaction necessary to achieve the safety goals

The flow and sequence of the safety requirement development is illustrated in the figure below.

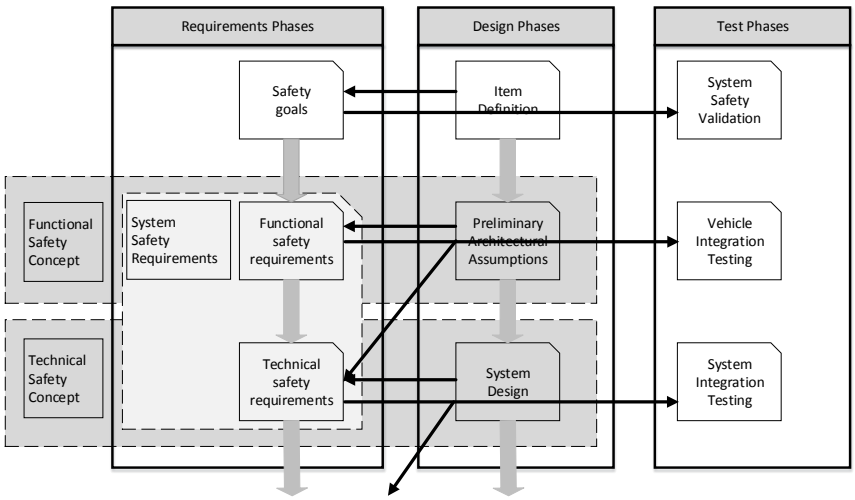

Fig. 1. Safety requirements, design and test flow. Excerpt from ISO 26262

A critical factor for achieving a safe ADS is a proper division of responsibility between sensors and sensor fusion blocks. In a sensor fusion system a significant effort has to be put into evaluating the redundancy needed to take into account in order to get a balanced and efficient design fulfilling all safety requirements. This is necessary when for example evaluating operational capabilities (redundancy and degradation concepts). Effort is put into identifying a methodology bridging the discrete domain of ASIL and the continuous domain of probabilities.

Different Sensors have different efficiency in interpreting different objects, and are even further differentiated with environmental conditions (rain, snow, dust). Radar for instance may show absence of objects (soft tissue) even when there is an object present, but if it shows presence of an object it is very accurate.

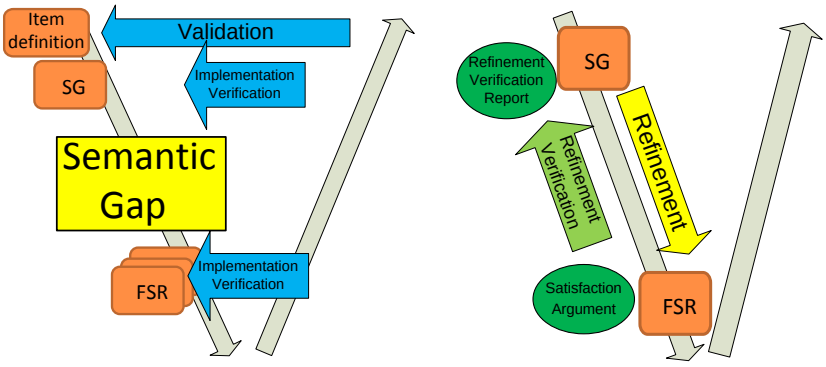

Fig. 2. The semantic gap (left) & Activities between adjacent requirement levels (right)

The "distance" between two requirement levels e.g. SG (Safety Goal) and FSRs (Functional Safety Requirement), or any other adjacent requirements levels, is denoted the Semantic Gap, see Fig. 2 (left) above. The concept phase of ISO 26262 describes how SGs are determined from the results of the HA&RA (Hazard Analysis and Risk Assessment). SGs are refined into FSRs, which implies that the SGs can be interpreted as top-level safety requirements in a layered requirement hierarchy. An SG is a high-level description of an objective on vehicle level, and the refinement of the SGs into an FSC (Functional Safety Concept, i.e. a set of FSRs allocated to architectural elements) may need a substantial amount of assumptions, domain knowledge and other input. If no or only weak arguments for the refinement of SGs to FSC exist, then verification to argue correctness and completeness is at best difficult.

A requirement (the upper level of two adjacent requirement levels) is refined into a composition of lower level requirements and rationale, known as satisfaction arguments. The satisfaction arguments shall be collected for the composition, see Fig. 2 (right). This bridge of information should "fill" the semantic gap. Satisfaction arguments may be e.g. assumptions, domain knowledge, design patterns. This is essential in almost every nontrivial refinement. The rationale justifies the "refinement path taken" through the semantic gap and improves to traceability. This is further discussed in [10].

Filling the semantic gap is particular challenging with systems that use sensor fusions due to the incomplete and overlapping redundancy that different sensors may have. To prove the correctness of the refinement, the satisfaction argument must document the overlaps and gaps between sensors. To attain this knowledge deep understanding needs to be collected with e.g. field tests and simulation of sensor properties.

3 Arguing Safety for an ADS

The safety analysis procedure, as defined by the ISO26262, is to define the Item (i.e. the overall system providing a specific vehicle-level function), carry out a hazard analysis and risk assessment (HA&RA) of the Item's potential failure modes, assess

the risk, and define associated safety goals with appropriate ASIL values. A safety goal expresses the goal of preventing or avoiding a hazardous event, i.e. a combination of a specific failure mode and a situation in which this failure mode may be hazardous. Each hazardous event shall be covered by a safety goal. The task of performing hazard analysis for an ADS and make sure the set of safety goals is complete and correct, taking into account all driving scenarios the function is designed for, is a big challenge in itself. This is still subject to research and not addressed in this paper, but initial attempts have been made [11] and it is reasonable to expect more guidance for this task in the future. In the section below we assume that safety goals have been identified, discuss what kind of safety requirements that might be allocated to the sensors, and how to argue that we can get evidence that all such safety requirements are met at all times under all conditions. As an example, we present some safety goals that might be hard to show that they are always fulfilled.

The functional safety concept (FSC), including functional safety requirements (FSRs), is derived from the safety goals, typically based on function analysis of the Item, evaluation of suitable architectural design patterns and other system design factors. FSRs are typically allocated to sensors, control system and actuators. The arguing we propose is illustrated by one typical architectural pattern for the functional safety concept, and it is then further described what is needed from the different blocks of that architecture. Note that the general argumentation in this paper is not limited to the architectural choice of the example.

3.1 Some Safety Goals for an ADS

In an ADS-equipped vehicle, the responsibility of driving is moved from the human driver to an ADS. In order to replace the driver, the ADS must be able to sense the surroundings (road objects, road conditions, weather conditions), control the vehicle behavior, assess any risks that lie ahead and make appropriate decisions. It must be at least equally capable as the human at operating the vehicle without causing any accidents and, as far as reasonable, mitigate dangerous situations caused by other road users. The phrase “Avoid collision” in our example safety goals below shall be interpreted with this definition in mind. Introducing the ADS function, means that we need to analyze and specify the functional capability maneuvering the vehicle equal to an experienced driver, taking in to account potential hazards caused by the limitations of the ADS and its ability in sensing the surroundings. The resulting safety goals needs to be specific and shall express what shall be avoided, e.g. a condition or functional limit, which could cause a hazard. For instance, it is difficult to assign a fair ASIL to a very general safety goal like ‘never collide’. The severity of collisions with different objects and at different impact speeds varies. Hence, we need to refine the safety goal on avoiding collisions into a number of specific ones. In this paper, we do not discuss further how to best elicit a set of safety goals, we merely use a few examples to illustrate the reasoning needed to show the safety goals we have are fulfilled. Here we chose a set of safety goals, for the ADS, addressing the importance not to collide with another vehicle in front.

SG I. Avoid collision with a higher impact speed than 65 *km/h* with a vehicle being on the road in front, ASIL D

SG II. Avoid collision with a higher impact speed than 40 *km/h* with a vehicle being on the road in front, ASIL C

SG III. Avoid collision with a higher impact speed than 25 *km/h* with a vehicle being on the road in front, ASIL B

SG IV. Avoid collision with a higher impact speed than 15 *km/h* with a vehicle being on the road in front, ASIL A

The difference between the safety goals in this example comes from the severity factors of the HA&RA. The numbers used are just examples, and could of course be determined differently. The important thing is that different impact speeds and collisions with different object types will result in different ASIL values.

3.2 Generic Architectural Pattern for ADS

On a very generic high level we could derive safety requirements on the sensors from the safety goals based on a conceptual architecture like the one in Fig. 3 below.

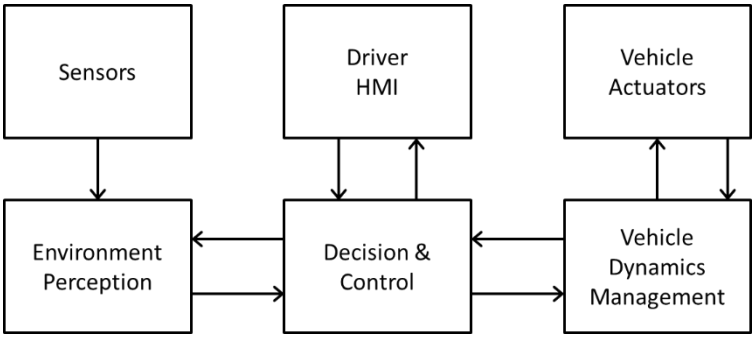

Fig. 3. Conceptual architecture for ADS

When defining the functional safety concept it is critical to show that all safety goals are completely and correctly refined into functional safety requirements allocated to the elements of the FSC architecture. This means that the division of responsibility between the Environment Perception block and the Decision & Control block must support the fulfilment of all the safety goals.

We assume that the Environment Perception block has a functional requirement: "The Environment Perception block shall detect vehicles in front of the subject vehicle".

What is needed as output from the Environment Perception block depends on the needed input to the Decision & Control block to fulfill its duties. Let us for now assume that the Decision & Control block would be satisfied if the Environment Percep-

tion block always could fulfill the following functional safety requirements, still just addressing the four safety goals in the example above:

FSR I. No omission of detected vehicles in front up to a distance of 30 *m*, ASIL D
FSR II. No omission of detected vehicles in front up to a distance of 40 *m*, ASIL C
FSR III. No omission of detected vehicles in front up to a distance of 50 *m*, ASIL B
FSR IV. No omission of detected vehicles in front up to a distance of 60 *m*, ASIL A

Note that these functional safety requirements might be needed to show that the safety goals are completely and correctly refined. They are requirements that need to be fulfilled regardless of what set of sensors that are chosen in the detailed implementation. In further refinement into a technical safety concept, it is important to show that these FSR always are fulfilled. As discussed above in the paragraph on ISO26262 refinement, each refinement step needs a convincing argumentation.

Assuming that the above set of functional safety requirements are identified as applicable for the aggregated and fused sensor information employed in the realization of the Environment Perception block above, we need to argue that the chosen sensors made explicit on lower levels, can provide evidence to always and under all applicable conditions fulfill these functional safety requirements. If we choose a set of sensors for which we cannot show compliance with all these FSRs, we will not be able to fulfill the refinement verification step. And we could hence regard such a design candidate as introducing a design fault: the E/E functionality will not fulfill all safety goals. This is not due to the presence of random hardware faults, but rather due to systematic faults in the system design process.

A way to describe the challenge of showing that the sensors of an ADS have been tested enough to show capabilities of all applicable scenarios according to its operational design domain, ODD, [3] is to say that we need to address all system design failures for the sensor blocks. Of course, we still have to show that the contributions from random hardware faults and from systematic hardware and software faults are restricted according to the ASIL levels, but we also need to obtain sufficient confidence that systematic system design faults are avoided. This is nothing new for sensors. It is the same problem whenever we choose to implement a given architectural element with a specific component. If we for example chose a CPU with insufficient processing power for the tasks allocated to it, this is to be regarded as a systematic system design fault.

To summarize the big challenge for sensor systems for ADS, it is to show sufficient confidence in the avoidance of systematic design faults, such as selecting sensors that do not meet the required capabilities according to the ODD, and hence violate the functional safety requirements. In order to argue for such freedom of systematic design faults, we need evidence covering all ODD applicable scenarios and environmental situations. To collect such a large set of evidence is very hard, and it is also very hard to argue that the set of performed test cases etc, is complete with respect to this purpose.

3.3 Decision Hierarchy

One task that is characteristic for an ADS equipped vehicle is the need to perform tactical decisions, which in this example, is part of the Decision & Control block (Fig. 3). Examples of those are what target speed to aim for, what distance to aim for to the vehicles in front, what lane to choose, whether to perform an overtake, etc. Such tasks are generally left to the driver for manually driven vehicles, even when having ADAS functionality. ADAS is in general focusing on operational decisions. For an adaptive cruise control (ACC) it might be expected from the driver to choose a target speed and some preferred distance to vehicle in front. Then the ACC in the operational decisions determines what vehicle speed this implies in the actual situation, taking care of how fast the vehicles in front are driving etc.

Returning to the example safety goals above, they can be more easily shown to be fulfilled if we introduce requirements on the tactical decisions. Let us for example assume four FSRs to be allocated to the Decision & Control block:

FSR V. Vehicle speed request shall be limited to a vehicle speed implying a maximal impact speed of 65 *km/h* in any vehicle objects on road, ASIL D

FSR VI. Vehicle speed request shall be limited to a vehicle speed implying a maximal impact speed of 40 *km/h* in any vehicle objects on road, ASIL C

FSR VII. Vehicle speed request shall be limited to a vehicle speed implying a maximal impact speed of 25 *km/h* in any vehicle objects on road, ASIL B

FSR VIII. Vehicle speed request shall be limited to a vehicle speed implying a maximal impact speed of 15 *km/h* in any vehicle objects on road, ASIL A

This is functional safety requirement expression based on the four safety goals. But we also say that we can give the full responsibility to the tactical decision to what is considered safe in the current situation. Instead of giving an absolute static task to the Environment Perception block, we now say that if the Environment Perception block reports what it can promise, we can then compensate for this by means of the tactical decisions.

Assume the Decision & Control block has a functional requirement: “The Decision & Control block shall output world coordinates identifying where there is a higher/lower risk of a hazard”.

Given that we have the above FSRs on the Decision & Control block, we can revise the FSRs for the Environment Perception block to:

FSR I. No omission of detected vehicles in front inside the stated coordinates, ASIL D

FSR II. No omission of detected vehicles in front inside the stated coordinates, ASIL C

FSR III. No omission of detected vehicles in front inside the stated coordinates, ASIL B

FSR IV. No omission of detected vehicles in front inside the stated coordinates, ASIL A

At a first glance this might look like redundant requirements, but this is not the case. What it says is that the Environment Perception block needs to state inside what boundaries that each of the ASIL attributes apply. Instead of giving static horizon distance requirements as in the example above (30m ASILD, 40m ASILC, 50m ASILB, 60m ASILA), we ask the block to dynamically determine the borders for each ASIL with respect to systematic design fault.

On the one hand we give a possibility for the sensor system to cut the assumed very high extra cost for the probability tail for higher ASIL, but on the other hand we ask it to determine itself what the ASIL capability is. What we say is that we can handle temporary reductions of sensor capabilities as long as the system is continuously aware of the current sensor capability. If we can determine the capability limit in run time, we can by the tactical decision assure to avoid systematic design faults.

3.4 Dynamic ASIL Capability Maps

Above we list four different FSRs to be allocated to an Environment Perception block all formulated as: 'No omission of detected vehicles in front of ego vehicle inside the stated coordinates' with different ASIL values. This implies that for each of these requirements the EP block is expected to provide a map of the boundaries that apply for the respective ASIL value. Collecting such map data into one map showing information of the integrity of claiming absence of a certain object type can look like what is depicted in Fig. 4 below. The white area in this map, closest to the vehicle, is where the perception block can provide the highest confidence (ASILD) that there is no object of the specified kind. The black area in the map, most far away from the vehicle, shows that this is beyond the high-integrity sensing horizon. There might be object here; at least there is no possibility for the Environment Perception block to claim the absence of the actual object.

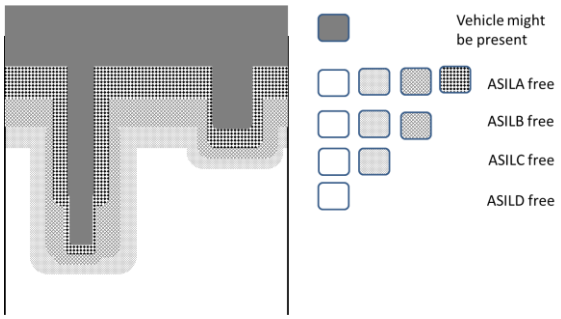

Fig. 4. Example map w.r.t. claiming absence of object of type Vehicle. The ego vehicle is in the bottom of this figure

For each type of object this kind of dynamic ASIL capability map is provided by the Environment Perception block. Furthermore there is also one such map telling the ASIL level of claiming the presence of each object kind.

By defining statically in design time exactly what object types to consider, the task for the Environment Perception block is to provide all the dynamic ASIL capability maps for claiming the absence and presence respectively of each object type. Then the Decision & Control block, by means of tactical decision, can make sure that no systematic design fault w.r.t. sensor capability will violate any safety requirement.

It is beyond the scope of this article to show in detail how an Environment Perception block in run-time can produce such maps showing the ASIL-limits of absence of each object category. The message from this paper is that by formulating the safety requirements for an Environment Perception block in this way, it is possible to argue safety by asking the tactical decisions to find a safe alternative for any set of maps. A detailed strategy how to construct these maps is subject to current research. However, basic means will be design-time information from each sensor based on extensive test under various conditions, geographical data and sensing of environmental conditions showing what conditions that the sensors are (not) facing, and run-time analysis of how the different (redundant) sensors in the Environment Perception block are more or less consistent in their observations.

4 Conclusions

This paper outlines a strategy that enables and supports the argumentation in a safety case that the ADS is sufficiently safe, despite the difficulty to show sufficiently safe environmental sensing capability. A key factor in such a safety case is to argue that at all instances the vehicle will balance the sensing capability with the tactical decisions. The higher sensor capability, the higher performance (e.g. vehicle speed on the road) can be shown safe.

We argue that the ADS may need different sensor capabilities depending on the driving conditions and the chosen style of driving. If the sensors in order to enable safety for the vehicle are required to have a capability higher than they actually have, we denote this systematic design fault. This means that when performing a refinement of safety requirements, we have not allocated such safety requirements on all blocks (for example a sensor block), that can be shown to be fulfilled at all conditions.

Instead of allocating a static set of safety requirements on the sensors w.r.t. systematic design faults, we propose to formulate the safety requirements in terms of dynamic ASIL capability maps. This means that for each type of object, the sensing systems are required to provide a map showing the limit for each ASIL confidence level. If such maps can be provided to the block in charge of tactical decisions, the overall ADS including both sensors and tactical decisions can be argued to behave safely at all times.

Acknowledgements. The research has been supported by the Swedish government agency for innovation systems (VINNOVA) in the ESPLANADE project (ref 2016-04268).

References

1. R. Sukthankar, “Situation Awareness for Tactical Driving”, Ph.D. thesis, Robotics Institute, Carnegie Mellon University, USA, January1997.
2. T. X. P. Diem and M. Pasquier, “From Operational to Tactical Driving: A Hybrid Learning Approach for Autonomous Vehicles”, 2008 10th Intl. Conf. on control, Automation, Robotics and Vision, Hanoi, Vietnam, December 2008.
3. SAE International, “SAE J3016: Taxonomy and Definitions for Terms Related to Driving Automation Systems for on-Road Motor Vehicle.” 2016.
4. ISO / AWI PAS 21448 Road vehicles -- Safety of the intended functionality.
5. IEC 62998 CD – Safety-related sensors used for protection of person.
6. IEC 61508 (2010): Functional safety of electrical/electronic/programmable electronic safety-related systems. The International Electrotechnical Commission.
7. IEC 62061 (2015): Safety of machinery – Functional safety of safety-related electrical, electronic and programmable electronic control systems. Ed. 1.2. The International Electrotechnical Commission.
8. ISO 13849 (2015): Safety of machinery -- Safety-related parts of control systems. International Organization for Standardization.
9. MIL-STD882(E): Department of Defense Standard Practice, System Safety, May 2012.
10. Carl Bergenhem, Rolf Johansson, Andreas Söderberg, Jonas Nilsson, Jörgen Tryggvesson, et al.. How to Reach Complete Safety Requirement Refinement for Autonomous Vehicles. Matthieu Roy. CARS 2015 - Critical Automotive applications: Robustness & Safety, Sep 2015, Paris, France. 2015.
11. Warg, F., Gassilewski, M., Tryggvesson, J., Izosimov, V., Werneman, A. and Johansson, R., Defining autonomous functions using iterative hazard analysis and requirements refinement. In *SAFECOMP Workshops, SASSUR* (pp. 286-297). Springer, 2016.