Using STPA in an ISO 26262 Compliant Process

Archana Mallya[1], Vera Pantelic[1], Morayo Adedjouma[2], Mark Lawford[1], and Alan Wassyng[1]

[1] McMaster Centre for Software Certification, Department of Computing and Software, McMaster University, Hamilton, ON, Canada
{mallya,pantelv,lawford,wassyng}@mcmaster.ca
[2] CEA LIST, LISE, Gif-sur-Yvette Cedex, France
morayo.adedjouma@cea.fr

Abstract. *ISO 26262 is the de facto standard for automotive functional safety, and every automotive Original Equipment Manufacturer (OEM), as well as their major suppliers, are striving to ensure that their development processes are ISO 26262 compliant. ISO 26262 mandates both hazard analysis and risk assessment. Systems Theoretic Process Analysis (STPA) is a relatively new hazard analysis technique, that promises to overcome some limitations of traditional hazard analysis techniques. In this paper, we analyze how STPA can be used in an ISO 26262 compliant process. We also provide an excerpt of our application of STPA on an automotive subsystem as per the concept phase of ISO 26262.*

Keywords: Hazard analysis, risk assessment, STPA, ISO 26262, ASILs, automotive industry, Battery Management System

1 Introduction

Systems Theoretic Process Analysis (STPA) [6] is a relatively novel hazard analysis technique, geared towards modern, software-intensive complex systems in which analyzing the interacting subsystems as separate entities could give misleading results. Recently, STPA has gained popularity in the automotive domain (e.g., [2], [4], [8]). Also, ISO 26262 [5] has become the de facto standard for automotive functional safety. Given the industry's gradual shift to compliance with ISO 26262, the topic of STPA's application in an ISO 26262 compliant process is highly relevant to enabling greater acceptance of STPA in the industry. The key difference between *STPA* and the *Hazard Analysis and Risk Assessment (HARA)* process of ISO 26262 is the risk assessment. While STPA has proven itself as an effective hazard analysis technique across industries, it does not—nor was it intended to—include risk analysis.

In this paper, we carefully explore how to use STPA to satisfy the hazard analysis requirements of the concept phase of development as per ISO 26262, and how to augment STPA with an appropriate risk analysis. First, we provide a detailed comparison of the standard and the technique: we note the major similarities and differences in the philosophical underpinnings of the two, and

provide a comparison of their key terms. Then, we build on the comparisons to check if and how every relevant artefact as required/recommended by ISO 26262 can be generated/supported by applying STPA. Finally, we illustrate the application of the approach on a real-world automotive subsystem provided by our industrial partner, a large automotive OEM.

Although the topic of using STPA in an ISO 26262 compliant process has been the subject of study [4], or at least its significance has been recognized [3], [8], to the best of the authors' knowledge, this paper represents the first detailed account of the topic. The closest to an investigation of the topic is presented in [4], where the author suggests that HAZOP (Hazard and Operability), STPA and FMEA (Failure Mode and Effects Analysis) could be used in the concept phase of development as per ISO 26262. Although [4] provides a rough, high-level view on the topic, our work presents a detailed analysis of the topic with illustrative examples. Also, Hommes [3] suggests investigating the effectiveness of STPA in the automotive domain as the ISO 26262 recommended hazard analysis techniques are not sufficient to handle the growing complexity of modern software intensive safety-critical systems.

This paper is organized as follows. Section 2 provides relevant background. Section 3 provides a comparison of the terminologies of STPA and ISO 26262, while Section 4 presents guidelines on how to use STPA in an ISO 26262 compliant process, illustrated with an excerpt from an automotive subsystem. Section 5 concludes the paper and provides suggestions for future work.

This paper is based on the Master's thesis of the first author [7], and we refer the reader to it for further details.

2 Preliminaries

In this section, we provide a brief introduction to STPA and ISO 26262.

2.1 Systems Theoretic Process Analysis (STPA)

STPA is based on the accident causation model called STAMP (Systems-Theoretic Accident Model and Processes), built on systems theory and systems engineering [6]. The main ideas behind systems theory are: 1) Emergence and Hierarchy and, 2) Communication and Control [6]. Safety is considered an *emergent* system property: the safety of the whole system cannot be guaranteed just by proving that the system's individual components are safe. Further, systems are modeled as a *hierarchy* of organizational levels, where each level is more detailed than the one above. Also, accidents are treated as a dynamic *control problem* (as opposed to the classical approach viewing accidents as caused by component failures only): accidents occur when inadequate or inappropriate control actions (commands issued by system's controllers) violate the safety constraints of the system.

The STPA technique follows three steps: Preliminary step (Step 0), Step 1 and Step 2. Step 0 deals with the identification of accidents, associated hazards,

safety constraints as a negation of those hazards, and drawing of the control structure (a functional abstraction of the system). Step 1 identifies the ways in which unsafe control actions could lead to accidents, and the corresponding safety constraints. Step 2, causal factor analysis, involves identifying the causes of previously identified unsafe control actions along with the corresponding safety constraints. The detailed steps to perform STPA are described in Section 4.

STPA promises to address some limitations of traditional hazard analysis techniques as it accounts for the interactions between the subsystems and the dynamics between the system and its environment, along with management issues and human factors. There exists related work in various domains that presents cases where STPA identified hazards previously not identified by ISO 26262 recommended hazard analysis techniques (a detailed review of the literature can be found in [7]). For example, Song [10] applied STPA on the Nuclear Darlington Shutdown system and compared the results with the original FMEA results. The author found that, when compared with FMEA, STPA identified more hazards, failure modes and causal factors, including inadequate control algorithms, missing feedback and an incorrect logic model [10].

There have been varied opinions on the ease of use of STPA and the learning curve involved. There is no strong evidence to suggest that STPA is harder to use than traditional hazard analysis techniques. According to the controlled experiment presented in [1], there is no significant difference in the ease of use and understandability between STPA, FMEA and FTA (Fault Tree Analysis).

2.2 ISO 26262 Standard

ISO 26262, published in late 2011, addresses functional safety of road vehicles, and applies to electric, electronic and software components within the vehicle [5]. ISO 26262 consists of ten parts, and our work focuses on the Hazard Analysis and Risk Assessment (HARA) clause of Part 3 of [5].

The *Item Definition* is a necessary prerequisite for the HARA. The *item* is defined as "system", "or array of systems to implement a function at the vehicle level, to which ISO 26262 is applied", [Part 1 of [5]]. It contains the requirements for the item under study, its dependencies and its interactions with the environment and other items. The HARA comprises Situation Analysis, Hazard Identification, Classification of *Hazardous Events*, and Determination of *Automotive Safety Integrity Levels (ASILs)* and *Safety Goals*. The Situation Analysis determines "the operational situations and operating modes in which an item's malfunctioning behaviour will result in a hazardous event". The *Operational situation* is defined as a "scenario that can occur during a vehicle's life" (e.g., driving), while the *operating mode* is a "perceivable functional state of an item or element" (e.g., system off, degraded operation, emergency operation), [Part 1 of [5]]. The Hazard Identification step involves identifying the *vehicle level hazards*, the *hazardous events* and *consequences of hazardous events*. Hazardous events are determined by considering the hazards in different operational situations identified during the situation analysis. In the Classification of Hazardous Events step, the hazardous events are classified using impact factors *Severity (S)*,

Probability of Exposure (E) and *Controllability (C)*. The severity is estimated based on the extent of potential harm to each person potentially at risk. The parameter ranges from S0 to S3. The probability of exposure is the duration or the frequency of occurrence of the operational situations and is valued from E0 to E4. The controllability factor, ranged from C0 to C3, is an estimation of the ability of the driver or other persons potentially at risk to control the hazardous event. ASIL levels help determine the stringency of the requirements and the safety measures needed to avoid what the standard considers to be unreasonable risks. The Determination of ASILs for each hazardous event is based on the estimated values of the severity, probability of exposure and controllability parameters in accordance with Table 4 of Part 3 of [5], and range from ASIL A to ASIL D (highest criticality). Another class called Quality Management (QM) exists to denote there is no safety requirement to comply with. For each hazardous event with an assigned ASIL, a safety goal shall be determined as a top-level safety requirement for the item. The ASIL identified for a hazardous event shall also be assigned to the corresponding safety goal.

Then, the Functional Safety Concept (FSC) clause helps derive the *Functional Safety Requirements (FSRs)* from the item's safety goals based on preliminary architectural assumptions. The standard suggests using traditional safety analyses like FMEA, FTA, and HAZOP to support the FSR specification.

3 STPA and ISO 26262

In this section, we first compare the foundations of STPA and ISO 26262, and then compare their central terminologies. Table 1 presents definitions of central terms used by STPA as defined in [6] and ISO 26262, as defined in Part 1 of [5].

Table 1: Definitions of terms in STPA and ISO 26262

Term	STPA [6]	ISO 26262 [Part 1 of [5]]
Hazard	A system state or set of conditions that, together with a particular set of worst-case environmental conditions, will lead to an accident (loss)	Potential source of harm caused by malfunctioning behaviour of the item
Malfunctioning behaviour	No explicit definition	Failure or unintended behaviour of an item with respect to its design intent
Failure	No explicit definition Note: A failure in engineering can be defined as the non-performance or inability of a component (or a system) to perform its intended function	Termination of the ability of an element, to perform a function as required Note: Incorrect specification is a source of failure
Accident	An undesired or unplanned event that results in a loss, including loss of human life or human injury, property damage, environmental pollution, mission loss, etc.	No explicit definition
Harm	No explicit definition	Physical injury or damage to the health of persons
Hazardous event	No explicit definition	Combination of a hazard and an operational situation

3.1 STPA and ISO 26262: Comparing Foundations

Both ISO 26262 and STPA are based on a systems engineering framework in which a system is considered to be more than merely the sum of its parts. Top-down analysis and development are common to both. However, while ISO 26262 emphasizes the importance of considering the context of a system in achieving safety (including the role of safety management and safety culture), there seems to be no consensus whether ISO 26262 considers a driver to be a part of the hazard analysis of an item. STPA on the other hand, includes all relevant aspects of the system's environment, including the driver.

The key difference between STPA and the HARA process of ISO 26262 is the risk assessment process. Risk assessment as per ISO 26262 involves determining the impact factors: Severity (S), Probability of Exposure (E) and Controllability (C). While ISO 26262 justifiably avoids using probabilities of failure of system components to estimate risk (estimating the probability of failures in modern systems is very hard given the prevalence of non-random failures and the lack of historical information), determining E and C factors is still rather subjective and not yet standardized in the industry. Although SAE J2980 presents a recommended practice to "provide guidance for identifying and classifying hazardous events, which are as per [5]", its current focus is limited to collision related hazards and not the wider scope of ISO 26262 [9]. Although some authors suggest using only severity for risk estimation [3], an ISO 26262 compliant process requires that S, E and C factors are determined and used.

3.2 STPA and ISO 26262: Comparing Basic Terminologies

Important terms are italicized in this section and defined in Table 1.

Starting from the definition of term *hazard*, the most notable difference between the definitions of the term is that STPA does not limit *hazards* to those caused by *malfunctioning behaviour*, while ISO 26262 does. Due to ISO 26262's ambiguity in the definition of the term of *malfunctioning behaviour* (both "design intent" and "unintended behaviour" are undefined in the standard), it is hard to determine what exactly the term is intended to mean. However, it seems that ISO 26262, by including "unintended behaviour of an item with respect to its design intent" in the definition of *malfunctioning behaviour*, departs from the notion that only component failures lead to accidents, but also includes unintended interactions of the system components often reflected in flawed requirements.

ISO 26262 does not have an explicit definition of *accident*. However, STPA's *accident* defines unacceptable losses; hence, it is very closely related to ISO 26262's *harm*. We note that STPA's concept of loss not only includes human injury or loss of human life, but also property damage, pollution, mission loss, etc., and hence is more general than ISO 26262's which considers only injury to people. We note, however, that STPA is meant for different stakeholders to adapt it as it suits them. Further, the consequences of hazardous events of ISO 26262, identified by considering the consequences of the *hazard* in various operational situations, could be mapped to *accidents* of STPA. However, STPA's *accidents* are more

general in nature—ISO 26262's consequences tend to be more fine-grained, since they are determined for different operational situations of the *hazard*.

4 Using STPA in an ISO 26262 Compliant Process

The foundation and terminology comparison presented in Sect. 3 lays the groundwork for our approach of using STPA in an ISO 26262 compliant process. In this section, we explore how every relevant artefact as required/recommended by ISO 26262 can be generated/supported by applying STPA. We illustrate the approach on a simplified version of the Battery Management System (BMS) of a Plug-in Hybrid Electric Vehicle (PHEV). Not all the details of the analysis and the control structure are shown due to their proprietary nature.

A PHEV is a hybrid electric vehicle (a vehicle combining Internal Combustion Engine (ICE) propulsion with electric propulsion) that has energy storage devices like rechargeable batteries that can be charged by connecting to the electrical grid using a plug. These rechargeable batteries are monitored and protected by the BMS. The primary functions of the example BMS are to: 1) enable charging and discharging of the battery back by closing the contactors; or disable charging and discharging of the battery back by opening the contactors, 2) provide accurate information on charge and discharge to the HPC (Hybrid Powertrain Controller) system, 3) equalize cell charge using passive cell balancing, 4) heat/cool the battery pack, 5) isolate the battery in case of emergency.

Figure 1 shows the STPA results (as shown on the right hand side of the figure) that can help generate the outputs required by the concept phase of ISO 26262 (as shown on the left hand side of the figure), illustrated with examples. The grey dashed box includes blocks numbered i2 to i10 corresponding to all the subclauses of the HARA. Rectangles denote outputs (artefacts) obtained as a result of either following the requirements of ISO 26262 or performing STPA steps. An oval inside a rectangle represents the output of HARA that is itself required to determine the output represented by the encompassing rectangle. The solid arrows from STPA blocks to ISO 26262 blocks denote that the specific STPA output can completely support the corresponding ISO 26262 block output, while the dashed arrows denote that the specific STPA output can partially support the corresponding ISO 26262 block output. The dotted arrows are used to represent cases where a result of STPA can help provide additional support in generating an output of ISO 26262. The numbers in the form of *w*-*x*-*y*-*z* point to the specific subclause of the standard where the requirements are specified. *w* corresponds to the specific part of the ISO 26262 standard, *x* corresponds to the specific clause and y-z corresponds to the subclauses where the requirements of the clauses and subclauses are mentioned.

As shown in Fig. 1, there exist outputs of the original STPA that have been mapped to ISO 26262's outputs represented by blocks i1, i2, i3, i5, i10 and i11. However, ISO 26262's outputs represented by blocks i4, i6, i7, i8 and i9 have no corresponding outputs of the original STPA—instead, we augment STPA's outputs with a set of new outputs (shaded rectangles in the figure) and present

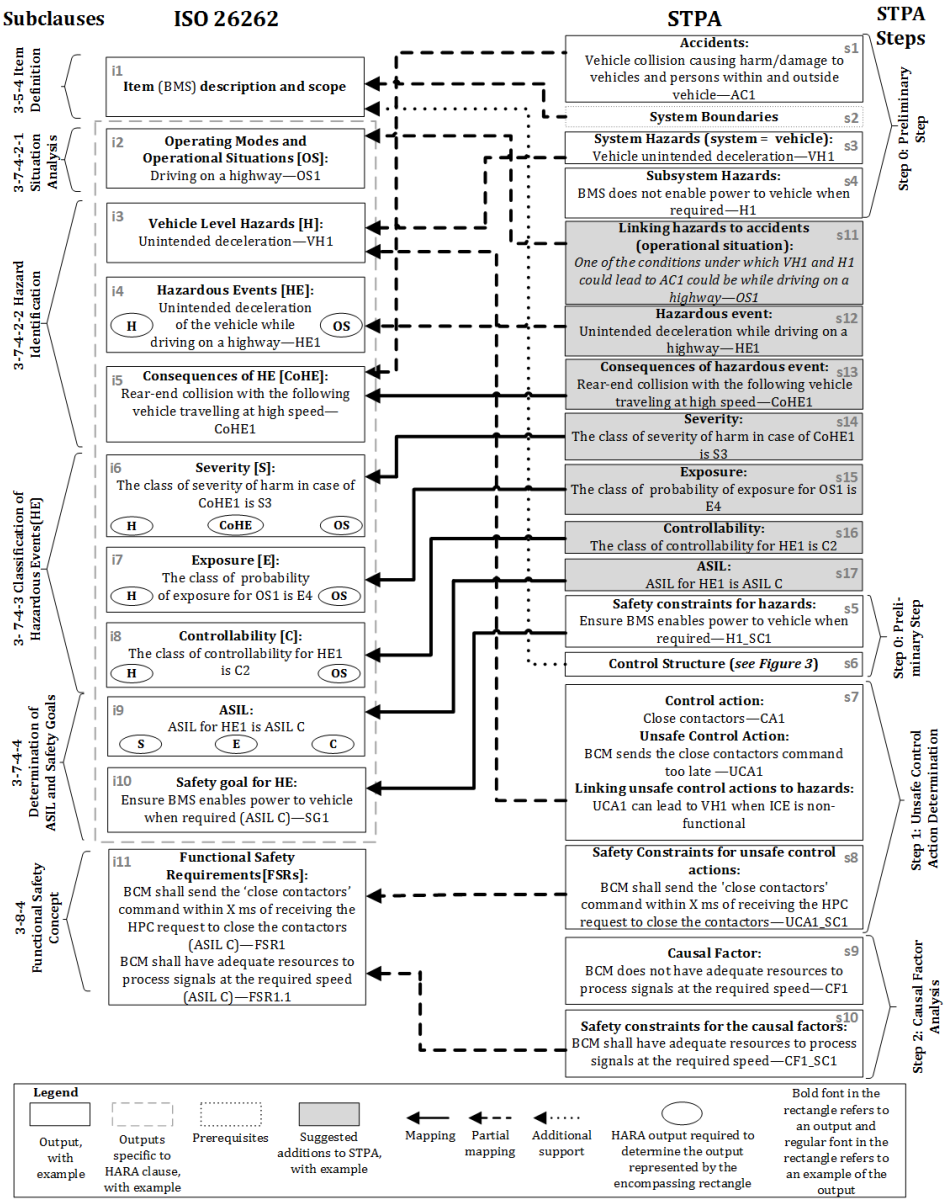

Fig. 1: STPA in compliance with ISO 26262, illustrated with excerpt of BMS example

guidelines on how they can be generated. The mappings from Fig. 1 are discussed next in the sequence in which STPA is performed.

First, in the initial step of STPA (Step 0), accidents, related hazards, corresponding safety constraints, and the control structure of the system under anal-

ysis are defined. *Accidents* of STPA are mapped to the standard's *consequences of hazardous events* (s1→i5 mapping). However, *accidents* are typically not as fine-grained as *consequences of hazardous events*, since STPA does not consider *accidents* in different *operational situations*. Thus, STPA can *partially* support derivation of the *consequences of hazardous events*. However, *consequences of hazardous events* are only to be determined later in the STPA process compliant with ISO 26262 after *hazardous events* have been determined. An example of a vehicle accident would be *AC1: Vehicle collision causing harm/damage to vehicles and persons within and outside vehicle*.

Next in STPA, vehicle level hazards are determined. If the hazards cannot be eliminated or controlled at the system level, the corresponding component hazards are identified. An example of a system (vehicle) level hazard for *AC1* is *VH1: Vehicle experiences unintended deceleration*, which maps to the vehicle level hazard (block i3) of ISO 26262. A corresponding BMS-level hazard is *H1: BMS does not enable power to vehicle when required*. *H1* is a refinement of *VH1* as *H1* is hazardous in the case when ICE is non-functional so the vehicle is completely dependent on the power from the battery pack.

In an ISO 26262-compliant process, *operational situations* should be documented. This is why, in our approach, we explicitly document *operational situations* during the process of linking hazards to accidents (block s11). When analyzing in what situations the hazards *VH1* and *H1* could lead to *AC1*, we can come up with situations like *driving on a highway, snow, ice on road, etc.* To help analysts in determining operational situations, the standard presents examples of operational situations (the examples have been summarized in [7]). As mentioned in ISO 26262, an overly detailed list of *operational situations* might result in "a very granular classification of *hazardous events*" and could eventually lead to "an inappropriate lowering" of an *ASIL* [Part 3 of [5]].

Once we have the list of *operational situations* and *hazards*, we can derive the *hazardous event* (block s12), as it is the combination of a *hazard* and an *operational situation*. Thus, we augment STPA to include this step (block s12). For *VH1*, a hazardous event identified is *HE1: Unintended deceleration while driving on highway*. It is at this point when the *consequence of hazardous event* is determined (block s13). For example, the consequence of *HE1* would be *CoHE1: Rear-end collision with the following vehicle travelling at high speed*.

Classification of hazardous events and *determination of ASILs* are the main subclauses of the concept phase of ISO 26262 that do not have a corresponding step in STPA. As per ISO 26262, the severity of potential harm is estimated using injury scales like the Abbreviated Injury Scale (AIS) and Maximum AIS in accordance with the Table 1 of Part 3 of [5]. According to ISO 26262, the *severity* is determined for a hazardous event, based on the *consequence of the hazardous event*. Using the results obtained from blocks i2, i3 and i5, i.e., *operational situations, hazards and consequences of hazardous events*, we can estimate the *S* impact factor (block i6) as per ISO 26262. For *HE1*, the class of severity of harm is *S3* (Life-threatening injuries). Further, the class of probability of exposure for driving on a highway is *E4* (greater than 10% of average operating time). Then,

the class of controllability for unintended deceleration while driving on a highway is *C2*, i.e. 90% or more of all drivers or other traffic participants are usually able to avoid harm. The impact factors *S3*, *E4*, and *C2* were estimated based on Tables 1, 2, and 3 of Part 3 of [5], respectively. *ASIL*s are then determined using the values of *S, E and C* from previous blocks and using Table 4 of Part 3 of [5]. The ASIL for the above chosen impact factors is *ASIL C*.

For each of the hazardous events with an assigned ASIL, a *safety goal* is determined. *Safety constraints* (block s5) derived as the negation of hazards in STPA's Step 0 can be used to support the *safety goal* determination subclause of ISO 26262 (as shown by the mapping s5→i10), as the safety goals are high-level requirements. An example of a safety constraint for *H1* is *H1_SC1: Ensure BMS enables power to the vehicle when required*, corresponding to ISO 26262's *safety goal*, denoted *SG1*. *SG1* is then assigned *ASIL C*, as determined for *HE1*.

The next task in STPA is the development of *control structure* (block s6) as a graphical representation of the functional model of the system [6]. Building from the vehicle view, a control structure for the entire vehicle (including driver and environment) is first built so that hazards due to interactions between the components can be identified (Fig. 2). Thus, an STPA control structure gives a holistic view of the entire system under study. We then zoom in on the BMS itself, with the control structure as given in Fig. 3. The *HPC (Hybrid Powertrain Controller)*, the contactors, the battery pack and the 12V battery are the external systems that interact with the BMS. Other systems in the BMS environment are the fan/pump components for the thermal management system, the on-board charger, and the external charger. The components of the BMS are shown inside the shaded dashed box in Fig. 3, namely, the *BCM (Battery Control Module)*, *BMM (Battery Monitoring Module)*, and the history log and cells specification module. The components and arrows as shown in Fig. 3 were identified based on the general functionalities of a BMS elicited through literature review and with the help of domain experts. The *control structure*, as a control-oriented diagram depicting the functionalities of an item, is additional information to help an analyst in accomplishing part of the item definition output. Thus it represents a valuable diagram that complements the existing item definition (block i1)—hence the dotted s6→i1 mapping in Fig. 1.

In STPA Step 1, the control actions from the control structure are categorized into four categories, i.e., four ways in which a control action can be unsafe (see Table 2). The control action selected for Step 1 of this example application is *CA1: Close contactors*. This control action is sent by the BMS after it receives the authorization or a request to close the contactors from the HPC. When the contactors are closed in the driving mode, the battery pack can receive power from HPC from regenerative braking and the HPC can receive power from the battery pack. Let us consider the control action *CA1: Close contactors* under the category 'Safe control action is provided too late, too early, wrong order'. An example of unsafe control action would be *UCA1: BCM sends the close contactors command too late*. When analyzing the ways in which a control action can be unsafe and linking them to the hazards, the analyst has to identify

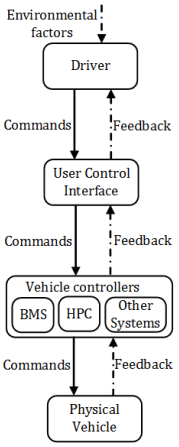

Fig. 2: High level control structure

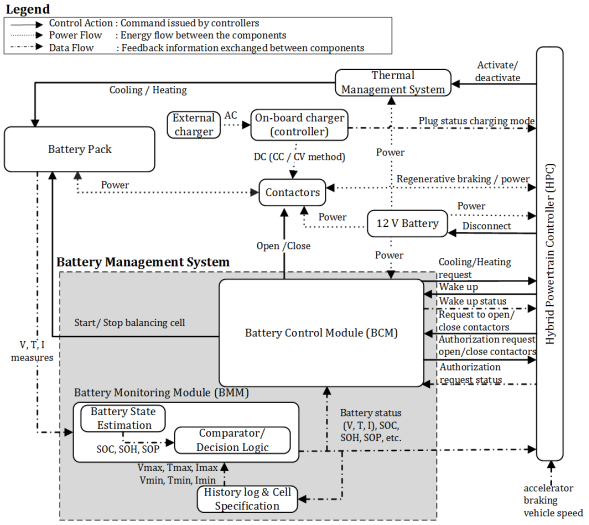

Fig. 3: Detailed BMS control structure [7]

the context which makes the control action hazardous. In this case, the *UCA1* can lead to *VH1: Vehicle experiences unintended deceleration*, when the *ICE is non-functional*. The assumption is that the HPC has already requested the BMS to close the contactors and that the HPC gives the command only when safe to do so. When linking UCAs to the hazards of Step 0, one can sometimes identify hazards that were not previously identified. Hence, Step 1 can be linked to the hazard identification step of ISO 26262 as well (see mapping s7 → i3). Step 1 also involves translating the UCAs into safety constraints and further refining the safety constraints from Step 0 (block s8). An example of a safety constraint for *UCA1* is *UCA1_SC1: BCM shall send the 'Close contactors' command within X ms of receiving the HPC request to close the contactors*. Since this safety constraint of Step 1 describes what needs to be done to achieve the *safety goal*, it represents ISO 26262's *functional safety requirement* (denoted by s8 → i11).

Table 2: Excerpt of results of STPA Step 1 for CA1: close contactors

Control action	Required control action not provided	Unsafe control action provided	Safe control action provided too late, too early, wrong order	Continuous safe control action provided too long or stopped too soon
CA1:Close contactors	...	...	UCA1:BCM sends the close contactors command too late [VH1]	...

Causal factor analysis (Step 2) of STPA involves examining the control loop of control actions and identifying the causes of unsafe controls (block s9 from Fig. 1). The control loop includes the controller that initiates the control action, the actuator, the sensor, and the controlled process [6]. A unique control loop

is identified and used for all identified Unsafe Control Actions (UCAs) of the selected control action. Then, a causal factors analysis diagram is defined for the UCAs based on the guide words provided by STPA [6]. Part of our causal factor analysis for *UCA1* is shown in Fig. 4. Specific causes of UCAs that may lead to hazards are shown in *italics*.

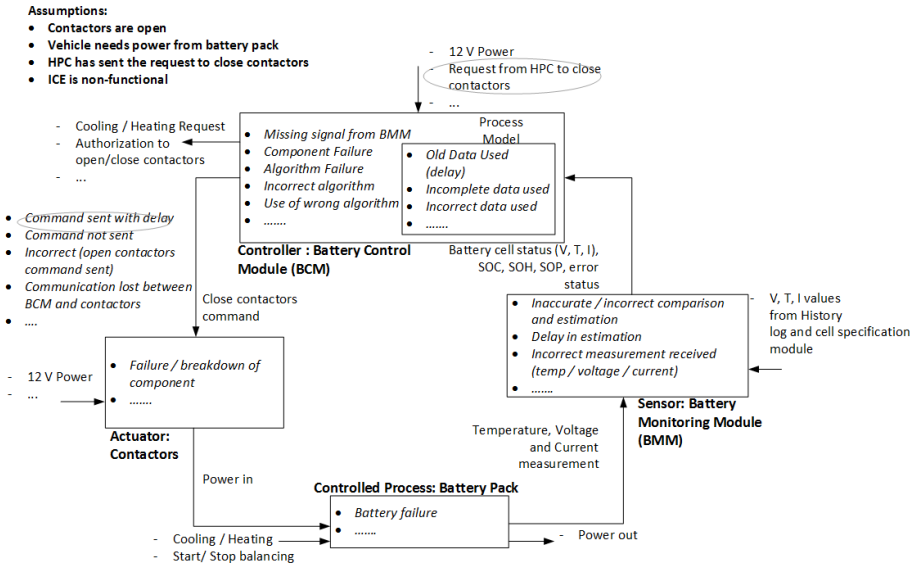

Fig. 4: Causal factor analysis for control action, CA1: Close contactors

For the loop in Fig. 4, the BCM, as a controller, should issue the control action *Close Contactors* to the actuators that will realize the command. The controlled process in the loop is the battery pack. During the causal factors analysis we assume that the HPC has already sent the request to close the contactors. For the sake of simplicity, in Fig. 4, we have only shown a few of the causal factors of the contactors, the battery pack, the BMM, the BCM including its process model and the ones between the BCM and the contactors. Other causal factors (e.g., 12 V power disconnected) are not shown here. The causal factors identified in STPA Step 2 can help fulfill one of the objectives of the safety analyses as per ISO 26262 (Clause 8 of Part 9 of ISO 26262), i.e., to identify the "causes that could lead to the violation of a safety goal or safety requirement". Once the causes are identified, the analyst needs to identify the safety constraints to mitigate or eliminate those causes. One of the causes which could result in *UCA1* is *BCM does not have adequate resources to process signals at the required speed.* An example of a Step 2 safety constraint identified for *CF1* is *CF1_SC1: BCM shall have adequate resources to process signals at the required speed.* This Step 2 safety constraint represents a *functional safety requirement* of ISO 26262—

hence the mapping s10→i11. While the level of details needed when defining the safety constraints during STPA is not pre-determined, ISO 26262 specifies the characteristics and parameters that a safety requirement should include, e.g., the fault tolerant time interval if available, the safe state, the operating mode, etc.

5 Conclusion and Future Work

While STPA represents a promising hazard analysis technique that addresses some limitations of traditional techniques, it does not attempt to provide the risk analysis component sometimes included in traditional hazard analysis techniques. By careful investigation of the requirements of the HARA clause of Part 3 of ISO 26262, we conclude that STPA does not interfere with the ISO 26262's risk analysis in any way—instead, STPA was shown to only require modest augmentation in order to be used in a HARA process compliant with ISO 26262. Thus, the augmented STPA presented here can support all the outputs of ISO 26262 generated as a result of satisfying the standard's HARA requirements. Consequently, we can utilize STPA's advantages in an ISO 26262-compliant process.

There are now a number of examples in the literature on how to use STPA in an automotive context. However, what seems to be lacking is principles for performing STPA. For example, finding an appropriate abstraction level for the control loop seems to be extremely important. Future work on documenting such principles and their rationale would be extremely useful.

References

1. Abdulkhaleq, A., Wagner, S.: A controlled experiment for the empirical evaluation of safety analysis techniques for safety-critical software. pp. 16:1–16:10. EASE'15, ACM (2015)
2. D'Ambrosio, J., Debouk, R., Hartfelder, D., Sundaram, P., Vernacchia, M., Wagner, S., Thomas, J., Placke, S.: Application of STPA to an automotive shift-by-wire system. In: STAMP Workshop, Cambridge, MA (2014)
3. Hommes, Q.V.E.: Review and assessment of the ISO 26262 Draft Road Vehicle - Functional Safety. SAE technical paper (2012)
4. Hommes, Q.V.E.: Safety analysis approaches for automotive electronic control systems (2015), `http://www.nhtsa.gov/DOT/NHTSA/NVS/Public%20Meetings/SAE/2015/2015SAE-Hommes-SafetyAnalysisApproaches.pdf`
5. ISO 26262: Road vehicles – Functional safety, International Organization for Standardization (ISO) (2011)
6. Leveson, N.G.: Engineering a Safer World: Systems Thinking Applied to Safety (Engineering Systems). The MIT Press, Cambridge, Massachusetts (2012)
7. Mallya, A.: Using STPA in an ISO 26262 compliant process. M.A.Sc., McMaster University, Canada (Oct 2015)
8. NHTSA: Request for comment on automotive electronic control systems safety and security. `https://federalregister.gov/a/2014-23805` (2014)
9. SAE J2980: Considerations for ISO 26262 ASIL Hazard Classification, SAE International (2015)
10. Song, Y.: Applying System-Theoretic Accident Model and Processes (STAMP) to Hazard Analysis. M.A.Sc., McMaster University, Canada (2012)