

# Simulation-based Safety Assessment of High-level Reliability Models

Simon József Nagy

Budapest University of Technology and Economics, Budapest, Hungary  
Department of Measurement and Information Systems

simon.jozsef.nagy@gmail.com

Bence Graics

Kristóf Marussy

András Vörös

Budapest University of Technology and Economics, Budapest, Hungary  
Department of Measurement and Information Systems

MTA-BME Lendület Cyber-physical Systems Research Group

graics@mit.bme.hu

marussy@mit.bme.hu

vori@mit.bme.hu

Systems engineering approaches use high-level models to capture the architecture and behavior of the system. However, when safety engineers conduct safety and reliability analysis, they have to create formal models, such as fault-trees, according to the behavior described by the high-level engineering models and environmental/fault assumptions. Instead of creating low-level analysis models, our approach builds on engineering models in safety analysis by exploiting the simulation capabilities of recent probabilistic programming and simulation advancements. Thus, it could be applied in accordance with standards and best practices for the analysis of a critical automotive system as part of an industrial collaboration, while leveraging high-level block diagrams and statechart models created by engineers. We demonstrate the applicability of our approach in a case study adapted from the automotive system from the collaboration.

## 1 Introduction

Safety-critical cyber-physical systems, such as embedded control systems in the automotive domain, must satisfy numerous stringent extra-functional requirements, such as safety, reliability, and availability, in addition to functional requirements. The ISO 26262 [16] standard for automotive systems requires addressing these issues during system design by following recommended development practices and demonstrating compliance. To comply with the standards and stakeholder requirements, the certification process of the system uses top-down deductive reliability modeling and analysis of the architecture and behaviors to verify safety properties.

Increasingly complex behaviors of embedded automotive systems, especially in the case of fault avoidance and other safety mechanisms, pose significant challenges in these analyses. Not only classical hardware redundancy mechanisms but also fault-tolerant sensor fusion algorithms and adaptive reconfiguration are applied. Thus, we need expressive modeling approaches to precisely describe the system behavior and faults, as well as tool support for calculating reliability metrics.

Our goal in this work is to use high-level statechart-based languages in the safety and reliability analysis process and support the reuse of engineering models defined in the system design phase, such as state machines and block diagrams. Therefore, we developed an integrated automotive reliability modeling and analysis approach, which is able to capture the architecture and the behavior of the system using high-level engineering models. We provide a statechart language in the Gamma framework for the engineers to design not only the engineering models but also the error propagation and fault/environmental

assumption models. Our approach exploits the simulation capabilities of recent probabilistic programming and simulation advancements (e.g., [4, 8, 14, 15]), and we provide the tool support to generate the input for these techniques. Therefore, we can calculate the compliance of the safety requirements and also support the decision-making process in all phases of the system development process.

We demonstrate our modeling and analysis approach in the context of a safety-critical automotive electric and electronics (E/E) system. Nevertheless, our approach is general and may be used for other adaptive cyber-physical systems and in other domains too.

## 1.1 Case study

We demonstrate our top-down modeling and analysis method in the context of a real-life example from the automotive industry, namely on an electronic control unit (*ECU*) of an electronic power-assisted steering (*EPAS*) unit [6] developed by Thyssenkrupp Components. Due to its critical role in the car, it has several safety functions, including adaptive reconfiguration and high redundancy in the hardware. In this case study, we calculate the probabilistic measures required by the ISO 26262 standards. Note that the case study was slightly changed compared to the real system for protecting intellectual property. However, the changes do not substantially affect its structure and behavior.

The simplified structural model of a typical, widely-used *EPAS ECU* is depicted in Figure 1a. It has two microcontrollers (*uC*) that are separated and can operate individually. A *uC* has three redundant sensors, each of which can have an operational state and two failure states, namely, *Shutdown*, and *Drift*. If the sensor stops due to a failure, the model enters the *Shutdown* state, which can be detected every time. In contrast, the *Drift* state represents a latent failure mode: in this state, the sensor seems to work correctly, yet it has an erroneous output. Thus, the detection of this failure mode requires redundant sensors and a voting mechanism. With the help of the sensors, the *uCs* provide the steering-actuation functionality using a closed control loop (a safety-critical function, *SCF*) that must operate during the lifetime of the car continuously.

Figure 1/b) depicts the simplified statechart<sup>1</sup> describing the behavior of a *uC*. The initial state is *Normal operation*. The model goes to state *AssistLoss* if the *uC* fails or all sensors fail (go to state *Shutdown*). In contrast, the *uC* goes to state *LatentError* if at least two sensors have latent error as the bad sensors will vote down the good one. Therefore, the *uC* will use wrong input data in the control loop, causing the *EPAS* (the states of which are modeled in an evaluation statechart) to go to state *Uncontrolled self-steering*. The *EPAS* system can fail in two ways as it can be seen in Figure 1/c):

- If both *uCs* go to state *AssistLoss*, the steering assistance will also stop operating, resulting in a troublesome steering experience. This state of the *EPAS* is called *Loss of assist* (LoA).
- If a *uC* goes to state *LatentError*, the whole *EPAS* will go to state *Uncontrolled self-steering* (SS). This situation is extremely dangerous and must be prevented by any means, as the driver loses control over the car.

---

<sup>1</sup> Variables in the statechart:

- drift num : number of drifted sensors
- ok num : number of normal sensors
- on num : number of sensors that seem to be working

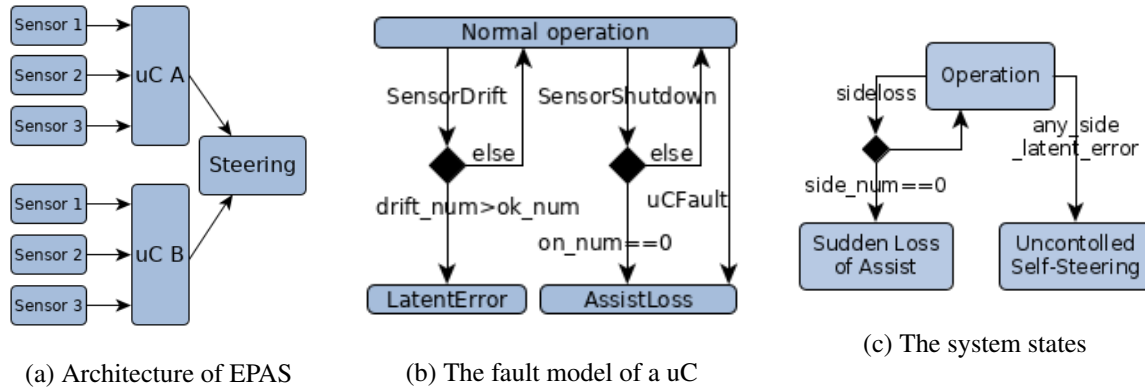


Figure 1: Structural and behavioral models of the EPAS

## 1.2 Challenges of the safety analysis

In the design of complex electronics and electronics (E/E) automotive systems, practitioners are well versed in classical top-down fault modeling with fault trees, as well as bottom-up analysis with failure modes, effects and diagnostic analysis (FMEDA). These techniques allow demonstrating system safety according to the applicable standards by manually creating and reviewing analysis models that are subsequently evaluated by software tools.

However, the increasing complexity of systems, such as distributed error detection and fail-over mechanisms in the EPAS, pose significant challenges for the classical approach.

Firstly, classical models may become very large and cumbersome to handle. For example, in the EPAS case study, a fault tree model required 92 logic gates even after simplifying the system behavior, while describing the full behavior and diagnostics is impossible with fault trees, as they do not support the definition of state-based behavior. Even though there are significantly more expressive formalisms for modeling stochastic systems [17], they often require specific expert knowledge, and cannot be applied widely by engineers. Therefore, experts must create (either automatically or by code generation) and review stochastic models based on high-level architecture models.

Secondly, different stochastic analyses are performed throughout the system design not only to demonstrate safety but also to inform design decisions and fix errors. For example, Pareto and sensitivity analysis find components most often responsible for safety goal violations, which are candidates for changes. Additionally, fitting parametric distributions in Weibull analysis [1] provides information about the product life cycle. These analyses often require purpose-built models with tailored abstractions, such as a different fault tree in the classical approach for every analysis question, which multiplies expert effort.

Lastly, the high number of hardware and software components poses a challenge for analysis tools due to the state space explosion. With more than four orders of magnitude difference between rare fault modes, iterative numerical analysis methods can be rendered ineffective, while simulation-based methods require support for rare event sampling.

These challenges made us seek an integrated modeling and analysis approach that can be applied by engineers working on the EPAS project. At the same time, it is effective enough for answering varied analysis questions throughout the development and certification process.




Figure 2: Activities and artifacts in the safety assessment approach

### 1.3 Overview of the approach

To overcome these challenges, we created an integrated, top-down safety modeling and analysis method (illustrated in Figure 2) for complex critical systems, such as automotive E/E components. It supports the safety assessment of the system defined by a functional model using high-level modeling languages, such as block diagrams and statecharts with which engineers are familiar.

The first step of our approach is the construction of the high-level reliability model, which specifies the failure modes of the hardware components, the behavior of safety mechanisms and the propagation of errors as interconnected statecharts. The resulting three-layered, state-based reliability model (presented in Section 2.3) defines the possible states of the hardware components and safety mechanisms, whereas the system configuration specifies the connections between the software components as well as their interaction modes. Even though most of these models, such as the behavioral and structural models of the safety mechanisms are available at the beginning of the analysis in the form of design artifacts, e.g., SysML state machines and internal block diagrams, some of these models have to be created during deductive safety analysis. Thus, the same systems engineering tools can be used for creating the analysis models that are used for the design.

Secondly, the fault distributions (presented in Section 3.3) specify the stochastic behavior of the state-transitions in the state-based reliability model. Similarly to more traditional approaches, distributions are obtained from standards or FMEDA analysis.

To facilitate analysis, we developed a *Probabilistic Runtime Environment (PRE)* that processes the high-level models for reliability analysis using the deep probabilistic programming [4] paradigm. Various analyses can be specified as probabilistic programs, and the inference engine can handle models with large state spaces and rare events. Therefore, we can use our analysis approach both during design with complex analysis questions and for demonstrating system safety.

To our best knowledge, this is the first application of the deep probabilistic programming to the reliability analysis of safety-critical systems.

## 2 Background and related work

### 2.1 Safety metrics

Critical systems have to satisfy various functional and extra-functional requirements. Dependability-related extra-functional requirements ensure that the system provides its functionalities/services steadily. In our paper, we focus on *reliability* and *availability* analysis. The reliability of a system represents the continuity of the correct service [3], and the availability of a system represents readiness for the correct service.

During safety assessment, all these concepts have to be analyzed with several techniques, e.g., life-time prediction, which discovers when and how the system can fail. These analysis methods can produce several properties of the system, e.g., *mean-time-to-first-failure* (MTFF) and the *failure-in-time* (FIT). MTFF determines the average time until the first system failure, and FIT determines the probability of the system failure during a given period of time. To calculate these measures, deductive, top-down analysis has to be applied.

## 2.2 Analysis solutions

One of the most popular deductive analysis methods is the fault tree analysis (FTA) [9]. It supports the modeling of static systems; however, modeling components with an inner state is impossible with this method. This poses limitations for modeling reconfigurable and self-diagnosing systems. Continuous-time Markov-chains (CTMCs) [18] and related formalisms can model a wide range of behaviors. The explicit solution of CTMCs is hindered by the state-space explosion, i.e., the number of states becomes extremely large due to the complexity of the system-under-analysis. Various abstraction techniques were developed to tackle this challenge [17], but they typically require significant modeling and analysis expertise.

Statistical model checking methods and tools, such as UPPAAL-SMC [5] were also developed: these approaches rely on random sampling to simulate the behavior of the system. The resulting data set can be analyzed using standard statistical analysis methods, e.g., statistical tests and Weibull analysis [1]. UPPAAL-SMC is widely used in the railway [21], automotive [10] and aerospace [23] domains. Even though it can model a wide range of behaviors by providing a low-level modeling language, direct modeling of large systems may be cumbersome. Translation from higher-level models to UPPAAL is possible [19], but may substantially increase model size.

In addition, simulation-based approaches are also widely used for FTA in the field of nuclear power [27], electric power distribution systems [26] and wastewater treatment [24]. The common disadvantage of these approaches is that they have to model every failure mode and operational mode of the system separately. As a result, they have to create a distinct fault tree for every failure mode, and they are unable to analyze the joint distribution of the failure modes. Unfortunately, the standard statistical methods are unable to analyze conditional models, even though they are used extensively during the safety assessment of critical systems.

## 2.3 State-based modeling

State machines provide a convenient formalism to model the behavior of reactive systems. State machine models process incoming events and react to them in accordance with their internal states. Statecharts [12] are a popular extension of state machines providing complex constructions to support the high-level design of reactive systems. This formalism contains several expressive modeling elements, such as variables, arithmetic expressions, and interfaces with parametric events.

The Gamma Statechart Composition Framework<sup>2</sup> [19] is an open-source, integrated modeling toolset to support the semantically sound composition of heterogeneous statechart components [11]. The framework reuses statechart modeling languages of third-party tools and their code generators, e.g., MagicDraw<sup>3</sup> and Yakindu Statechart Tools<sup>4</sup>, for separate components. As a core element, the framework

---

<sup>2</sup><http://gamma.inf.mit.bme.hu/>

<sup>3</sup><https://www.nomagic.com/products/magicdraw>

<sup>4</sup><https://www.itemis.com/en/yakindu/state-machine/>

provides the Gamma Composition Language, which supports the interconnection of components hierarchically based on precise semantics. Gamma provides automated code generators as well as test case generators for the analysis of interactions between components. Gamma also supports system-level formal verification and validation (V&V) functionalities, i.e., the system model can be exhaustively analyzed with respect to formal requirements, by mapping statechart and composition models into the input representations of verification back-ends.

## 2.4 Probabilistic programming

In order to formulate complex statistical inference problems, such as Bayesian machine learning and differential privacy applications as computer programs, a new programming paradigm called *probabilistic programming* has been developed, which explicitly allows sampling from probability distributions as part of a program. In the last ten years, many tools have appeared to support this paradigm, e.g., Stan [7], Edward [25], Anglican [2], PyMC [22] and Pyro [4].

In comparison to other tools and approaches, the greatest advantage of probabilistic programming tools is the general *inference algorithm*, which can analyze conditional models and calculate posterior distributions independently from the particular formulation of the model. These conditional models are created by placing *observe* statements in the program, which define the observed (conditional) distributions of random variables. By running the probabilistic program in an inference environment, we obtain the corresponding posterior distributions.

Most inference algorithms use a gradient-based Monte-Carlo method, such as the Hamiltonian Monte-Carlo [8] and the No-U-Turn algorithms [15]. Recently, a new approach called deep probabilistic programming [4] has emerged relying on deep learning algorithms for fast and efficient computations. In this work, we relied on the stochastic variational inference algorithm (SVI) [14] from this family, which fits a parametric distribution to the output of the probabilistic program. It optimizes the following evidence lower bound (ELBO) between the parametric *guide* function  $q_\varphi(\mathbf{z})$  and the sampled posterior  $p_\theta(\mathbf{x}, \mathbf{z})$  (dependent on the parameters  $\varphi$  and  $\theta$ ) distribution of the output of the probabilistic program, where  $\mathbf{x}$  and  $\mathbf{z}$  are the conditioned and latent random variables, respectively:  $\text{ELBO} = \mathbb{E}_{q_\varphi(\mathbf{z})} [\log p_\theta(\mathbf{x}, \mathbf{z}) - \log q_\varphi(\mathbf{z})]$ .

By choosing an appropriate guide  $q_\varphi(\mathbf{z})$  and condition  $\mathbf{x}$ , different aspects of the program can be analyzed. In the probabilistic programming environment Pyro [4], both the guide, the program, and the conditions can be specified as Python programs in a user-friendly manner.

## 3 Modeling the EPAS ECU

### 3.1 State-based system modeling

Our approach introduces a fault-modeling technique based on high-level state-based (statechart) models. We describe the system under analysis as a composition of software and hardware elements with Yakindu statecharts and Gamma modeling languages using the Gamma Composition Framework. Statecharts provide an expressive language to capture complex error propagation scenarios of systems. Therefore, in our approach, we can apply statechart-based modeling (interconnected statecharts) to describe the faulty behavior in the system, e.g., latent, cascading, multi-point, and common-cause failures.

With the expressive modeling elements of the aforementioned languages, we can model complex safety mechanisms. We use variables and arithmetic expressions to describe the complex decision-making algorithms in the reconfiguration strategies and the sensor fusion algorithms. In addition, para-

metric events in statechart interfaces can model both the communication and the fault propagation between the components.

We propose a top-down, three-step approach for the modeling of the EPAS system based on both the system models and safety requirements. In Step I, the high-level states and components of the system are defined, which is followed by Step II, the specification of safety mechanisms. Finally, in Step III, elementary hardware components are modeled.

**I. System-level behavior** In the first step, we model the high-level state of the system from the perspective of functional and extra-functional requirements. Thus, we create the *system statechart* (depicted in Figure 4), which contains the operating and failure states of the EPAS, namely, *Normal*, *Loss of Assist* and *Self-steering*.

After that, we define a *composition model* of the system, which contains communication channels between the communication ports of each statechart model. Error events and reconfiguration events propagate from the lower-level components (e.g., hardware and safety mechanisms) to the higher-level components (e.g., the system-level statechart) according to the semantics of the Gamma Composition Language.

Components are represented by statecharts, which will be acquired or defined in the subsequent modeling steps. The composed model describes the behavior of the full system, while engineers can focus on the different statechart models at the appropriate levels of abstraction for top-down (deductive) analysis. This model contains every component of the system, the failure of which can contribute to the failure of the whole system and defines how they can interact with each other. The error propagation model of the EPAS ECU is depicted in Figure 5.

**II. Safety-level behavior** In the second step, we specify the behavior of the safety mechanisms using the statechart models created by the systems engineers. In our analysis, these models were already available as part of the engineering models of the system.

The state of the EPAS ECU is determined by the state of the motor controllers. Consequently, the system statechart is connected to the statechart models of the controllers (depicted in Figure 7). This controller switches off the actuation if either the uC is faulty or the sensor diagnostic sends an error status message. Consequently, the behavior of a controller is affected greatly by the behavior of the diagnostic function. Therefore, the controller statechart is connected to a diagnostic statechart (depicted in Figure 8), which defines the voting mechanism in the diagnostic function. Note that the expressive modeling elements of statecharts, such as orthogonal regions, variables, and arithmetic expressions greatly facilitate the specification of the complex behavior of the sensor diagnostics.

**III. Hardware-level behavior** In the last step, we model the elementary hardware subcomponents of the system, which are the smallest portion of hardware components considered in the safety analysis. These hardware components have an independent and distinct functionality from the perspective of the safety analysis. We can model these hardware elements with statecharts, which contain the operating and failure modes of the hardware components. Such statecharts can be created based on standards [13], supplier information, and the result of other analysis methods, e.g., FMEDA.

Generally, the statecharts of most hardware components, such as the uC in the ECU (Figure 10), have only two states: a *Good* and a *Faulty*. Although, in the case of complex adaptive systems, the failure mode greatly influences the state of the system. The sensor (Figure 9) has an operational state and two

vastly different failure states: *Shutdown* and *Drift*. The *Shutdown* state can cause only *Loss of assist* failure whereas the *Drift* state is the only source of the *Self-steering* failure.

### 3.2 Composite model

The statechart models introduced in Steps I, II, and III are composed into a single, deterministic, state-based model of the system according to the composition model defined in Step I.

The semantics of the composition are defined by the Gamma framework and are similar to composition facilities provided by SysML block diagrams, Extended Timed Automata, and tools like UPPAAL, albeit with support for ports, interfaces, and various synchronous and asynchronous communication semantics. For a description of the semantics, we direct the interested reader to [11, 19], and Appendix A.

As an example of event propagation in the composite model of the system, consider the transitions to the states named *SelfSteering* in the system level statechart in Figure 4 and motor controller in Figure 7. Entry to the *SelfSteering* state in any of the two instances of the motor controller statechart, which correspond to the two redundant uCs of the EPAS, raises the event *selfsteering* on their respective *Monitor* ports. The system-level statechart listens to these events on its *MonitorA* and *MonitorB* ports, which are connected to outputs of the two controllers. Any of these events trigger the transition from the *Operation* state to the *SelfSteering* state, and raise the *SS* event on the *Eval* port, signalling the safety goal violation.

### 3.3 Fault distributions

In order to model not only the functions and services of the system but also its extra-functional aspects (including reliability and availability), we have to affix fault occurrence distributions to the low-level hardware faults in our model.

The Yakindu and Gamma modeling languages cannot express stochastic behaviors. Therefore, we model the stochastic transitions separately, by annotating them in an external table (see Table 2). During simulation-based analysis, the distributions are sampled to generate (timed) fault sequences, which are input to the composite state-based model as a sequence of low-level hardware events. Safety goal violations can be ascertained as output events from the model.

## 4 Analysis of the system

In this section, we present an analysis method for deterministic, state-based composite models annotated by fault distribution, which can be obtained using the methodology in the previous section.

Our analysis is based on deep probabilistic programming. While probabilistic programs, such as the PRISM guarded command language, are widely used in reliability analysis, to our best knowledge, this is the first use of an inference engine with *observe* statements for conditional reliability analysis and distribution fitting.

### 4.1 Probabilistic Runtime Environment

In order to analyze the composite model of the system with randomly sampled fault sequences using deep probabilistic programming, we created a bridge between the state-based model and the probabilistic programming environment, called the Probabilistic Runtime Environment (PRE). This environment allows running the state-based model as part of a deep probabilistic program. The PRE can be executed



on its own to compute the mean lifetime of the system, or as part of a probabilistic program containing *observe* statements to facilitate conditional analysis and distribution fitting for inference.

PRE is implemented in Python to be compatible with the most popular stochastic analysis tools [4, 7, 22]. The implementation is built on top of the Pyro language, since Pyro includes the inference algorithms needed for the analysis (summarized in Section 2.4), as well as state-of-the-art inference tools, such as GPU accelerated deep probabilistic programming.

## 4.2 Translation to a probabilistic program

Discrete models are defined in the Gamma framework, out of which a Java implementation is generated. This implementation can evaluate the system-level effects of the hardware faults via interface functions. The model of the system is then directly translated into a probabilistic program that simulates random hardware faults and the system behavior.

Our approach is based on continuous-time, asynchronous simulation. The probabilistic program generates the component failure events randomly for all failure modes of all components in the system at once. Thus, we sample each failure-mode distribution for each component and arrange the resulted events in chronological order to produce a stochastic event series for the discrete model, which determines the failure-time and failure-mode. However, the number of possible failure combinations grows exponentially with the number of hardware components in the system. Pyro provides several algorithms that can mitigate this problem, such as interest sampling and SVI.

The pseudo-code of the probabilistic simulator program is shown in Algorithm 1. This algorithm consists of the following steps:

1. First, the fault distribution of every failure mode of every hardware component is sampled and collected into the fault set.
2. Then, the fault events are arranged in chronological order to get fault series. This procedure is essential for our reliability assessment since, during a simulation, each fault event is sent to the model of the software components individually.
3. Finally, the result of a component failure series is evaluated in a *while* loop. If the system reaches a failure state during the evaluation cycle, the *while* loop stops, and the simulation returns with both the failure state and the elapsed time to reach the failure state.

All of these steps are integrated into the *simulate* function of the *PRE*. As a result, the generated probabilistic program is available for the analysis scripts, which can both validate the system requirements and identify the weak points in the system.

## 4.3 Stochastic analysis of the case study

### 4.3.1 Lifetime prediction

To predict the time to failure (TFF) of the system, the probabilistic program in Algorithm 1 can be run repeatedly to sample from the TFF distribution and estimate the MTTFF by stochastic model checking methods [5]. The results were visualized in a histogram (see Figure 3).

To reveal the aging of the failures of the EPAS, we performed Weibull analysis [1] using stochastic variational inference (SVI) from Pyro as outlined in Section 2.4. We set  $q_{\{\eta, \beta\}}(time) = Weibull(\eta, \beta; time)$  as the guide function, where  $Weibull(\eta, \beta; time)$  is the pdf of a Weibull distributed random variable with  $\eta$  and  $\beta$  scale and shape parameters respectively. In this case, the set of conditioned variables is empty, i.e.,  $\mathbf{x} = \emptyset$ , since we use neither conditioning nor the *observe* statement.

The SVI optimization algorithm in the *PRE* explores the optimal Weibull shape ( $\beta$ ) and scale ( $\eta$ ) parameters, which reveals the fitted lifetime distribution. This method helps to understand the changes in the failure rate of the system over time, as indicated by the shape parameter  $\beta$ .

### 4.3.2 Conditional analysis

For safety assessment, we can examine even complex, low-probability events in the system-under-analysis, such as the MTTF in case of a specified failure mode. By utilizing the Pyro deep probabilistic programming algorithm inside the *PRE*, we are able to analyze joint as well as conditional distributions. We can use all the inference algorithms implemented in Pyro (e.g., importance sampling, Hamiltonian Mont-Carlo, and SVI) to analyze efficiently complex conditional and parametric distributions even with low occurrence probability [20].

To apply these inference algorithms, we first have to create a posterior model in the *PRE*, namely, a *conditional simulation model*. In this model, we assume that some random variables in the simulation will have a specified value (e.g., we assume that the system will go to a specified failure mode state). In order to create such a model, we put *observe* statements for each conditioned random variable in the generated *simulate* function. Thereafter, we either put the *conditional simulation model* into a Pyro inference algorithm and run Monte Carlo simulations or we use the model fitting (SVI) algorithm of Pyro (introduced in the previous subsection) with an appropriate *guide* function.

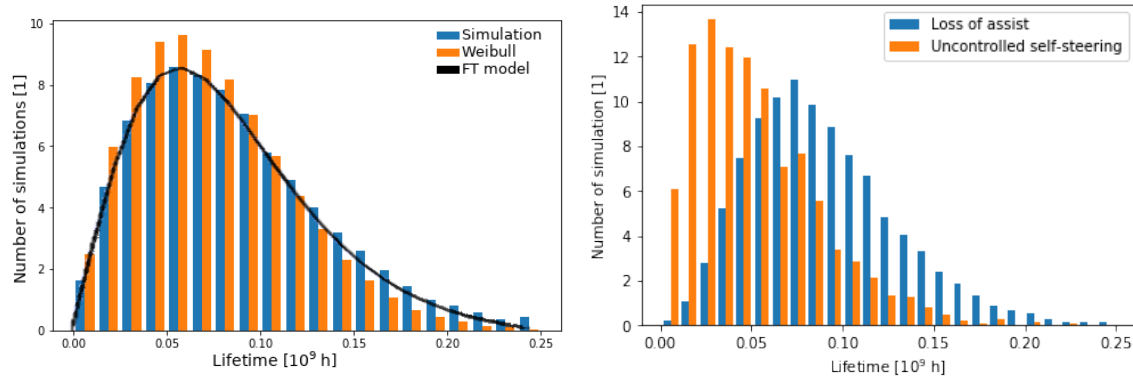
The conditional analysis can be used for component sensitivity calculations. The main objective of this method is to investigate how the lifetime of the system changes if a given component fails. We created a *conditional model* where we assume (*observe*) that a given component fails during the mission time. Thus, the weaknesses of the system design could be identified and remedied.

## 5 Evaluation

### 5.1 Compatibility with ISO 26262 safety analysis requirements

The safety analysis of critical components has to be approved by several industry-specific standards. Therefore, we defined our analysis method in accordance with the ISO 26262-2018 standard as it is one of the most modern and relevant standards. In the following, we present the conformity between our method and the ISO 26262:

- Our analysis technique supports system development in the full life-cycle, including the examination of new ideas with modular analysis even when the system model is incomplete.
- If the safety requirements are not fulfilled, we can construct a conditional analysis, which can identify the weaknesses of the system design. Moreover, with the help of the Gamma framework, we are able to functionally analyze the components of the system separately.
- The use of our analysis method does not require any special competences. Thus the safety analysis can be an integrated part of the development process. As a result, statechart-based component models can be reviewed directly by the engineers, and the engineers by themselves are able to create component models without any training.



(a) Comparison of the simulated and fitted lifetime distribution

(b) Comparison of the lifetime distribution of the failure modes

Figure 3: Lifetime prediction and conditional lifetime prediction with simulation

## 5.2 Analysis of the EPAS ECU

### 5.2.1 Lifetime prediction

We ran 10,000 independent simulations in the *PRE* and visualized the results in Figure 3a with the blue histogram. As can be seen, the results resemble the characteristics of the Weibull distribution. Thus, using the Pyro optimizers and SVI, we can fit a Weibull distribution to the system behavior, as shown in Listing 3. We ran 10,000 steps with the SVI algorithm to fit the model. The comparison of the real and the fitted Weibull model is depicted in Figure 3a, where the fitted Weibull model is represented with the orange chart. As can be seen, the resulting model is not perfect, but it gives a good approximation for the scale and the shape parameters. These parameters give an insight into how the system changes over time. In addition, the parameters may be used for system-level analysis, including all subsystems of the car.

To validate the results, we also conducted a fault tree analysis (FTA) manually on the EPAS model. As it is depicted in Figure 3a, the results of the FTA matched those of simulation closely.

### 5.2.2 Conditional lifetime prediction

The ISO26262 standard requires a detailed analysis of the failure modes, therefore we have to analyze the *conditional behavior*, the *Uncontrolled self-steering* and the *Loss of assist* failure modes from the perspective of the expected occurrence time.

We created two conditional models in the *PRE* that give us the posterior distribution of the lifetime, assuming that the failure mode is *Loss of assist* or *Uncontrolled self-steering*. The associated Pyro probabilistic program is illustrated in Listing 4. We ran the inference algorithm 10,000 times. The comparison of the two resulted posterior lifetime distributions are depicted in Figure 3b. The results meet the expectations, *Uncontrolled self-steering* occurs much earlier than *Loss of assist* due to the following reason: if a sensor has a latent failure and any other sensor has any kind of problem, the system will go to state *Uncontrolled self-steering* immediately. The *Loss of assist* failure mode occurs when both sides have a uC fault or three sensors have *Shutdown* type faults. Reaching *Loss of assist* takes a much longer time than *Uncontrolled self-steering*.

#microcontrollers	2	2	4	4
#sensors	6	12	24	48
Estimated state-space size	$10^{20}$	$10^{40}$	$10^{60}$	$10^{80}$
TTF analysis time (sec)	20.8	33.3	49.6	89.7
Conditional TTF analysis time (sec)	82.9	122.6	201.3	449.1

Table 1: Running time of the benchmark model simulations in seconds

### 5.3 Run-time measurements of the algorithms

In order to validate our analysis method, we made several run time measurements. We examined the time-to-failure analysis as well as the conditional time-to-failure analysis. We ran all analysis scripts five times and calculated the median and checked the consistency of the results.

We evaluated our analysis method on several different models. We used the EPAS, defined in Section 1.1, as a template of the benchmark models. Thus we created three new versions of this model. Each new model version was extended with some additional sensors and some new uCs. Similarly to the original EPAS model, each uC has a voting mechanism for their sensors, and the system goes to state *Loss of assist* if no uC is operational and goes to state *Uncontrolled self-steering* if at least one of them uses bad sensor-data for the steering control. For the measurements, we used an average PC configuration<sup>5</sup>. Median running times of the analyses are presented in Table 1.

The results show that the run-time of a single simulation scales well as the size of the model increases and the analysis method can be applied even for large models.

The results of the run time measurements are depicted in Table 1. Note that all analysis scripts run successfully within 10 minutes.

## 6 Conclusion

The top-down analysis of adaptive critical embedded automotive systems is a challenging task due to the high redundancy in the hardware and the applied complex reconfiguration strategies in the software. Traditional modeling and analysis methods do not support the sensor fusion algorithms, the fail-operational adaptation, and the dependent failures. In this work, we introduced a new top-down analysis approach based on the Gamma framework and extended it with stochastic distributions. For the analysis, we created a *PRE*, which provides an easy-to-use interface even for complex analysis techniques, e.g., SVI. Finally, we applied the implemented algorithms on a power-steering model from the automotive industry, and the results show that our algorithm scales well even for large models. Subject to future work, we plan to include the analysis of the sensor fusion algorithms and system effects on the hardware failures, where the modeled controller can optionally modify the failure/environmental distributions.

**Acknowledgment** This work was partially supported by the ÚNKP-19-3 New National Excellence Program of the Ministry for Innovation and Technology and the European Union, co-financed by the European Social Fund (EFOP-3.6.2-16-2017-00013). We thank Dr. Pter Gyrke and ThyssenKrupp Presta Kft. for the fruitful collaboration and the case study.

<sup>5</sup>Ubuntu 18.04.3 LTS; AMD Athlon X4 860K CPU; 8 GB RAM; GeForce GTX 1050Ti; Pyro-0.4.1; PyTorch-0.4.1

## References

- [1] R B Abernethy, J E Breneman, C H Medlin & G L Reinman (1983): *Weibull Analysis Handbook*. AFWAL-TR 83-2079, Air Force Wright Aeronautical Laboratories. Available at <https://apps.dtic.mil/dtic/tr/fulltext/u2/a143100.pdf>.
- [2] Christophe Andrieu, Arnaud Doucet & Roman Holenstein (2010): *Particle Markov chain Monte Carlo methods*. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 72(3), pp. 269–342, doi:10.1111/j.1467-9868.2009.00736.x.
- [3] Algirdas Avizienis, Jean-Claude Laprie, Brian Randell & Carl Landwehr (2004): *Basic Concepts and Taxonomy of Dependable and Secure Computing*. Technical Report 2004-47, University of Maryland. Available at [https://drum.lib.umd.edu/bitstream/handle/1903/6459/TR\\_2004-47.pdf](https://drum.lib.umd.edu/bitstream/handle/1903/6459/TR_2004-47.pdf).
- [4] Eli Bingham, Jonathan P. Chen, Martin Jankowiak, Fritz Obermeyer, Neeraj Pradhan, Theofanis Karaletsos, Rohit Singh, Paul Szerlip, Paul Horsfall & Noah D. Goodman (2019): *Pyro: Deep Universal Probabilistic Programming*. *J. Mach. Learn. Res.* 20(1), p. 973978, doi:10.5555/3322706.3322734.
- [5] Peter Bulychev, Alexandre David, Kim Gulstrand Larsen, Marius Mikučionis, Danny Bøgsted Poulsen, Axel Legay & Zheng Wang (2012): *UPPAAL-SMC: Statistical model checking for priced timed automata*. In: *QAPL 2012, Elec. Proc. Theor. Comput. Sci.* 85, pp. 1–16, doi:10.4204/EPTCS.85.1.
- [6] A. W. Burton (2003): *Innovation drivers for electric power-assisted steering*. *IEEE Control Systems Magazine* 23(6), pp. 30–39, doi:10.1109/MCS.2003.1251179.
- [7] Bob Carpenter, Andrew Gelman, Matthew D Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li & Allen Riddell (2017): *Stan: A probabilistic programming language*. *J. Stat. Softw.* 76(1), doi:10.18637/jss.v076.i01.
- [8] Tianqi Chen, Emily Fox & Carlos Guestrin (2014): *Stochastic Gradient Hamiltonian Monte Carlo*. In Eric P. Xing & Tony Jebara, editors: *Proceedings of the 31st International Conference on Machine Learning, Proceedings of Machine Learning Research* 32, PMLR, Beijing, China, pp. 1683–1691.
- [9] Joanne Bechta Dugan, Kevin J Sullivan & David Coppit (2000): *Developing a low-cost high-quality software tool for dynamic fault-tree analysis*. *IEEE Trans. Reliab.* 49(1), pp. 49–59, doi:10.1109/24.855536.
- [10] Predrag Filipovikj, Nesredin Mahmud, Raluca Marinescu, Cristina Seceseanu, Oscar Ljungkrantz & Henrik Lnn (2016): *Simulink to UPPAAL Statistical Model Checker: Analyzing Automotive Industrial Systems*. In: *FM 2016*, pp. 748–756, doi:10.1007/978-3-319-48989-6\_46.
- [11] Bence Graics & Vince Molnár (2018): *Mix-and-Match Composition in the Gamma Framework*. In: *25th Minisymposium, Department of Measurement and Information Systems*, Budapest, Hungary.
- [12] David Harel (1987): *Statecharts: A Visual Formalism for Complex Systems*. *Sci. Comput. Program.* 8(3), pp. 231–274, doi:10.1016/0167-6423(87)90035-9.
- [13] J. W. Harms (2010): *Revision of MIL-HDBK-217, Reliability Prediction of Electronic Equipment*. doi:10.1109/RAMS.2010.5448046.
- [14] Matthew D. Hoffman, David M. Blei, Chong Wang & John Paisley (2013): *Stochastic Variational Inference*. *J. Mach. Learn. Res.* 14(1), p. 13031347, doi:10.5555/2567709.2502622.
- [15] Matthew D. Homan & Andrew Gelman (2014): *The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo*. *J. Mach. Learn. Res.* 15(1), p. 15931623, doi:10.5555/2627435.2638586.
- [16] ISO/TC 22/SC 32 (2018): *Road vehicles — Functional safety — Part 9: Automotive Safety Integrity Level (ASIL)-oriented and safety-oriented analyses*. ISO 26262-9:2018, International Organization for Standardization. Available at <https://www.iso.org/standard/51365.html>.
- [17] Joost-Pieter Katoen (2016): *The Probabilistic Model Checking Landscape*. In: *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 16*, Association for Computing Machinery, New York, NY, USA, p. 3145, doi:10.1145/2933575.2934574.

- [18] Marta Kwiatkowska, Gethin Norman & David Parker (2011): *PRISM 4.0: Verification of Probabilistic Real-time Systems*. In: *CAV 2011, LNCS 6806*, Springer, pp. 585–591, doi:10.1007/978-3-642-22110-1\_47.
- [19] Vince Molnár, Bence Graics, András Vörös, István Majzik & Dániel Varró (2018): *The Gamma Statechart Composition Framework*. In: *ICSE 2018, ACM*, pp. 113–116, doi:10.1145/3183440.3183489.
- [20] Yura N. Perov, Logan Graham, Kostis Gourgoulis, Jonathan G. Richens, Ciarán M. Lee, Adam Baker & Saurabh Johri (2019): *MultiVerse: Causal Reasoning using Importance Sampling in Probabilistic Programming*. *ArXiv abs/1910.08091*.
- [21] Enno Ruijters & Mariëlle Stoelinga (2016): *Better Railway Engineering Through Statistical Model Checking*. In: *Leveraging Applications of Formal Methods, Verification and Validation: Foundational Techniques*, Springer, pp. 151–165, doi:10.1007/978-3-319-47166-2\_10.
- [22] J. Salvatier, T. V. Wieckiâ & C. Fonnesbeck (2016): *PyMC3: Python probabilistic programming framework*. Astrophysics Source Code Library. Available at <https://ui.adsabs.harvard.edu/abs/2016ascl.soft10016S>.
- [23] Lijun Shan, Yuying Wang, Ning Fu, Xingshe Zhou, Lei Zhao, Lijng Wan, Lei Qiao & Jianxin Chen (2014): *Formal Verification of Lunar Rover Control Software Using UPPAAL*. In Cliff Jones, Pekka Pihlajasaari & Jun Sun, editors: *FM 2014: Formal Methods*, Springer, pp. 718–732, doi:10.1007/978-3-319-06410-9\_48.
- [24] Masoud Taheriyoun & Saber Moradinejad (2015): *Reliability analysis of a wastewater treatment plant using fault tree analysis and Monte Carlo simulation*. *Environmental Monitoring and Assessment* 187(1), p. 4186, doi:10.1007/s10661-014-4186-7.
- [25] Dustin Tran, Alp Kucukelbir, Adji B Dieng, Maja Rudolph, Dawen Liang & David M Blei (2016): *Edward: A library for probabilistic modeling, inference, and criticism*. *arXiv preprint arXiv:1610.09787*.
- [26] P. Zhang & K. W. Chan (2012): *Reliability Evaluation of Phasor Measurement Unit Using Monte Carlo Dynamic Fault Tree Method*. *IEEE Transactions on Smart Grid* 3(3), pp. 1235–1243, doi:10.1109/TSG.2011.2180937.
- [27] Marko epin & Borut Mavko (2002): *A dynamic fault tree*. *Reliability Engineering & System Safety* 75(1), pp. 83 – 91, doi:10.1016/S0951-8320(01)00121-1.

## A Gamma composition semantics

Gamma is a modeling framework for the semantically precise composition of statechart components. Statecharts (considered as atomic components) can be composed in the Gamma Composition Language (GCL), which supports the definition of *synchronous* and *asynchronous* composite components, two fundamentally different system types determining how their constituent components receive events and how they are executed. In the subsequent sections, we informally introduce the communication elements in GCL as well as the cascade composition mode of synchronous systems, as we used this composition mode for defining the EPAS configuration. Additional information on the synchronous and cascade composition modes, as well as asynchronous systems, can be found in [11].

### A.1 Communication Elements

In GCL, components (both atomic and composite) communicate through *ports*. Each port defines a point of service through which certain *event notifications* can be sent or received. An event notification (or event for short) is a piece of information passed between components, which can also have *parameters* to forward data. An event is called *message* in the case of asynchronous components and *signal* in the case of synchronous components. Events are declared on *interfaces*, which may be realized by ports. An event may be declared as *input* or *output*. The declared directions are reversed, however, if the port does not *provide*, but *require* the interface, which are the two possible modes in which a port can realize an interface.

### A.2 Synchronous components

Synchronous components represent models that communicate in a synchronous manner using *signals*. They are executed in a lockstep fashion, triggered by an enclosing component (synchronous or asynchronous) or an external actor from the environment. When executed, synchronous components process incoming signals and produce output signals in accordance with their internal states. Input signals are not queued but sampled: upon execution, the component can access the most recent signal for each event on every port since the last execution (if there is any). Similarly, output signals are reset at the beginning of every execution and each output event on every port can get a new signal assigned to it.

Synchronous components in Gamma are *statechart definitions*, which are considered atomic components as well as *synchronous* and *cascade composite components*, which can be freely mixed in hierarchically composed synchronous systems.

### A.3 Cascade composite components

Conceptually, components in a cascade composite model represent a set of “filters” through which inputs are transformed into outputs. Therefore, constituent components immediately see the output signals of other components in the same composite component during execution. By default, constituent components are executed once in the *order of their instantiation*. Alternatively, an execution list can be defined that determines the execution order of instantiated components. The execution list can contain a particular constituent component many times, supporting repeated execution. The typical arrangement of a cascade composite component definition is illustrated in the following snippet.

```

cascade Epas [
  // System port declarations
  port S1AFault: requires SensorFault
  // ...
] {
  // Component instances
  component S1A: SensorStatechart
  component DiagA: DiagnosticStatechart
  // ...
  // Binding composite model ports to internal ports
  bind S1AFault->S1A.HWFFault
  // ...
  // Channel definitions connecting internal ports:
  // A provided and required realization of the same interface
  channel [S1A.SensorFault] -o)- [DiagA.S1HW]
  // ...
}

```

Listing 1: The textual representation of a typical Gamma cascade composite component

## B Models

### B.1 System layer

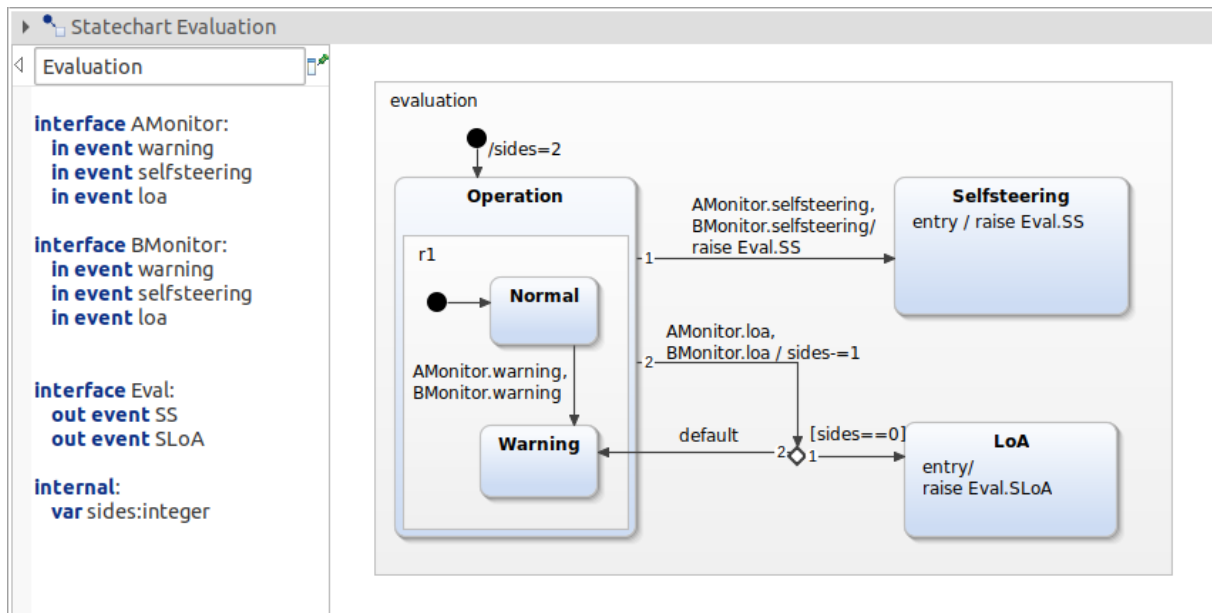


Figure 4: System level statechart of the EPAS



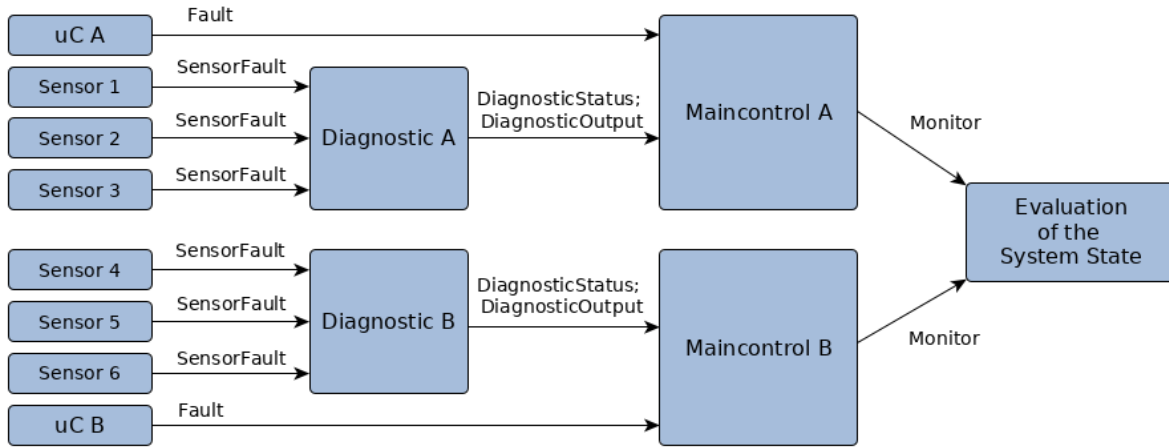


Figure 5: Layered structure of the EPAS model. Labels on arrows refer to the interfaces which define the propagated error and reconfiguration events

```

package epas
import "interfaces/Interfaces.gcd"

cascade Epas [
  port State: provides Eval
  port S1AFault: requires SensorFault
  port S2AFault: requires SensorFault
  port S3AFault: requires SensorFault
  port S1BFault: requires SensorFault
  port S2BFault: requires SensorFault
  port S3BFault: requires SensorFault
  port UCAFault: requires UCFault
  port UCBFault: requires UCFault
] {
  component S1A: SensorStatechart
  component S2A: SensorStatechart
  component S3A: SensorStatechart

  component S1B: SensorStatechart
  component S2B: SensorStatechart
  component S3B: SensorStatechart

  component DiagA: DiagnosticStatechart
  component DiagB: DiagnosticStatechart

  component UCA: UCStatechart
  component UCB: UCStatechart

  component ACTRL: MainctrlStatechart
  component BCTRL: MainctrlStatechart

  component Ev: EvaluationStatechart

  bind S1AFault->S1A.HWFFault
  bind S2AFault->S2A.HWFFault
  bind S3AFault->S3A.HWFFault

  bind S1BFault->S1B.HWFFault
  bind S2BFault->S2B.HWFFault
  bind S3BFault->S3B.HWFFault
}
  
```

```

bind UCAFault->UCA.HWFault
bind UCBFault->UCB.HWFault

bind State->Ev.Eval

channel [S1A.SensorFault] -o)- [DiagA.S1HW]
channel [S2A.SensorFault] -o)- [DiagA.S2HW]
channel [S3A.SensorFault] -o)- [DiagA.S3HW]

channel [S1B.SensorFault] -o)- [DiagB.S1HW]
channel [S2B.SensorFault] -o)- [DiagB.S2HW]
channel [S3B.SensorFault] -o)- [DiagB.S3HW]

channel [DiagA.DiagnosticOutput] -o)- [ACTRL.DiagnosticOutput]
channel [DiagA.DiagnosticStatus] -o)- [ACTRL.DiagnosticStatus]
channel [DiagB.DiagnosticOutput] -o)- [BCTRL.DiagnosticOutput]
channel [DiagB.DiagnosticStatus] -o)- [BCTRL.DiagnosticStatus]

channel [UCA.Fault] -o)- [ACTRL.UCHW]
channel [UCB.Fault] -o)- [BCTRL.UCHW]

channel [ACTRL.Monitor] -o)- [Ev.AMonitor]
channel [BCTRL.Monitor] -o)- [Ev.BMonitor]
}

```

Listing 2: Gamma textual model of the EPAS configuration

```

interface UCFault {
  out event shutdown
}

interface SensorFault {
  out event det
  out event latent
}

interface DiagnosticStatus {
  out event Error
  out event Warning
}

```

```

interface Eval {
  out event SS
  out event SLoA
}

interface DiagnosticOutput {
  out event WrongOutput
}

interface Monitor {
  out event warning
  out event loa
  out event selfsteering
}

```

Figure 6: Gamma textual models of the EPAS interfaces

## B.2 Safety layer

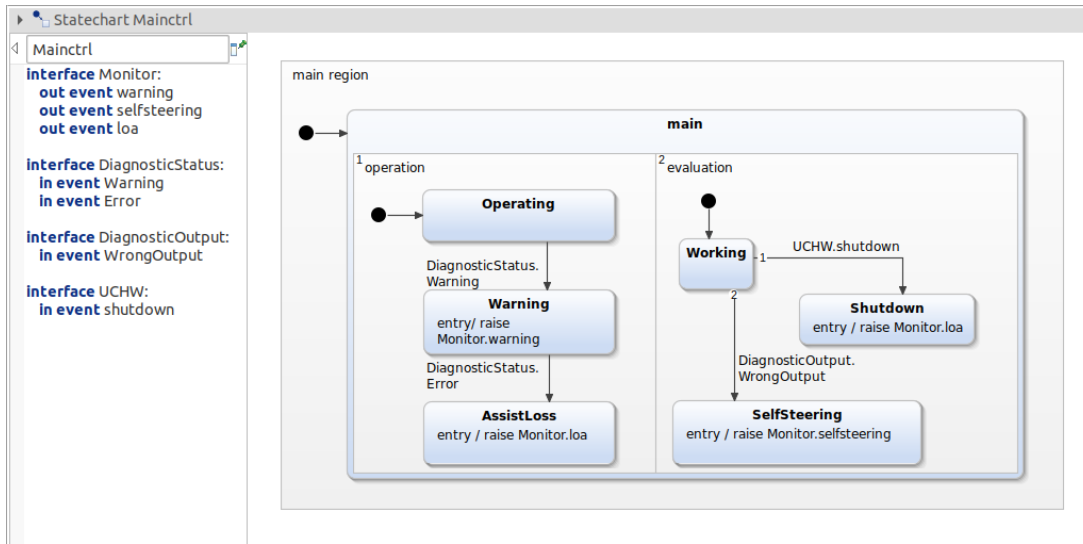


Figure 7: Statechart of the motor controller

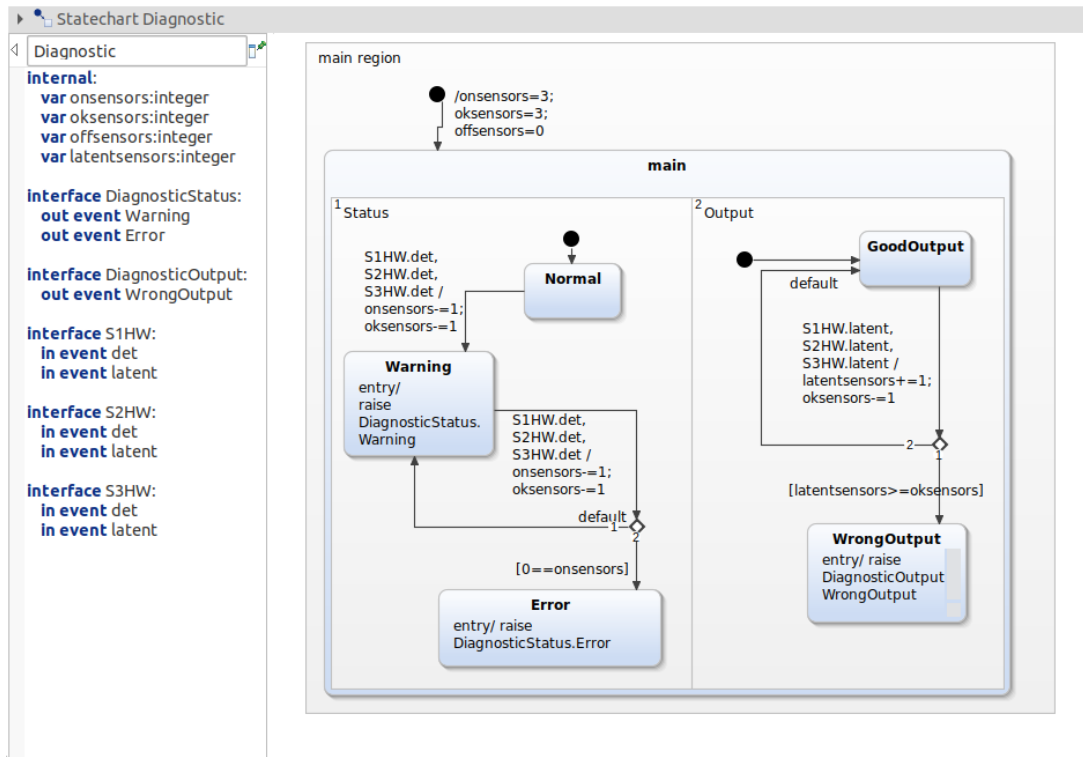


Figure 8: Statechart of the sensor diagnostics

### B.3 Hardware layer

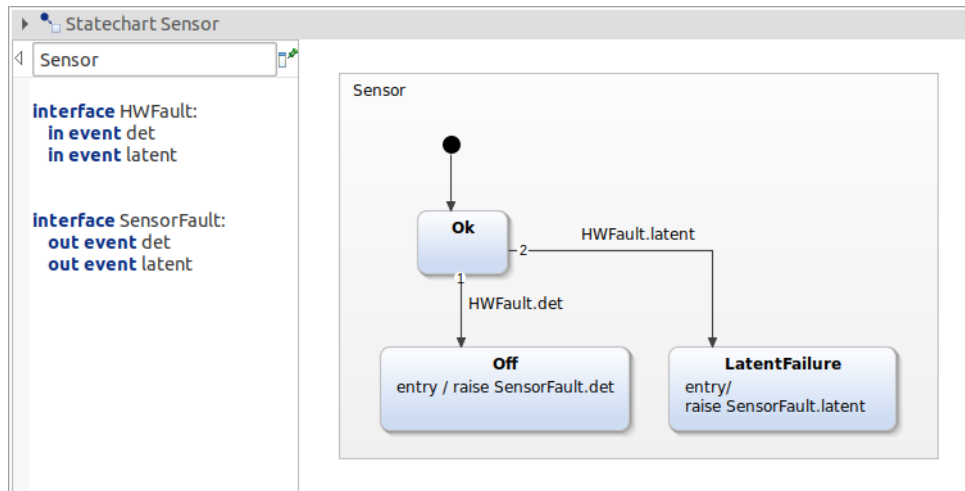


Figure 9: Statechart based failure model of the sensor

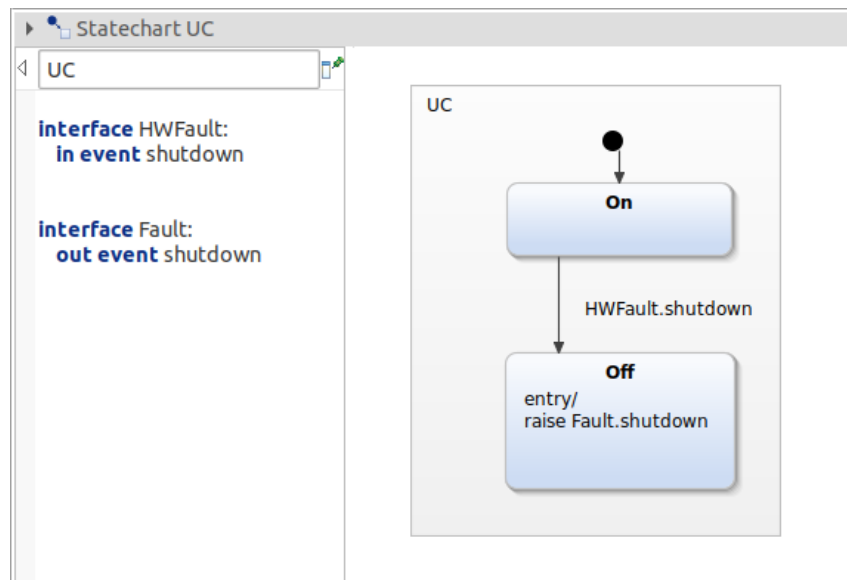


Figure 10: Statechart based failure model of the uC sensor

Failure distribution	Distribution parameters	HW statechart	From state	To state
Weibull	concentration=1.5, scale= $0.1 \cdot 10^{-9} \frac{1}{h}$	uC	On	Off
Exponential	rate= $10.0 \cdot 10^{-9} \frac{1}{h}$	Sensor	Ok	Off
Exponential	rate= $1.0 \cdot 10^{-9} \frac{1}{h}$	Sensor	Ok	LatentFailure

Table 2: Connection between the distributions and the state transitions in the case-study

## C Probabilistic Runtime Environment

---

**Algorithm 1** Pseudo code of the translated probabilistic program

---

```

faults  $\leftarrow \emptyset$ 
for all component  $\in$  system do
  for all failure_mode  $\in$  component do
    fault_time  $\leftarrow$  sample(failure_mode.distribution)
    faults  $\leftarrow$  {⟨fault_time; component; failure_mode⟩}  $\cup$  faults
  end for
end for
faults  $\leftarrow$  orderByTime(faults)
state  $\leftarrow$  "Normal"
while state == "Normal" do
  fault  $\leftarrow$  faults.next()
  state  $\leftarrow$  GammaModel.getResultOf(fault.component, fault.failure_mode)
end while
return {fault.time; state}

```

---

```

# Instantiate the generated Probabilistic Runtime Environment
simulate = create_pru()
# The guide function encodes the fitted distribution
def guide():
    scale = pyro.param("scale", torch.tensor(1.0), constraint=constraints.positive)
    shape = pyro.param("shape", torch.tensor(1.0), constraint=constraints.positive)
    pyro.sample("failure_time", dist.Weibull(scale, shape))
# Set up SVI with the Adam optimizer
optimizer = pyro.optim.Adam({"lr": 0.05, "betas": (0.9, 0.999)})
svi = pyro.infer.SVI(simulate, guide, optimizer, loss=pyro.infer.Trace_ELBO())
n_steps = 500
for step in range(n_steps):
    svi.step()
# Extract numerical solution
scale = pyro.param("scale").item()
shape = pyro.param("shape").item()

```

Listing 3: Probabilistic program for Weibull distribution fitting with SVI

```

# Instantiate the generated Probabilistic Runtime Environment
simulate = create_pru()
# Failure mode 0 (passed to torch.tensor) corresponds to self-steering
conditional = pyro.condition(simulate, {"failure_mode": torch.tensor(0.0)})
sampler = pyro.infer.Importance(conditional, num_samples=10000)
# Generate a histogram of the conditional distribution
marginal_dist = pyro.inter.EmpiricalMarginal(sampler.run(), sites="failure_time")

```

Listing 4: Probabilistic program for conditional analysis