

# Run-Time Risk Mitigation in Automated Vehicles: A Model for Studying Preparatory Steps

Mario Gleirscher

Faculty of Informatics  
Technical University of Munich  
Munich, Germany  
mario.gleirscher@tum.de

We assume that autonomous or highly automated driving (AD) will be accompanied by tough *assurance obligations* exceeding the requirements of even recent revisions of ISO 26262 or SOTIF. Hence, automotive control and safety engineers have to (i) comprehensively analyze the driving process and its control loop, (ii) identify relevant hazards stemming from this loop, (iii) establish feasible automated measures for the effective mitigation of these hazards or the alleviation of their consequences.

By studying an example, this article investigates some achievements in the modeling for the steps (i), (ii), and (iii), amenable to formal verification of desired properties derived from potential assurance obligations such as the *global existence of an effective mitigation strategy*. In addition, the proposed approach is meant for step-wise refinement towards the automated synthesis of AD safety controllers implementing such properties.

## 1 Introduction

For many people, driving a car is a difficult task even after guided training and many years of driving practice: This statement gets more tangible when driving in dense urban traffic, complex road settings, road construction zones, unknown areas, with hard-to-predict traffic co-participants (i.e., cars, trucks, cyclists, pedestrians), bothersome congestion, or driving a defective car. Consequently, hazards such as drivers misjudging situations and making poor decisions have had a long tradition. Hence, road vehicles have been equipped with many sorts of safety mechanisms, most recently, with functions for “safety decision support” [7], driving assistance, and monitor-actuator designs, all aiming at reducing operational risk or making a driver’s role less critical. In AD, more highly automated mechanisms will have to contribute to risk mitigation at run-time and, thus, constitute even stronger assurance obligations [7].

In Sec. 1.1, we introduce basic terms for *risk analysis and (run-time) mitigation (RAM)* for automated vehicles (AV) as discussed in this work. Sec. 1.2 elaborates some of these assurance obligations.

### 1.1 Background and Terminology

According to control theory, a *control loop*  $\mathcal{L}$  comprises a (physical) *process*  $P$  to be controlled and a *controller*  $C$ , i.e., a system in charge of controlling this process according to some laws defined by an application and an *operator* [20]. The engineering of safety-critical control loops typically involves reducing *hazards* by making controllers safe in their intended function (SOTIF), resilient (i.e., tolerate disturbances), dependable (i.e., tolerate faults), and secure (i.e., tolerate misuse).

In automated driving, the process under consideration is the *driving process* which we decompose into a set of *driving situations*  $\mathcal{S}$ , see Sec. 2. Taxonomies of such situations have been published in, e.g.

[32]. Based on recommendations of SAE and U.S. DoT [27, 34], we distinguish the following *modes* or *levels* (L) of automation for AD: AD assistance (ADAS, L1-2), highly automated (HAD, L3-4), and fully automated or autonomous (FAD, L5). Yet, in level 4, a human driver is supposed to stay “weakly” in-the-loop for an occasional machine-to-human hand over.

By the term *causal factor* (CF), we denote any concept ranging from a *root cause* to a *hazard*, to a *near-mishap*.<sup>1</sup> Causal factors can form causal chains describing effect propagation through  $\mathcal{L}$ , particularly, a near-mishap can lead to another root cause.<sup>2</sup> By *causal factor model* (also: hazard or risk model), we refer to the collection of causal factor descriptions used for a RAM application.

A *rare undesired event* (RUE, also: hazardous event or situation) denotes any *state of*  $\mathcal{L}$  (i) with an occurrence of one or more causal factors and (ii) arguably increasing the risk of mishaps. Moreover, by (*operational*) *risk* [19], we refer to the likelihood—quantitatively, the probability or frequency—of reaching a specific set of RUEs from any current state. Then, a *safe state* is any state with acceptably low risk—quantitatively, risk below a *budget*  $b$ . Although we use further information for the modeling in RAM, any safe state can be transcribed into a corresponding behavioral invariant or a safety property of  $\mathcal{L}$  in the sense of [17, 23].

Elaborating on the notion of *behavioral safety* in [9], we consider two types of actions in  $\mathcal{L}$ : *endangerments* are actions whose performance leads to RUEs, and *mitigations* are actions whose performance represents countermeasures for specific sets of RUEs. For any causal factor  $cf$ , we consider two types of endangerments:  $e^{cf}$  denotes activation resulting in a change to phase  $cf$  (i.e., *active*), and  $e_m^{cf}$  describes any *mishap*  $cf$  potentially following  $cf$ . Furthermore, we consider three types of mitigations:  $m_s^{cf}$  initiates mitigation by a change to phase  $\overline{cf}$  (i.e., *mitigated*),  $m_c^{cf}$  completes mitigation by a change to phase  $0^{cf}$  (i.e., *inactive*), and  $m_e^{cf}$  directly deactivates  $cf$  and completely restores its consequences, again by a change to  $0^{cf}$ . Shown in Fig. 1, these notions lead to what we call *phase model*<sup>3</sup> [11]. Phase models can be composed. Then, by 0 we refer to the “safest” state in a model composed of phase models for several causal factors, see Sec. 3.

For the description of these concepts, we employ *labeled transition system* (LTS) modeling. An LTS is a tuple  $(\Sigma, \mathcal{A}, \Delta)$  with a set of *states*  $\Sigma$ , a set of *action* labels  $\mathcal{A}$ , and a transition relation  $\Delta \subseteq \Sigma \times \mathcal{A} \times \Sigma$ . For a *transition*  $(\sigma, \alpha, \sigma') \in \Delta$ , we also say that the state  $\sigma'$  is the *event* observed from the performance of the action  $\alpha$ . We will work with the (usually dense) *loop state space*  $\Sigma_{\mathcal{L}}$ , the finite *situation state space*  $\Sigma_{\mathcal{S}}$ , and the finite *risk state space*  $\Sigma$ . The symbols used to model CFs in  $\Sigma$  (Fig. 1) can refer to predicates over  $\Sigma_{\mathcal{L}}$ . Further details on this terminology will follow in Secs. 2 and 3.

## 1.2 Motivation and Challenges

Let us take the viewpoint of a safety engineer dealing with the assurance of AVs not only compliant with available standards (e.g. ISO 26262) but also trying to achieve the level of safety desired by any traffic participant. What does that mean? Well, our engineer might try to argue towards the following claim from manual driving:

**Claim 1** *The rational driver always tries to reach and maintain a safe state w.r.t. all RUEs recognizable and reacted upon by the driver in any driving situation.*

<sup>1</sup>For example, a road accident for which airbags successfully alleviated harm.

<sup>2</sup>This helps fixing the “Swiss cheese problem” raised by LEVESON [18].

<sup>3</sup>Gray transitions indicate that we might employ more expressive variants of this phase model as, e.g. discussed in [16] for a different context. Operationally, the phase model would perform non-observable (also: silent) transitions if none of the described actions is enabled. Similar abstractions are used for testing, e.g. in [33].

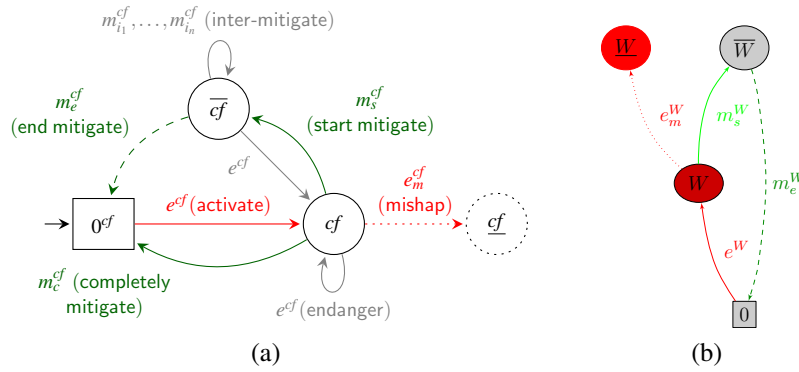


Figure 1: Generic phase model for causal factors (a) exemplified for the causal factor  $W$  for badWeather (b).

The reader might not agree with this claim or require further claims to be established in this context. However, for sake of simplicity, the following discussion is restricted to this claim. Next, we list three **tasks** of our safety engineer responsible for RAM for AD:

**Task S:** investigate the domain of *situations* in the driving process,

**Task E:** investigate the domain of *endangerments*, identify causal factors and potential mishaps,

**Task M:** investigate the domain of *mitigations*, prepare effective countermeasures.

These tasks are error prone, largely subject to expert judgment and require rigorous method because

- driving is a complex task (see beginning of Sec. 1)
- we aim at FAD,<sup>4</sup>
- we might need to automate mitigations as well,<sup>5</sup>
- we have to *argue* that the measures developed (**M**) can mitigate the RUEs identified (**E**) in the considered driving process (**S**) and residual risk stays acceptably low, e.g. below an upper bound  $b$ . We investigate these tasks using the notions from Sec. 1.1, discuss some challenges as well as a RAM approach based on this model:

For **Task S**, we have to come up with a model of driving situations reflecting the driving process in the road environment (see Sec. 2), abstract but comprehensive enough for RAM purposes.

**Running Example Pt. 1** *Throughout the paper, we will work with an example highlighting aspects of RAM for AD. From our knowledge about the driving domain, we can quickly derive a large number of situations such as, e.g. leaving a parking lot, entering a highway, halting somewhere, driving through road construction zones, overtaking in urban traffic, and, more generally, AD at level 4, called drivingAtL4Generic below.*

For the automation of some analyses, we use a newly developed tool called YAP based on concepts discussed in [9, 11]. YAP stands for “yet another planner.” The complete example for this paper, a preliminary version of YAP, and a user’s manual are available online.<sup>6</sup>

*Questions from Task S:* (S1) Which situations form equivalence classes? (S2) Which situations do we have to discriminate in our model? (S3) How are the discriminated situations related?

For **Task E**, given a set of situations from Task S, we have to understand how a *safe state* in each situation looks like as well as all the undesirable ways of leaving this state,<sup>7</sup> i.e., endangerments leading to some RUE. Note that RUEs can model simultaneously occurring causal factors.

<sup>4</sup>See, e.g. <http://www.cnbc.com/2017/06/07/robot-cars-cant-count-on-us-in-an-emergency.html>.

<sup>5</sup>See, e.g. <https://www.theguardian.com/technology/2017/apr/04/uber-google-waymo-self-driving-cars>.

<sup>6</sup>Demonstration artifacts can be downloaded from <http://gleirscher.de>.

<sup>7</sup>For behavioral safety we have to deal with leaving safety invariants in many ways. Our interest, hence, lies on accepting

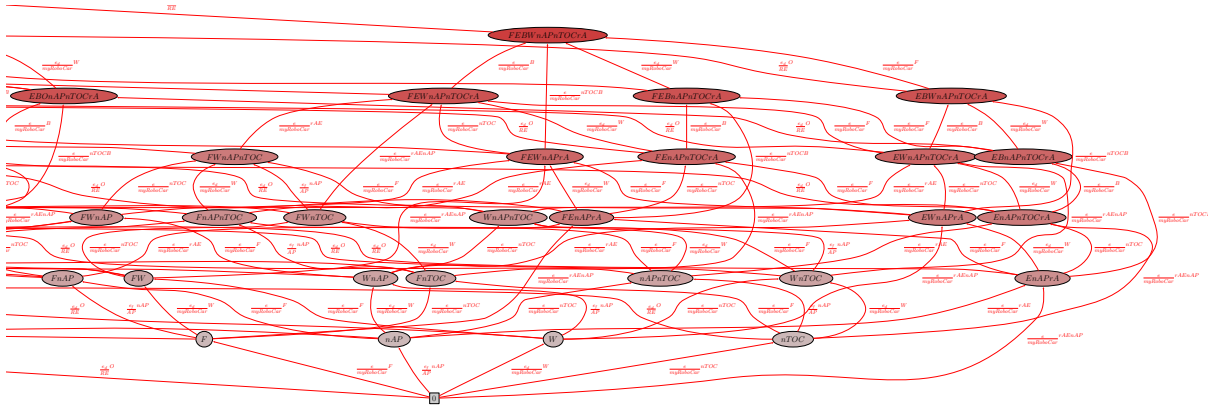


Figure 2: Cutout of the risk state space  $\Sigma$  indicating endangerment complexity: *nodes* indicate RUEs to be regarded in a situation; *edges* indicate potential endangerments leading to these RUEs.

**Running Example Pt. 2** *Safety engineers of the driving domain would derive many causal factors such as, e.g. sensor fault, inattentive driver, electric power shortage, low fuel. For the situation driveAtL4Generic, Fig. 2 depicts risk complexity in terms of RUEs hypothetically reachable and predictable from the state 0 for which we might be required to provide mitigations in any AD mode. Combining 10 causal factors (incl. the mentioned ones) using the phase model from Sec. 1.1, with initial state 0, we get 7128 states reachable via 51984 transitions. By only regarding ways of how we directly get into a RUE, we are left to consider 112 states and 372 transitions. Further details see Secs. 3 and 4.*

*Questions from Task E:* (E1) Which causal factors form equivalence classes? (E2) Which ones do we have to include in our model? (E3) How can we further classify these factors? (E4) Which factors have to be focused? (E5) How are the remaining factors related to each other?

For **Task M**, given the causal factors from Task E, we have to select and prepare optimal ways of how each of the RUEs can be mitigated.

*Questions from Task M:* (M1) Which mitigation is required for which RUE in which situation? (M2) How can we represent this in our model? (M3) Does it help to classify RUEs according to the available mitigations? (M4) Can we construct mitigations applicable to classes of RUEs to get a minimal set of effective mitigations for a given set of situations?

Aligned with the challenges summarized in [12], the questions S1-3, E1-5, and M1-4 define our general research setting.

**Contributions and Outline.** We show how we get to a model for RAM for AD in a systematic way, particularly, keeping a link to the detailed models used by control engineers for the controlled process. We discuss • the modeling of the driving process as a controlled process over driving situations (Task S, Sec. 2), • a refinement of the transition system model discussed in [11], • the classification and modeling of multiple causal factors comprising RUEs in the AD control loop specific to a set of driving situations (Secs. 2.2 and 3), • the state space exploration based on this model together with constraints for reducing and shaping this state space (Task E, Secs. 3.4 and 4.2), • a taxonomy of mitigation actions (Sec. 4.1),

to leave an “inner invariant,” i.e., allowing causal factors to occur, and trying not to leave an “outer invariant,” i.e., reducing the likelihood of an accident. This way, we employ LAMPORT’s [17] notion of safety property in a layered manner, similar to the discussion of “safety envelopes” [14] or layer-of-protection analysis (LOPA) [1].

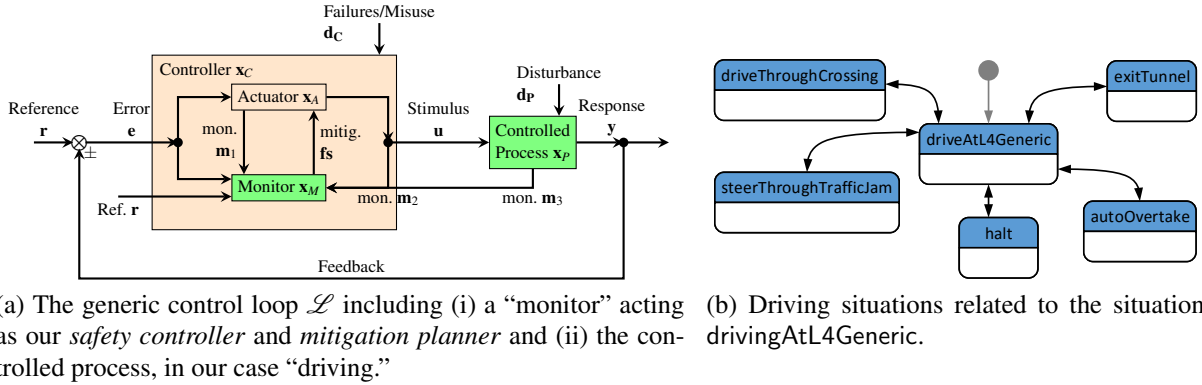


Figure 3: Generic structure of the control loop (a) and cutout of the abstraction of the controlled process (b).

and  $\bullet$  the usage of this model for RAM (Task M, Sec. 4). We discuss related work in Sec. 5, further issues and a research vision in Sec. 6, and conclude in Sec. 6.5.

## 2 Modeling Driving Situations, Processes, and Scenarios

Based on Sec. 1.1, we provide an abstraction of the driving process in  $\mathcal{L}$ , and a notion of driving scenario for the symbolic execution of this process as shown in Sec. 4.3.

### 2.1 Preliminaries on Driving Processes

**State of the Control Loop.** Given the process output  $\mathbf{y}$ , reference  $\mathbf{r}$ , and control input  $\mathbf{u}$ , the *process*  $P$  (open loop) can be described by  $\mathbf{y}_{new} = P(\mathbf{y}, \mathbf{u})$ . We close this loop using  $\mathbf{u} = C(\mathbf{e})$  for the controller  $C$ ,  $\mathbf{e} = \mathbf{r} - \mathbf{y}_{new}$  for the error  $e$ . At this level, we have the state  $(\mathbf{r}, \mathbf{u}, \mathbf{y}) \in \Sigma_{\mathcal{L}} \triangleq \mathbb{R}^{dim(\mathbf{r})+dim(\mathbf{u})+dim(\mathbf{y})}$  of the control loop  $\mathcal{L}$  [20]. However, as shown in Fig. 3a, based on the controller-internal state  $\mathbf{x}_C$ , controller-related causal factors  $\mathbf{d}_C$ , process-internal state  $\mathbf{x}_P$ , and process-related causal factors  $\mathbf{d}_P$ , we get  $\mathbf{u} = C(\mathbf{x}_C, \mathbf{e}, \mathbf{d}_C)$  and state  $(\mathbf{r}, \mathbf{x}_C, \mathbf{d}_C, \mathbf{u}, \mathbf{x}_P, \mathbf{d}_P)$  which we can refine [4] to get  $\mathbf{u} = A(\mathbf{x}_A, \mathbf{e}, \mathbf{f}_S, \mathbf{d}_C)$  for the *actuator*  $A$  with  $\mathbf{f}_S = M(\mathbf{x}_M, \mathbf{m}_{1,2,3}, \mathbf{r}, \mathbf{e})$  for the *monitor*  $M$ . For the process, we then consider  $\mathbf{y}_{new} = P(\mathbf{x}_P, \mathbf{u}, \mathbf{d}_P)$  with the *loop state*  $\sigma_{\mathcal{L}} = (\mathbf{r}, \mathbf{x}'_C, \mathbf{x}_A, \mathbf{x}_M, \mathbf{d}_C, \mathbf{u}, \mathbf{x}_P, \mathbf{d}_P)$  where  $\mathbf{x}_P$  includes  $\mathbf{y}$  and  $\mathbf{x}_C = (\mathbf{x}'_C, \mathbf{x}_A, \mathbf{x}_M)$ .

**Driving Processes.** Inspired by process algebra (see, e.g. [24, 13]), for a set of *driving situations*  $\mathcal{S}$ , we represent *driving processes*  $S$  or  $T$  by expressions of the form

$$S, T ::= s \mid S \parallel T \mid S | T \mid S; T \mid S^* \quad (\text{ordered by operator precedence}) \quad (1)$$

where  $s \in \mathcal{S}$ ,  $S; T$  denotes the *sequential composition* of  $S$  and  $T$ ,  $S | T$  *non-deterministic choice* between  $S$  and  $T$ ,  $S \parallel T$  their *parallel composition*, and, for convenience,  $S^*$  the (possibly empty) *repetition* of  $S$ . Given the set  $\mathcal{S}^*$  of tuple<sup>8</sup> sequences over  $\mathcal{S}$ , a process  $S$  defines a set of *driving scenarios* (also: situation traces or process runs) denoted by  $\llbracket S \rrbracket$  and defined by

$$\begin{aligned} \llbracket S \rrbracket &= \{S\} \iff S \in \mathcal{S} & \llbracket S; T \rrbracket &= \{st \in \mathcal{S}^* \mid s \in \llbracket S \rrbracket \wedge t \in \llbracket T \rrbracket\} \\ \llbracket S | T \rrbracket &= \llbracket S \rrbracket \cup \llbracket T \rrbracket & \llbracket S \parallel T \rrbracket &= \{(s, t) \in \llbracket S \rrbracket \times \llbracket T \rrbracket \mid (s, t) \models \Phi\} & \llbracket S^* \rrbracket &= \emptyset \cup \llbracket S; S^* \rrbracket \end{aligned}$$

<sup>8</sup>We use tuples to represent parallel composition.

where  $\Phi$  is a formula used to constrain  $\llbracket S \parallel T \rrbracket$  to admissible sequences of situation tuples in  $\mathcal{S}^*$ . A discussion of algebraic process properties and instances of  $\Phi$  would go beyond the scope of this paper.

We visualize a *driving process*  $S$  by a graph: Nodes represent situations and edges the successor relation for envisaged transitions between such situations. Bidirectional edges abbreviate two correspondingly directed edges. The anonymous initial situation is indicated by a gray dot. Such a graph can visualize the execution of *driving scenarios* compliant with the successor relation.

**Running Example Pt. 3** *Fig. 3b shows driving situations related to situation drivingAtL4Generic. Note that the model in Fig. 3b simplifies reality, e.g. halt can follow autoOvertake only indirectly after visiting driveAtL4Generic for an instant. We do not consider this as a restriction of expressiveness.*

## 2.2 Identification of Hazards for a Driving Situation

Knowing a fragment of the driving process, we can perform hazard identification and selection for the situations we identified. We assume that early-stage and process-level forward and backward analysis [8] can be applied, e.g. domain-specific checklists, hazard identification (HazId) or analysis (HazAn), hazard operability study (HazOp), failure-mode and effects (FMEA), event-tree (ETA), fault-tree (FTA), cause-consequence (CCA) [26], or system-theoretic process (STPA) [18] analysis.

**Running Example Pt. 4** *For the situation driveAtL4Generic, we conducted a light-weight, HazOp-like forward and backward analysis resulting in 10 causal factors, e.g.*

- *noAutoPilot (nAP)*, denoting the class of failures of an “autopilot” function, and
- *nearCollision (nC)*, denoting the event of a near-collision.

*Based on  $\sigma_{\mathcal{L}}$  (Sec. 2.1), nearCollision can be a predicate relating longitudinal and angular acceleration and direction of the AV and distance between the AV and other objects nearby.*

## 2.3 Grouping and Composing Driving Situations

Carrying on with process analysis can lead to a large  $\mathcal{S}$  representing the driving process for which we have to perform RAM. Hence, we group situations using regions. We call a process  $S$  a *region* (also: aspect) if it superimposes (by  $\parallel$ ) properties of a causal factor model on a subset of  $\mathcal{S}$ . Grouping criteria can be (i) specific modes of vehicle operation, e.g. the level of automation at which the vehicle is operated (see Sec. 1.1), (ii) situations across whom similar AD functions are used, e.g. requestTakeOverByDr, (iii) equivalence classes over process parameters, e.g.  $x_P$ .speed.


**Running Example Pt. 5** *Following Eq. (1), Fig. 4a models a driving process. Visualized in Fig. 4b, regions are shown in gray. For the region containing supplyPower, we applied forward reasoning (e.g. FMEA) to power-related vehicle components of the loop. We identified 3 sub-system failure modes as causal factors:*

- *lowOrNo-Fuel (F)*, denoting low fuel level or no more fuel,
- *lowOrNo-Energy (E)*, denoting reduced or outage of primary electric power supply, and
- *lowOrNo-Battery (B)*, denoting reduced or outage of secondary electric power supply.

*The results from FMEA provide information about the phases  $\underline{F}$ ,  $\underline{E}$ , and  $\underline{B}$  of these phase models.*

$$\begin{aligned}
P0 &\equiv \text{start}; (\text{basic} \parallel \text{supplyPower} \parallel D1) \\
D1 &\equiv (PP1 \parallel \text{autoLeaveParkingLot}); P1 \mid (PP1 \parallel \text{leaveParkingLot}); P2 \\
P1 &\equiv (PP2 \parallel \text{driveAtL4Generic}); D11 \\
P2 &\equiv (PP3 \parallel \text{driveAtL1Generic}); D12 \\
D11 &\equiv (PP2 \parallel (\text{driveThroughCrossing} \mid \text{exitTunnel} \mid \text{autoOvertake})); (P1 \mid P2) \mid \\
&\quad (PP2 \parallel \text{halt}); (P1 \mid P2 \mid (PP1 \parallel \text{parkWithRemote}; D1)) \mid \\
&\quad (PP1 \parallel \text{steerThroughTrafficJam}); P1 \\
D12 &\equiv P1 \mid (PP3 \parallel \text{manuallyOvertake}); P2 \mid (PP1 \parallel \text{manuallyPark}); D1 \\
PP1 &\equiv \text{driveAtLowSpeed} \\
PP2 &\equiv \text{drive} \parallel \text{driveAtL4} \parallel \text{requestTakeoverByDr} \\
PP3 &\equiv \text{drive} \parallel \text{driveAtL1} \parallel \text{operateVehicle}
\end{aligned}$$

(a) Despite “ $\parallel$ ” precedes “ $\mid$ ” and “ $\mid$ ” precedes “ $;$ ” we use parentheses for clarity.



(b) In blue, the *driving situations* comprising the controlled process; in gray, *regions* grouping properties of the *causal factor model* built for this example.

Figure 4: A *driving process* constructed according to Eq. (1): recursive expressions (a) and their visualization (b).

### 3 What Can Happen in a Specific Driving Situation?

In this section, we will investigate a structure for identifying and characterizing run-time risk.

#### 3.1 Risk Structures

For any situation  $s \in \mathcal{S}$ , we use the LTS model (Sec. 1.1) to create a *risk structure*  $\mathfrak{R}_s = (\Sigma_s, \mathcal{A}_s, \Delta_s)$  [10, 11] from parallel composition of *phase models* for the set of causal factors factorizing risk in  $s$ . The graphs in Figs. 1b, 2, and 5 show applications of the phase model:

- nodes depict *risk states*  $\subseteq \Sigma_s$  composed of phases, 0 denoting the “safest” state,
- red edges are transitions  $\subseteq \Delta_s$  to states of higher risk (*endangerments*  $\subseteq \mathcal{A}_s$ ), particularly, from 0,
- green edges are transitions  $\subseteq \Delta_s$  to states of lower risk (*mitigations*  $\subseteq \mathcal{A}_s$ ), preferably, towards 0.

In  $\mathfrak{R}_s$ , RUEs are states in  $\Sigma_s$  related with mishaps of high severity reached by corresponding combinations of endangerments. Although we can apply both, backward and forward analysis, to construct  $\mathfrak{R}_s$  using the phase model, this work goes along the lines of [11] and carries on with forward construction. For RAM, we consider this to be natural: We want to predict the *step-wise approach of states of higher risk reachable from any current state* (usually 0) and identify early mitigation steps (i.e., points of interception) rather than starting with mishap analysis (see, e.g. [31]).

**Running Example Pt. 6** *The situation drive represents a common situation composed in parallel with the situations basic, supplyPower, and the process (driveAtL4 || requestTakeoverByDr | driveAtL1 || operateVehicle)\* including its subordinate driving situations. We can understand drive as an aspect of these situations and, consequently, of the scenarios composed with drive. With drive, we associate the following CFs: badWeather (W), obstacleInTrajectory (O), nearCollision (nC), and collision (C).*

*Fig. 5a depicts a constrained composition of the phase models for W, O, nC, and C. In the following, particularly, in Sec. 3.4, we will discuss constraints that led to  $\mathfrak{R}_{\text{drive}}$ .*

#### 3.2 Selecting Hazards According to Driving Situations

As we have already seen, situations help focus on a specific part of the controlled process and, consequently, our hazard analysis. Particularly, situations can induce specific relationships between CFs. Moreover, relationships between a fixed pair of CFs can change from one situation to another. However, the final decision on the set of relevant CFs for a specific situation depends on many criteria addressed by the techniques mentioned in Sec. 2.2.

**Running Example Pt. 7** *A CF stemming from undesired behavior of human operators only has influence on a driving situation where an operator is actually in the loop. In our model, the causal factor dumbDriving (D) causes nC in any situation composed with the aspect operateVehicle, however, in any situation composed with the aspect driveAtL4, D is not part of the causes relationship with nC.*

#### 3.3 Which Events Increase Risk? What is an Endangerment Comprised of?

Based on the model from Secs. 1.1, 2.1, and 3, risk is assumed to be only increased by *endangerments*. Endangerments are those actions performed in  $\mathcal{L}$  that result in *activations* of causal factors, i.e., exactly those risk-increasing events we want our *safety controller* (Fig. 3a) to observe.






Figure 5: The situation drive as an aspect: (a) The whole drive aspect, (b) endangerment paths in drive starting from 0, (c) phase model for collision (C) and (d) for obstacleInTrajectory (O).

**Running Example Pt. 8** Fig. 5b shows combinations of endangerments in the situation drive for which the AV has to be prepared, e.g. when in 0. For example, the action  $e_d^O$  performs the activation of the causal factor O. Index d denotes that this action belongs to the class of disturbances.

**Classifying Endangerments and Causal Factors.** We classify a CF by the type of action that activates it, associated with the part of  $\mathcal{L}$  which performs this action. For this, we distinguish between

- failures ( $f$ ), stemming from technical parts of the controller,
- disturbances ( $d$ ), stemming from some parts of the controlled process,
- misuse ( $mu$ ), stemming from anyone who has access to the controls (e.g. human operator), and
- near-mishap ( $nm$ ), a special class used to model endangerments that lead to what we understand as “a mishap for which we are able to prepare effective countermeasures.”

**Running Example Pt. 9** In Fig. 5b, we can see actions of type d denoting disturbances, of type mu denoting misuse, and of type nm. Here, collision is an nm-typed CF rather than an unacceptable mishap because we assume to have an airbag Ab as a mitigation. In Fig. 5a, the action  $e_{nm}^C$  leads to a state with C activated.

### 3.4 Which Constraints Reduce the State Space Reachable by Endangerments?

The reader might have noticed that the *unconstrained composition* of the phase models for C, nC, O, and W in Fig. 5a would have resulted in a much larger risk structure. In fact, we applied several constraints to

our model to avoid the investigation of transitions irrelevant or not meaningful for designing mechanisms for RAM (cf. Sec. 1.2 and Fig. 3a).

For any risk state  $\sigma \in \Sigma$  and causal factors  $cf_1, cf_2$ , we can apply the *constraints*:

- $cf_1$  **requires**  $cf_2$  : $\Leftrightarrow$   $cf_1$  can only be activated if  $cf_2$  has already been activated,
- $cf_1$  **causes**  $cf_2$  : $\Leftrightarrow$  the activation of  $cf_1$  is propagated to  $cf_2$  (if not yet occurred),
- $cf_1$  **denies**  $cf_2$  : $\Leftrightarrow$  the activation of  $cf_1$  denies the activation of  $cf_2$  (if not yet occurred), and
- $cf_1$  **excludes**  $cf_2$  : $\Leftrightarrow$  factor  $cf_1$  renders  $cf_2$  irrelevant ( $cf_2$  is deactivated by  $cf_1$ ).

The usage of constraints will generally reduce  $\Sigma$ .<sup>9</sup> The application of these constraints in a specific causal factor model is based on, e.g. expert knowledge of our safety engineer and on results of system identification experiments performed by control engineers. In particular, **causes** relationships might have to be confirmed by, e.g. FMEA, **requires** relationships by, e.g. FTA.

## 4 Which Countermeasures Can We Employ At Run-Time?

In this section, we discuss the risk mitigation part of RAM based on the hazard analysis part of RAM shown in Sec. 3.

### 4.1 Which Events Decrease Risk? What is a Mitigation Comprised of?

Based on the model in Sec. 1.1, risk is supposed to be decreased by *mitigations*. Mitigations are those actions performed in  $\mathcal{L}$  that result in a mitigation and, finally, in a deactivation of causal factors, i.e., exactly those risk-decreasing events we want our *safety controller* (Fig. 3a) to actuate.

**Classifying Mitigations.** We classify mitigations by the type of action that deactivates a CF. The following criteria contribute to a taxonomy of mitigations: (i) the type of endangerment (see Sec. 3.3), (ii)  $\mathcal{L}$ 's mode of operation after mitigation (e.g. vehicle automation level, speed level), (iii) completion of deactivation of a causal factor at run-time. According to this, we distinguish between

- *fail-safe*, encompassing reactions to AV failures by *fail-silent* or *fail-operational* behavior,
- *deescalation*, comprising mechanisms for rejecting *disturbances* (e.g. stabilization),
- *protection*, including *failure-independent* mechanisms for risk prevention or alleviation of harm,
- *uncontrolled*, representing mitigation mechanisms *not* in control of the AV,
- *repair and maintenance*, representing mechanisms capable of complete deactivation of a CF or restoration of its undesired consequences, not necessarily at run-time.

**Running Example Pt. 10** For  $C$ , our model contains the action  $\text{prt}_{\text{alv}}^C$  which represents a protection mechanism resulting in the alleviation of any occurrence of  $C$  and its consequences. Finally, Fig. 5c shows the parts of  $\mathcal{L}$  embodying these actions, e.g.  $\frac{\text{prt}_{\text{alv}}^C}{\text{Ab}}$  indicates that  $\text{prt}_{\text{alv}}^C$  is going to be conducted by an airbag Ab as the alleviating protection mechanism.

<sup>9</sup>This goes along the lines of state space reduction in model checking, see e.g. [3].

## 4.2 Which Constraints Shape the State Space Reachable by Mitigations?

We can apply *shaping directives* for further reduction of  $\Sigma$ :

- ***cf direct***: *cf* can be completely mitigated, intentionally at run-time (this corresponds to the action  $m_c^{cf}$  shown in Fig. 1a),
- ***cf offRepair***: The change of *cf* to the phase *inactive* ( $0^{cf}$ ) by an action  $m_e^{cf}$  requires putting the AV out of order.

**Running Example Pt. 11** We use the following constraints and shaping directives for *nC* and *C*:

- ***nC requires O***: We only consider worlds in which *nearCollision* events require *obstacle-InTrajectory* events to occur in advance of and last until at least *nC*. Note that this constraint inhibits our CF model to be used for passive *nearCollisions*. This issue could be addressed by introducing an extra causal factor.
- ***nC direct***: Any implementation of a mitigation mechanism for *nC* performs a complete deactivation of *nC*.
- ***C requires nC***: We only consider worlds in which any *nearCollision* event precedes any possible *collision* event.
- ***C excludes nC***: Once a *collision* event happens we stop taking care about any further *nearCollisions*.
- ***C offRepair***: Our understanding of *collisions* in road traffic implies that a complete restoration, if possible, has to take place off-line, i.e., not in the situations shown in Fig. 4b. This way, the safety controller *M* knows that repair actions have to take place as soon as possible. So, “emergency stop” and “limp home” might become the next run-time actions to be performed by *M*.

## 4.3 Can Endangerments and Mitigations Change the Driving Situation?

In general, mitigations change  $\Sigma_{\mathcal{L}}$ , particularly, the state of the process *P*. Consequently, jumps from a risk state of one driving situation into a specific risk state of another situation can occur in many cases.

**Running Example Pt. 12** An emergency stop triggered by *lowOrNo-Fuel* in the situation *auto-Overtake* would bring the loop into the situation *halt* or, possibly, into some yet unknown situation in the region *driveAtLowSpeed*, while carrying on with further mitigations as required.

## 4.4 Towards Assessment by Symbolic Execution

By *symbolic execution* with YAP, we can assess our model for plausibility and refinement.

**Running Example Pt. 13** Tab. 1 shows a driving scenario starting from the situation *start* with brief statistics per situation describing the number of causal factors regarded in the current situation (#CFs), the initial state to which the random execution jumps at every step (only listing activated CFs), number of reachable risk states (particularly RUEs), and number of transitions (indicating some choice among mitigation options for *M*). Tab. 2 shows another driving scenario from the situation *steerThroughTrafficJam*. Both scenarios were cut after five steps.

Step	<i>Driving Situation</i>	#CFs	Initial State	#States	#Trans.
1	<i>start</i>				
2	<i>leaveParkingLot</i>	3	<i>B</i>	9	17
3	<i>driveAtL1Generic</i>	10	<i>nCoLI</i>	502	2165
4	<i>driveAtL4Generic</i>	10	<i>FEBWC</i>	1323	5726
5	<i>steerThroughTrafficJam</i>	3	<i>E</i>	9	15

Table 1: Scenario  $\rho \in \llbracket P0 \rrbracket$  starting from situation *start* (cf. *P0*, Fig. 4a).

Step	<i>Driving Situation</i>	#CFs	Initial State	#States	#Trans.
1	<i>steerThroughTrafficJam</i>	3	<i>B</i>	14	23
2	<i>driveAtL4Generic</i>	10	<i>EBOWnAP</i>	2190	8432
3	<i>exitTunnel</i>	10	<i>BOWrA</i>	3606	15337
4	<i>driveAtL1Generic</i>	10	<i>EBOWnCI</i>	276	868
5	<i>manuallyOvertake</i>	10	<i>EO</i>	756	2953

Table 2: Scenario  $\rho \in \llbracket D11 \rrbracket$  starting from situation *steerThroughTrafficJam* (cf. *D11*, Fig. 4a).

## 5 Synergies and Improvements Over Related Work

The questions S2 (Task S), E1-5 (Task E), and M1 (Task M) posed in Sec. 1.2 are typically addressed by *risk assessment techniques* such as, e.g. FTA, FMEA [8]. However, because of the specifics of RAM in AD, the performance of these techniques heavily depends on loop and causal factor models available and employed for RAM. The mentioned techniques rarely supply specialized and formally verifiable models. The causal factor model described in Sec. 1.1 incorporates the principle of “modeling what can go wrong” as employed in practical risk models such as CORAS [19]. Similar to, e.g. FTA, FMEA, and HazOp, the CORAS method itself is not specific to RAM for AD and the CORAS tool set has not yet been equipped with automation support for generating and analyzing risk state spaces. However, we consider a mapping from CORAS risk models into causal factor models and vice versa as a valuable extension. Similarly, the constraints described in Secs. 3.4 and 4.2 resemble some *gates* known from FTA, e.g. the AND gate can be modeled by a **requires** constraint.

Regarding the *modeling of driving processes*, [32] discusses a set of driving situations together with a specific safety argument based on hazards relevant in each situation. The approach at hand allows to transform such a situation taxonomy together with the arguments into a concise model.

For the *generation of risk state spaces*, we can mention several approaches: Volk et al. [35] provide an efficient technique to generate sparse MARKOV automata from dynamic fault trees and for the synthesis of parameters to calculate, e.g. the mean-time-to-failure measure. Regarding sparseness, YAP does not yet implement partial-order reduction to a comparable extent [3]. Bäckström et al. [2] describe an analytic approach to probabilistic, time-limited reachability of failure states based on continuous-time MARKOV chains derived from static and dynamic fault trees. Similarly, based on stochastic timed automata (STA), Kumar and Stoelinga [16] propose a time- and budget-limited analysis of a combined model of faults and attacks. They transform this model using event-class-specific STA templates rather than having a single generic phase model (cf. Fig. 1a). However, modeling the controlled process by situations combined with run-time planning over situation-specific risk state spaces provides an appropriate context for the employment of these algorithms and for quantitative assessment.

In the context of *safety standards for unmanned aircrafts*, [5] sketches a architecture where the

main controller is protected by a safety monitor and a recovery control. Regarding *controller synthesis*, Machin et al. [21] provide a systematic and formal procedure for the synthesis of intervention rules (i) identified through a UML-based HazOp technique and (ii) implementing a robot safety monitor. The rules are automatically verified based on (i) notions of warning and catastrophic states similar to our generic phase model in Fig. 1a and (ii) a measure of permissiveness. Mitsch and Platzer [25] present an intentionally similar approach, however, applicable to hybrid programs, a more expressive class of control systems.

A more specific line of work deals with *hybrid reachability analysis*, e.g. [29], where reachable sets are approximated over metric spaces. Such approaches are well-suited for performing optimal and safe *reach-avoid* mitigations for endangerments like, e.g. `nearCollisions`. Our approach aims at a model for developing mitigation strategies for multiple causal factors and performing mitigations including both, hybrid control actions and actions performing over non-metric spaces.

## 6 Discussion

This section highlights several aspects of the presented approach to be pursued by further research.

### 6.1 Notes on the Formalization of Constraints

Let each phase label in Fig. 1a be an atomic proposition. Then, in a timed extension of linear time temporal logic [22, 23, 15] and given causal factors  $cf_1$  and  $cf_2$ ,

- we translate **requires** into the formula  $\Box[cf_1 \rightarrow \blacklozenge_{\leq t} cf_2] \Leftrightarrow \Box[\neg cf_1 W_{\leq t} cf_2]$ ,
- **causes** into  $\Box[cf_1 \rightarrow \lozenge_{\leq t} (cf_2 U \neg cf_1)]$ ,
- **denies** into  $\Box[cf_1 \rightarrow \lozenge_{\leq t} (\neg cf_2 U \neg cf_1)]$ ,
- and **excludes** into  $\Box[cf_1 \rightarrow \lozenge_{\leq t} (0^{cf_2} U \neg cf_1)]$ .

By having one translation of these constraints into temporal formulas, we gain a way to translate (i) any causal factor into a temporal formula by logical conjunction ( $\wedge$ ), and, hence, (ii) the causal factor model into a corresponding formula. The formula resulting from (ii) characterizes the risk state space  $\Sigma$  without speaking of the action classes described in Secs. 3.3 and 4.1.

Furthermore, the **requires** constraint is a relaxed variant of the **globally precedes** pattern, e.g. discussed in [6]:  $cf_1$  **requires**  $cf_2 \Leftarrow$

$$cf_2 \text{ globally precedes } cf_1 \Leftrightarrow \lozenge cf_1 \rightarrow (\neg cf_1 U (cf_2 \wedge \neg cf_1)) \quad (2)$$

Eq. (2) does not hold for runs where  $cf_1$  and  $cf_2$  happen simultaneously and  $cf_2$  is required to happen.

### 6.2 Towards a Research Vision

Methodically, for RAM, it is important to ask why we need a dedicated causal factor model for AD? Well, let us refine Claim 1 from Sec. 1.2 by:

**Claim 2** *An optimal AV controller (AVC) always tries to reach and maintain a  $b$ -safe<sup>10</sup> state w.r.t. RUEs recognizable and acted upon by the AVC in any known and recognizable driving situation.*

<sup>10</sup> $b$  denotes a maximum “risk budget” to encode the notion of “acceptably safe,” see Sec. 1.1.

From Claim 2 and the answers to question M4, we can derive planning procedures and verification goals: are there (i) undesirable mitigations derivable from the model, (ii) undesirable combinations of situations and mitigations? How can we minimize exposure to mitigation-endangerment cycles? Can we prove from our model (given it is valid) the (global) existence of an (effective) *mitigation strategy* incorporating mechanisms from several parts of  $\mathcal{L}$ ? Technically, we aim at the instantiation of a monitor-actuator pattern [28] based on the resulting model of  $\mathcal{L}$  depicted in Fig. 3a.

### 6.3 About the Tasks S, E, and M

For Task S, we need to determine the physical events governing the driving process and its control, as well as the internal events in the implementation of the controller. For Task E, a predictive model of operational risk in any driving situation would be of great help. However, we only made a first step into the direction of discrete event model-predictive control. For Task M, our mitigation strategy always depends on the mechanisms and abilities feasible and controllable in  $\mathcal{L}$ .

### 6.4 Notes on Terminology, Assumptions, and Limitations

Our assumption for FAD, worth mentioning, is that *functional safety* of the controller is not considered in separation of *overall safety* of the control loop, which goes along the lines of [18]. It might seem cumbersome to speak of “rare” *undesired events*. However, *frequent* undesired events might stem from systematic defects or disturbances in the control loop for which general design-time measures might be desired before applying the discussed RAM approach.

We excluded the consideration of probabilities of causal factors and transitions in the driving process. However, stochastic reasoning helps in quantitative approximation for any planning algorithm employed for Claim 2. Moreover, probabilities will play a role in quantifying to which extent the model from Eq. (1) represents the controlled process. In Secs. 3.3 and 4.1, we left classifications of endangerments and mitigations shallow, though, YAP already incorporates a slightly more elaborate taxonomy.

### 6.5 Summary and Future Work

In this paper, we discussed *first steps* of an approach to risk analysis and run-time mitigation (RAM) suited for (i) the investigation of the driving process (Task S, S1-3), (ii) capturing multiple causal factors forming a set of rare undesired events in the risk state space (Task E, E1-5), and (iii) designing run-time mitigations to traverse the risk state space in a manner to establish Claims 1 and 2 (Task M, M1-4). Our model supports *safety engineering decision making* and the transfer of such decisions into a *safety controller*. This safety controller is supposed to implement the monitor in Fig. 3a as well as a *strategic mitigation planner* for run-time risk mitigation.

**Next Steps.** For the formalization of the discussed concepts, one next step is to provide MARKOV decision process semantics for the process model described in Sec. 2 and to map risk structures (i.e., the situation-specific LTS models described in Secs. 1.1 and 3) into equivalent KRIPKE structures. Extending these structures, we aim to equip our RAM approach for continuous update of real-time determined weights. Next, YAP which demonstrates part of the automated analyses for RAM (i.e., state space generation) has to be refined. Beyond these steps, we aim at the enhancement of the presented approach towards a controller design method integrating run-time risk mitigation for automated individual and collective driving.

**Acknowledgments.** I want to thank several engineers from German car makers and suppliers with whom I had many insightful discussions on this line of work. My sincere gratitude goes to my mentor and former advisor Manfred Broy for always putting highly valuable critique in a nutshell. I am grateful for a supportive work atmosphere in our collaboration with industry established by my project leader Stefan Kugele who has accompanied me with exciting previous work on this topic.

## References

- [1] Arthur M. Dowell III (1998): *Layer of protection analysis for determining safety integrity level*. *ISA Transactions* 37(3), pp. 155 – 165, doi:[10.1016/S0019-0578\(98\)00018-4](https://doi.org/10.1016/S0019-0578(98)00018-4).
- [2] Ola Bäckström, Yuliya Butkova, Holger Hermanns, Jan Krcál & Pavel Krcál (2016): *Effective Static and Dynamic Fault Tree Analysis*. In Skavhaug et al. [30], pp. 266–280, doi:[10.1007/978-3-319-45477-1\\_21](https://doi.org/10.1007/978-3-319-45477-1_21).
- [3] Christel Baier & Joost-Pieter Katoen (2008): *Principles of Model Checking*. MIT Press.
- [4] Manfred Broy & Ketil Stølen (2001): *Specification and Development of Interactive Systems: FOCUS on Streams, Interfaces, and Refinement*. Springer, Berlin, doi:[10.1007/978-1-4613-0091-5](https://doi.org/10.1007/978-1-4613-0091-5).
- [5] Stephen P. Cook (2017): *An ASTM Standard for Bounding Behavior of Adaptive Algorithms for Unmanned Aircraft Operations (Invited)*. AIAA SciTech Forum, American Institute of Aeronautics and Astronautics, doi:[10.2514/6.2017-0881](https://doi.org/10.2514/6.2017-0881).
- [6] Matthew B. Dwyer, G. S. Avrunin & J. C. Corbett (1999): *Patterns in property specifications for finite-state verification*. In: *ICSE*, pp. 411–20, doi:[10.1109/icse.1999.841031](https://doi.org/10.1109/icse.1999.841031).
- [7] R. Eastwood, T.P. Kelly, R.D. Alexander & E. Landre (2013): *Towards a safety case for runtime risk and uncertainty management in safety-critical systems*. In: *System Safety Conference incorporating the Cyber Security Conference 2013, 8th IET International*, pp. 1–6, doi:[10.1049/cp.2013.1713](https://doi.org/10.1049/cp.2013.1713).
- [8] Clifton A. Ericson (2015): *Hazard Analysis Techniques for System Safety*, 2nd edition. Wiley.
- [9] Mario Gleirscher (2014): *Behavioral Safety of Technical Systems*. Dissertation, Technische Universität München, doi:[10.13140/2.1.3122.7688](https://doi.org/10.13140/2.1.3122.7688).
- [10] Mario Gleirscher & Stefan Kugele (2017): *Defining Risk States in Autonomous Road Vehicles*. In: *High Assurance Systems Engineering (HASE), 18th Int. Symp.*, pp. 112–115, doi:[10.1109/hase.2017.14](https://doi.org/10.1109/hase.2017.14).
- [11] Mario Gleirscher & Stefan Kugele (2017): *From Hazard Analysis to Hazard Mitigation Planning: The Automated Driving Case*. In C. Barrett et al., editor: *NASA Formal Methods (NFM) – 9th Int. Symp., Proceedings, LNCS 10227*, Springer, Berlin/New York, pp. 310–326, doi:[10.1007/978-3-319-57288-8\\_23](https://doi.org/10.1007/978-3-319-57288-8_23).
- [12] Jeremie Guiochet, Mathilde Machin & Helene Waeselynck (2017): *Safety-critical Advanced Robots: A Survey*. *Robots and Autonomous Systems*, doi:[10.1016/j.robot.2017.04.004](https://doi.org/10.1016/j.robot.2017.04.004).
- [13] Charles A. R. Hoare (1985): *Communicating Sequential Processes*, 1st edition. Int. Series in Comp. Sci., Prentice-Hall.
- [14] Phil Koopman & Michael Wagner (2016): *Challenges in Autonomous Vehicle Testing and Validation*. In: *SAE World Congress*, doi:[10.4271/2016-01-0128](https://doi.org/10.4271/2016-01-0128).
- [15] Ron Koymans (1990): *Specifying real-time properties with metric temporal logic*. *Real-Time Syst.* 2(4), pp. 255–99, doi:[10.1007/bf01995674](https://doi.org/10.1007/bf01995674).
- [16] Rajesh Kumar & Mariëlle Stoelinga (2017): *Quantitative Security and Safety Analysis with Attack-Fault Trees*. In: *18th IEEE International Symposium on High Assurance Systems Engineering, HASE 2017, Singapore, January 12-14, 2017*, IEEE, pp. 25–32, doi:[10.1109/HASE.2017.12](https://doi.org/10.1109/HASE.2017.12).
- [17] Leslie Lamport (1977): *Proving the Correctness of Multiprocess Programs*. *IEEE Trans. Software Eng.* 3(2), pp. 125–43, doi:[10.1109/TSE.1977.229904](https://doi.org/10.1109/TSE.1977.229904).
- [18] Nancy Gail Leveson (2012): *Engineering a Safer World: Systems Thinking Applied to Safety*. Engineering Systems, MIT Press.

- [19] Mass Soldal Lund, Bjørnar Solhaug & Ketil Stølen (2011): *Model-Driven Risk Analysis: The CORAS Approach*, 1st edition. Springer, doi:[10.1007/978-3-642-12323-8](https://doi.org/10.1007/978-3-642-12323-8).
- [20] Jan Lunze (2010): *Regelungstechnik 1: Systemtheoretische Grundlagen, Analyse und Entwurf einschleifiger Regelungen*, 8th edition. Lehrbuch, Springer, doi:[10.1007/978-3-642-13808-9](https://doi.org/10.1007/978-3-642-13808-9).
- [21] Mathilde Machin, Jérémie Guiochet, Hélène Waeselynck, Jean-Paul Blanquart, Matthieu Roy & Lola Masson (2016): *SMOF – A Safety MONitoring Framework for Autonomous Systems* 99, pp. 1–14. doi:[10.1109/tsmc.2016.2633291](https://doi.org/10.1109/tsmc.2016.2633291).
- [22] Zohar Manna & Amir Pnueli (1991): *The Temporal Logic of Reactive and Concurrent Systems: Specification*, 1st edition. Springer.
- [23] Zohar Manna & Amir Pnueli (1995): *Temporal Verification of Reactive Systems: Safety*, 1st edition. Springer, doi:[10.1007/978-1-4612-4222-2](https://doi.org/10.1007/978-1-4612-4222-2).
- [24] Robin Milner (1995): *Communication and Concurrency*. International Series in Computer Science, Prentice Hall.
- [25] Stefan Mitsch & André Platzer (2016): *ModelPlex: Verified Runtime Validation of Verified Cyber-Physical System Models*. doi:[10.1007/978-3-319-11164-3\\_17](https://doi.org/10.1007/978-3-319-11164-3_17).
- [26] D.S. Nielsen (1971): *The cause/consequence diagram method as basis for quantitative accident analysis*. Technical Report RISO-M-1374, Danish Atomic Energy Commission.
- [27] On-Road Automated Driving Committee (2016): *Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles*. Technical Report SAE J 3016, SAE International, doi:[10.4271/j3016\\_201609](https://doi.org/10.4271/j3016_201609).
- [28] Christopher Preschern, Nermin Kajtazovic & Christian Kreiner (2013): *Building a safety architecture pattern system*. In Uwe van Heesch & Christian Kohls, editors: *Proceedings of the 18th European Conference on Pattern Languages of Programs (EuroPLoP), Irsee, Germany, July 10-14, 2013*, ACM, p. 17, doi:[10.1145/2739011.2739028](https://doi.org/10.1145/2739011.2739028).
- [29] Hendrik Roehm, Jens Oehlerking, Thomas Heinz & Matthias Althoff (2016): *STL Model Checking of Continuous and Hybrid Systems*, pp. 412–27. Springer, doi:[10.1007/978-3-319-46520-3\\_26](https://doi.org/10.1007/978-3-319-46520-3_26).
- [30] Amund Skavhaug, Jérémie Guiochet & Friedemann Bitsch, editors (2016): *Computer Safety, Reliability, and Security - 35th International Conference, SAFECOMP 2016, Trondheim, Norway, September 21-23, 2016, Proceedings*. Lecture Notes in Computer Science 9922, Springer, doi:[10.1007/978-3-319-45477-1](https://doi.org/10.1007/978-3-319-45477-1).
- [31] I. Svedung & J. Rasmussen (2002): *Graphic representation of accident scenarios: Mapping system structure and the causation of accidents*. *Safety Science* 40(5), pp. 397–417, doi:[10.1016/s0925-7535\(00\)00036-9](https://doi.org/10.1016/s0925-7535(00)00036-9).
- [32] Transport Systems Catapult (2017): *Taxonomy of Scenarios for Automated Driving*. Technical Report, Transport Systems Catapult.
- [33] Jan Tretmans (2008): *Model Based Testing with Labelled Transition Systems*. In: *Formal Methods and Testing*, pp. 1–38, doi:[10.1007/978-3-540-78917-8\\_1](https://doi.org/10.1007/978-3-540-78917-8_1).
- [34] U.S. Department of Transportation (2016): *Federal Automated Vehicles Policy*. Technical Report, U.S. Department of Transportation.
- [35] Matthias Volk, Sebastian Junges & Joost-Pieter Katoen (2016): *Advancing Dynamic Fault Tree Analysis - Get Succinct State Spaces Fast and Synthesise Failure Rates*. In Skavhaug et al. [30], pp. 253–265, doi:[10.1007/978-3-319-45477-1\\_20](https://doi.org/10.1007/978-3-319-45477-1_20).