ISO 26262: SEooC - Safety Engineering out of Control ?

Nicholas Mc Guire and Andreas Gabriel-Platschek

[1] OSADL, Safety Critical Linux Working Group
Im Neuenheimer Feld 583, D-69120 Heidelberg, GERMANY
<safety@osadl.org>
[2] OpenTech EDV Research GmbH
Augasse 21, 2193 Bullendorf, AUSTRIA
<andi@opentech.at>

Abstract. Autonomous driving is a fascinating technology - no doubt about that. But aside from functional considerations, bringing such technologies into the public space mandates due diligence on the safety side of the automotive industry. 1.5 metric tones bolting down a highway at 30m/s are hazardous - and controlling this hazard requires highly-reliable cognitive systems to keep this hazard at a tolerable level.
Humans might not be the most reliable cognitive systems for simple tasks but they are darn good compared to computers for complex tasks. A naive "computers are better - Deep Blue even beat Kasparov" with some marketing pressure added and unrealistic expectations raised in public are a good recipe for trouble.
If the plan is to bring autonomous driving to the public space a sound and agreed on definition of safety for such systems and their elements is needed. Currently the lack of such a standard is a fundamental issue that needs to be addressed if autonomous vehicles are to be a benefit for society.

1 Introduction

This is not an attempt at an exhaustive description of IEC 61508 Ed 2 or any of its derived standards, but rather it is an attempt to present one of the core ideas of the approach - managing complexity - and put it into context of the on-going systematic, violation of this concept by abusing ISO 26262 [1] for autonomous vehicles which was conceived in the context of typical usage at time of writing - effectively pre-2011 and thus more along the lines of classical automotive systems. From the information that is publicly available up to now - the upcoming Ed 2 will **not** solve this fundamental problem either.

Now, some folks in the industry will cry-out on the above statement that ISO 26262 is abused for autonomous vehicles and this may seem like a unnecessarily harsh statement. We hope though, that our concerns become clear after reading this paper, as the major part (see Section 3) of it is dedicated to present

some thoughts on this statement. But first the basic knowledge that is the pre-requisite for creating a new standard is discussed in Section 3, as it has to be assumed that many readers are not familiar with this process.

The conclusion (Section 4) not only summarizes issues highlighted by this paper, but also suggests what the automotive industry should do in the authors opinion.

2 Creating a new standard

To get started, a summary of the basis that is necessary for creating a new standard is given. The necessary known elements are then mapped against the state-of-the-art of autonomous driving as well as (the context of) ISO 26262 [1].

Not too surprisingly, guidance on how to approach writing a new standard can be found in a standard, namely ISO/IEC 51 [3] which lists the following pre-requisites that have to apply in order to being even able to write a standard for a certain domain:

- *detailed working knowledge of the product or system;*
- *requirements and guidelines from various origins, both general and specific to the standard development;*
- *human behavior studies and anthropometric data;*
- *injury/incident data of defects, and the recall history of the product or system;*
- *knowledge of the potential health and environmental effects of the product or system;*
- *feedback based on experience of end users of the product or system;*
- *knowledge of the potential risk reduction measures (protective measures);*
- *knowledge of the implications of possible future developments of the product or system;*
- *industry standards and guidelines;*
- *best available expertise and scientific advice from relevant stakeholders;*
- *legal requirements.*

[ISO/IEC GUIDE 51:2014 7.3.1]

Of this list, not one item can be statisfied by ISO 26262 in the context of autonomous vehicles. More importantly, not one of these aspects has been considered in the context of autonomous vehicles when ISO 26262 was created as a summary by Ward [4] suggests.

The implications of this are quite straight forward - if autonomous vehicles were not considered during writing of ISO 26262, then the context does not fit and ISO 26262 is not applicable - it is simply the wrong standard.

So what can be done when there is no applicable (domain) standard? In that case the first thing to do is to go up the standard hierarchy (i.e. at the standard from which ISO 26262 was derived) and look at IEC 61508 Ed 2 which is suitable

for complex systems. The core idea of IEC 61508 with respect to managing complexity is best expressed in the way IEC 62061 [5] harnesses IEC 61508:

The subsystem shall be realized by either selection (see 6.7.3) or design (see 6.7.4) in accordance with its safety requirements specification (see 6.6.2.1.7), taking into account all the requirements of 6.2. Subsystem(s) incorporating complex components shall comply with IEC 61508-2 and IEC 61508-3 as appropriate for the required SIL.

[IEC 62061 Ed 1 6.7.2.1]

The core idea is to ensure safety by adequate analysis and testing to the point where the element can actually be understood in the specific context sufficiently well to derive statements about the relevant subset of behavior affecting overall safety. It is not about compliance in any way - it is about knowledge.

The "trick" IEC 62061 [5] (and other derived standards like IEC 61511 [6]) do is then to allow this understanding of the element to be used to divide and conquer:

Where the design of a subsystem incorporates a complex component (as a subsystem element) which satisfies all relevant requirements of IEC 61508-2 and IEC 61508-3 in relation to the SILCL, it can be considered as a low complexity component in the context of a subsystem design since its relevant failure modes, behavior on detection of a fault, rate of failure, and other safety-related information are known. Such components shall only be used in accordance with its specification and the relevant information for use provided by its supplier.

[IEC 62061 Ed 1 6.7.4.2.3]

The complex component was transformed to a well defined and understood - managed - set of properties, not by compliance. This allows to build up complex systems without loss of control. The complexity reduction is not some kind of magic but achieved by applying a well structured framework that allows encoding commonalities in a well defined context.

The idea of standards is to encode such generic knowledge in a consolidated form agreed upon by the industry, the goal being that future engineering efforts can build on this foundation and by doing so utilize this knowledge to manage the complexity of systems. Compliance to abstract measures will not do much good to generate the necessary knowledge and without the understanding it is not possible to perform any reasonable analysis or testing of systems. With other words, without adequate knowledge safety is not attainable.

3 Whats going wrong?

In the automotive industry (but not exclusively there), there still is this somewhat absurd notion of table-driven safety in the heads of safety engineering departments - and some clauses in ISO DIS 26262 [2] seem to even advocate the same - where it somewhat unfortunately states:

"A rationale shall be given that the selected combination of methods complies with the corresponding requirement. If all highly recommended methods listed for a particular ASIL are selected a rationale needs not to be given."

[ISO DIS 26262 2009]

This is slightly less prominent in the final version of ISO 26262 2011 [1] where it then states:

For consecutive entries, all methods shall be applied as recommended in accordance with the ASIL. If methods other than those listed are to be applied, a rationale shall be given that these fulfil the corresponding requirement.

[ISO 26262-4 2011 4.2 Interpretations of tables]

Where it then also permits a "rational" *"based on the methods listed in the table is sufficient"* in the ensuing note:

NOTE A rationale based on the methods listed in the table is sufficient. However, this does not imply a bias for or against methods not listed in the table.

Even though the relevant software safety life-cycle requirements IEC 61508 Ed 2 part 3 [7] very clearly states, that this "table-driven safety" is a no-go:

For each life-cycle phase, appropriate techniques and measures shall be used. Annexes A and B provide a guide to the selection of techniques and measures, and references to IEC 61508-6 and IEC 61508-7. IEC 61508-6 and IEC 61508-7 give recommendations on specific techniques to achieve the properties required for systematic safety integrity. Selecting techniques from these recommendations ***does not guarantee by itself*** *that the required safety integrity will be achieved.*

[IEC 61508-3 Ed 2 7.1.2.7]

The problem is that safety is about encoding and tracing "why X was done" and even if it is the most obvious of things to do a rational captures this "why" and only with this "why" documented can we say in retrospect that someone actually knew what they were doing. Dropping the rational cuts this ability to detect high-level organizational faults in the process and is a fundamental no-go for any reasonable safety process.

As a secondary reason - even well established methods can go out-of-scope during retrofitting and maintenance — a rational would allow this misfit to be at least detectable.

Finally, if one prefers to ignore other standards, the safety community also has expressed similar views in less formal terms, very clearly. When Kelly [8] states that *The responsibility should then be placed on the software and systems developer to* ***present an argument as to why*** *their systems are safe and how they* ***meet the fundamental intent*** *of the standards* - compliance in an ***unthinking manner*** [8] is not going to help much.

Deriving approved generic combinations (e.g. as found in EN 50128 Ed 2) merely by SIL is possible, if and only if, the context is well enough defined. This only is achieveable if the product family is well understood in a known

context. With these pre-requisites a sector/product standard can actually be encoded faithfully describing reasonable assumptions and constraint — again these conditions do not hold for autonomous driving.

We do not yet understand it and even foundational questions are still up in the air [9]. Preventing bit-flips in memory if we are using algorithms we do not understand in environments we do not fully comprehend is a quite useless exercise - maybe even counterproductive as efforts are wasted to no avail.

Even at a very formal level ISO 26262 is hardly suitable for autonomous vehicles. ISO 26262 is titled "Road vehicles – Functional safety –" parts 1 through 9 being normative and part 10 being informative. The scope section - reproduced in all 10 parts states:

Scope
ISO 26262 is intended to be applied to safety-related systems that include one or more E/E systems and that are installed in series production ***passenger cars*** *with a max gross weight up...*

Part 1 "Vocabulary" of ISO 26262 goes on to the define relevant terms in a normative manner for the purpose of this standard:

1.85 passenger car
Vehicle designed and constructed primarily for the carriage of persons and their luggage, their goods, or both, having not more than a seating capacity of eight, ***in addition to the driver****, and without space for standing passengers.*

This standard was written for systems operated by a driver and applying this standard out-of-context will results in a "Safety Element out of Control".

ISO 26262 was derived from IEC 61508 (unfortunately largely Ed 1) for the then state-of-the-art of the automotive industry - micro-controllers and relatively simple OSEK-style software stacks - and it encapsulated, much as the intent of IEC 61508 is, the consolidated knowledge of the industry to allow streamlining terminology, techniques and procedures to improve the overall level of functional safety in that implied context. Applying this knowledge-base to systems that have orders of magnitude higher complexity and in some parts are not even understood well at the scientific level while claiming that ISO 26262 represents the relevant state-of-the-art is frankly absurd.

The authors of ISO 26262 Ed 1 were looking into the future as well while encoding the state-of-the-art. The introduction all parts of ISO 26262 identifies this mind-set and perceived future quite nicely:

Safety is one of the key issues of future automobile development. New functionality not only in the area of driver assistance but also in vehicle dynamics control and active and passive safety systems increasingly touches the domain of safety engineering. Future development and integration of these functionalities will even strengthen the need of safe system development processes and the possibility to provide evidence that all reasonable safety objectives are satisfied.

[ISO 26262 Ed 1 Introduction]

While driver assist systems undoubtedly can achieve quite significant complexity and malfunctioning of the same can be the initiating event for hazards of all kind, there is still a trained human operator who is expected to retain the necessary contextual bits and intervene where the assist function goes off into the wild. Consequently the assist functions have been rated at no more than ASIL-A - the lowest automotive safety integrity levels — most of them though actually are rated QM.

Autonomous driving is not a linear aggregation of assist functions in any way (even if this can be the perception to the untrained observer) — not technically, not legally, not with respect to the know-how needed at the organizational level. It is a new can of worms and it needs motivated systems engineers to treat it as such by systematically addressing the scientific, technical and regulatory challenges at hand. Or as Kelly [10] states under the heading "Theres no Substitute for Thinking": *Mere compliance with standards is a weak motivation.*

Building safe systems will mandate asking context specific questions, from accident models appropriate to address the given complexity [11], mitigation of fundamental technological hazards [9] to answering security and legal [12] issues — only moving forward when these are suitably answered.

4 Conclusion

This industry needs to come together, apply guidance (like that of ISO/IEC 51) to define and encode a suitable functional safety standard for autonomous vehicles along with a few other side-dishes like security and legal aspects - consolidate the knowledge base and then determine suitable, by then approved, combinations of state-of-the-art methods to ensure customer and public safety. IEC 61508 may be a suitable starting point with respect to functional safety objectives - in any case measures and techniques suitable for the target systems — autonomously piloted vehicles — need to be developed along with the overall life-cycle adjustments including organizational issues, to allow such systems to be fit for the public.

It is time for the automotive industry to accept the plain, simple (and normative) fact: formal conformance to a standard that is not suitable due to covering a completely different context does not entail safe systems. After all that should be our top-level goal.

References

1. ISO 26262: *Road vehicles – Functional safety – all parts*, Editon 1, IEC, 2011
2. ISO 26262 DIS: *Road vehicles – Functional safety – all parts*, Editon 1 DIS, IEC, 2009
3. ISO/IEC 51: *Safety aspects – Guidelines for their inclusion in standards*, Edition 3, IEC, 2014
4. David D. Ward,*ISO DIS 26262 - The new automotive functional safety standard*, in Safety Systems 19-1, 2009

5. IEC 62061: *Safety of machinery — Functional safety of safety-related electrical, electronic and programmable electronic control systems*, Edition 1, IEC, 2005
6. IEC 61511: *Functional safety – Safety instrumented systems for the process industry sector – 3: Guidance for the determination of the required safety integrity level*, Edition 1, IEC, 2007
7. IEC 61508:*Functional safety of electrical/electronic/programmable electronic safety-related systems – Part 3: Software requirements* , Editon 2, IEC, 2010
8. Tim Kelly*Software Safety - by prescription or argument ?*, in Safety Systems 7-2, 1998
9. Anish Athalye, Logan Engstrom, Andrew Ilyas, Kevin Kwok, *Query-efficient Black-box Adversarial Examples*, `https://arxiv.org/abs/1712.07113`, 2017
10. Tim Kelly*Theres no Substitute for Thinking*, in Safety Systems 20-3, 2011
11. Nancy G. Leveson,*The Need for New Paradigms in Safety Engineering*, in Safety Systems 17-2, 2008
12. Autonomes Fahren - ein Einblick in die rechtlichen Rahmenbedingungen, in Wirtschaftsinformatik und Management, 2017