

From Probabilistic Counterexamples via Causality to Fault Trees

Matthias Kuntz¹, Florian Leitner-Fischer², and Stefan Leue²

¹ TRW Automotive GmbH, Germany

² University of Konstanz, Germany

Abstract. In recent years, several approaches to generate probabilistic counterexamples have been proposed. The interpretation of stochastic counterexamples, however, continues to be problematic since they have to be represented as sets of paths, and the number of paths in this set may be very large. Fault trees (FTs) are a well-established industrial technique to represent causalities for possible system hazards resulting from system or system component failures. In this paper we suggest a method to automatically derive FTs from counterexamples, including a mapping of the probability information onto the FT. We extend the structural equation approach by Pearl and Halpern, which is based on Lewis counterfactuals, so that it serves as a justification for the causality that our proposed FT derivation rules imply. We demonstrate the usefulness of our approach by applying it to an industrial case study.

1 Introduction

In recent joint work [1] with our industrial partner TRW Automotive GmbH we have proven the applicability of stochastic formal analysis techniques to safety analysis in an industrial setting. In [1] we showed that counterexamples are a very helpful means to understand how certain error states representing hazards can be reached by the system. While the visualization of the graph structure of a stochastic counterexample [2] helps to analyze the counterexamples, it is still difficult to compare the thousands of paths in the counterexample with each other, and to discern causal factors during fault analysis. In safety analysis, fault tree analysis (FTA) [21] is a well-established industrial method and graphical notation to break down the hazards occurring in complex, technical systems into a combination of what is referred to as basic events, which represent system component failures. The main drawback of fault tree analysis is that it relies on the ability of the engineer to manually identify all possible component failures that might cause a certain hazard. In this paper we present a method that automatically generates a fault tree from a probabilistic counterexample. Our method provides a compact and concise representation of the system failures using a graphical notation that is well known to safety engineers. At the same time the derived fault tree constitutes an abstraction of the probabilistic counterexample since it focuses on representing the implied causalities rather than enumerating all possible execution sequences leading to a hazard. The causality expressed by

the fault tree is rooted in the counterfactual notion of causality that is widely accepted in the literature. Our approach can be described by identifying the following steps:

- Our fault tree computation method uses a system model given in the input language of the PRISM probabilistic model checker [14].
- For this model we compute counterexamples for stochastic properties of interest, representing system hazards, using our counterexample computation extension of PRISM called DiPro [2]. The counterexamples consist of potentially large numbers of system execution paths and their related probability mass information.
- In order to compute fault trees from these counterexamples we compute what is commonly referred to as basic events. Those are events that cause a certain hazard. The fault tree derivation is implemented in a tool called CX2FT.
- The justification for the fault tree computation is derived from a model of causality due to Halpern and Pearl [12] that we modify and extend to be applicable to our setting.
- The path probabilities computed by the stochastic model checker are then mapped on the computed fault tree.
- Finally, the obtained fault tree is represented graphically by an adapted version of the FaultCAT tool³.

All analysis steps are fully automated and do not require user intervention. We demonstrate the usefulness of our approach by applying it to a selection of case studies known from the literature on stochastic model checking.

This paper is organized as follows: In Section 2 we briefly introduce the concepts of counterexamples in stochastic model checking and fault trees. In Section 3 we describe the model of causality that we use, and how probabilistic counterexamples can be mapped to fault trees. In Section 4 we demonstrate our approach on a case study known from the literature. A discussion of related work follows in Section 5. Finally, Section 6 concludes the paper.

2 Counterexamples and Fault Trees

In stochastic model checking, the property that is to be verified is specified using a variant of temporal logic. The temporal logic used in this paper is Continuous Stochastic Logic (CSL) [4]. Given an appropriate system model and a CSL property, stochastic model checking tools such as PRISM [14] can verify automatically whether the model satisfies the property. Stochastic model checkers do not automatically provide counterexamples, but the computation of counterexamples has recently been addressed in, amongst others, [3, 13]. For the purpose of this paper it suffices to consider only upper bounded probabilistic timed reachability properties. They require that the probability of reaching a certain state, often corresponding to an undesired system state, does not exceed a certain upper

³ <http://www.iu.hio.no/FaultCat/>

probability bound p . In CSL such properties can be expressed by formulae of the form $\mathcal{P}_{\leq p}(\varphi)$, where φ is path formula specifying undesired behavior of the the system. A counterexample for an upper bounded property is a set Σ_C of paths leading from the initial state to a state satisfying φ such that the accumulated probability of Σ_C violates the probability constraint $\leq p$. If the CSL formula $\mathcal{P}_{=?}(\varphi)$ is used, the probability of the path formula φ to hold is computed and the counterexample contains all paths fulfilling φ . The probability of the counterexample is computed using a stochastic model checker, in our case PRISM. Notice that in the setting of this paper the counterexample is computed completely, i.e., all simple paths leading into the undesired system state are enumerated in the counterexample. Fault trees (FTs) [21] are being used extensively in industrial practice, in particular in fault prediction and analysis, to illustrate graphically under which conditions systems can fail, or have failed. In our context, we need the following elements of FTs, for an in-depth discussion of FTs we refer the reader to [21].

1. Basic event: represents an atomic event.
2. *AND*-gate: represents a failure, if all of its input elements fail.
3. *OR*-gate: represents a failure, if at least one of its input elements fails.
4. Priority-*AND* (*PAND*): represents a failure, if all of its input elements fail in the specified order. The required input failure order is usually read from left to right.
5. Intermediate Event: failure events that are caused by their child nodes. The probability of the intermediate event to occur is denoted by the number in the lower right corner. A top level event (TLE) is a special case of an intermediate event, representing the system hazard.

3 Computing Fault Trees from Counterexamples

Inferring Causality. Fault Trees express causality, in particular they characterize basic events as being causal factors in the occurrence of the top-level event in some Fault Tree. The counterexamples that we use to synthesize these causal relationships, however, merely represent possible executions of the system model, and not explicitly causality amongst event occurrences. Each path in the counterexample is a linearly ordered, interleaved sequence of concurrent events. The question is hence how, and with which justification, we can infer causality from the sets of linearly ordered event sequences that we obtain in the course of the counterexample computation. We use the concept of *structural equations* as proposed by Halpern and Pearl [12] as a model of causality. It is based on *counterfactual* reasoning and the related *alternative world* semantics of Lewis [17, 9]. The counterfactual argument is widely used as the foundation for identifying faults in program debugging [22] and also underlies the formal fault tree semantics proposed in [20]. The "naive" counterfactual causality criterion according to Lewis is as follows: event A is causal for the occurrence of event B if and only if, were A not to happen, B would not occur. The testing of this condition hinges upon the availability of alternative worlds. A causality can be inferred if there is

a world in which A and B occur, whereas in an alternative world neither A nor B occurs. The naive interpretation of the Lewis counterfactual test, however, leads to a number of inadequate or even fallacious inferences of causes, in particular if causes are given by combinations of multiple events. The problematic issues include common or hidden causes, the disjunction and conjunction of causal events, the non-occurrence of events, and the preemption of failure causes due to, e.g., repair mechanisms. A detailed discussion of these issues is beyond the scope of this paper, and we refer to the critical literature on counterfactual reasoning, e.g., [9]. Since we are considering concurrent systems in which particular event interleavings may be the cause of errors, e.g., race conditions, the order of occurrence of events is an potential causal factor that cannot be disregarded. Consider a railroad crossing model in which G denotes the gate closing, O the gate opening, T the train crossing the road, and C the car crossing the tracks. A naive counterfactual test will fail to show that the event sequence $\langle G, O, T, C \rangle$ is a potential cause of a hazard, whereas $\langle G, T, O, C \rangle$ is not. In addition, the naive counterfactual test may determine irrelevant causal events. For instance, the fact that the train engineer union has decided not to call for a strike is not to be considered a cause for the occurrence of an accident at the railroad crossing. Halpern and Pearl extend the Lewis counterfactual model in [12] to what they refer to as *structural equation model* (SEM). It encompasses the notion of causes being logical combinations of events as well as a distinction of relevant and irrelevant causes. However, the structural equation model does not account for event orderings, which is a major concern of this paper. We now sketch an actual cause definition adopted from [12]. An actual cause is a cause in which irrelevant events are factored out. A causal formula in the SEM is a boolean conjunction ψ of variables representing the occurrence of events. We only consider boolean variables, and the variable associated with an event is true in case that event has occurred. The set of all variables is partitioned into the set U of *exogenous* variables and the set V of *endogenous* variables. Exogenous variables represent facts that we do not consider to be causal factors for the effect that we analyze, even though we need to have a formal representation for them so as to encode the "context" ([12]) in which we perform causal analysis. An example for an exogenous variable is the train engineer union's decision in the above railroad crossing example. Endogenous variables represent all events that we consider to have a meaningful, potentially causal effect. The set $X \subseteq V$ contains all events that we expect jointly to be a candidate cause, and the boolean conjunction of these variables forms a causal formula ψ . The causal process comprises all variables that mediate between X and the effect that ψ is causing. Those variables are not root causes, but they contribute to rippling the causal effect through the system until reaching the final effect. Omitting a complete formalization, we assume that there is an actual world and an alternate world. In the actual world, there is a function val_1 that assigns values to variables. In the alternate world, there is a function val_2 assigning potentially different values to the variables. In the SEM, a formula ψ is an actual cause for an event represented by the formula φ , if the following conditions are met:

AC1: Both ψ and φ are true in the actual world, assuming the context defined by the variables in U , and given a valuation $val_1(V)$.

AC2: The set of endogenous events V is partitioned into sets Z and W , where the events in Z are involved in the causal process and the events in W are not involved in the causal process. It is assumed that $X \subseteq Z$ and that there are valuations $val_2(X)$ and $val_2(W)$ assigning values to the variables in X and W , respectively, such that:

1. Changing the values of the variables in X and W from val_1 to val_2 changes φ from true to false.
2. Setting the values of the variables in W from val_1 to val_2 should have no effect on φ as long as the values of the variables in X are kept at the values defined by val_1 , even if all the variables in an arbitrary subset of $Z \setminus X$ are set to their value according to val_1 .

AC3: The set of variables X is minimal: no subset of X satisfies conditions AC1 and AC2.

AC2(1) corresponds to the Lewis counterfactual test. However, as [12] argue, AC2(1) is too permissive, and AC2(2) constrains what is admitted as cause by AC2(1). Minimality in AC3 ensures that only those elements of the conjunction that are essential for changing φ in AC2(1) are considered part of the cause; inessential elements are pruned.

Formal Representation of Events and their Order. To logically reason about the causality of events in our setting we need to allow for the description of conjunctive and disjunctive occurrence of events and represent, at the same time, the order in which the events occur. In the common description of the structural equation model the occurrence of events is encoded as boolean formulae. In these formulae, boolean variables represent the occurrence of an event (true = event occurred, false = event did not occur). These variables are connected via the boolean and- or-or-operators to express conjunctive or disjunctive constraints on their occurrence. Note that this representation does not yet allow for expressing logical constraints on the order in which events need to occur. We first define a mathematical model that allows us to logically reason about the occurrence of events in sets of execution sequences forming counterexamples in stochastic model checking. Technical Systems evolve in discrete computation steps. A system state s is defining a valuation of the system state variables. In our setting, we limit ourselves to considering systems that only contain Boolean state variables representing the occurrence of events, as described above. We use a set of atomic propositions that represent the Boolean state variables we consider. A computation step is characterized by an instantaneous transition which takes the system from some state s to a successor state s' . The transition from s to s' will be triggered by an action a , corresponding to the occurrence of an event. Since we wish to derive causality information from sets of finite computations, which we obtain by observing a finite number of computation steps, our main interest will be in sets of state-action sequences. We define the following

model as a basis for our later formalization of the logical connection between events.

Definition 1. *State-Action Trace Model.* Let S denote a set of states, AP a finite set of atomic state propositions, and Act a finite set of action names.

- A finite sequence $s_0, a_0, s_1, a_1, \dots, a_{n-1}, s_n$ with, for all i , $s_i \in S$ and $a_i \in Act$, is called a state-action trace over (S, Act) .
- A State-Action Trace Model (SATM) is a tuple $M = (S, Act, AP, L, \Sigma)$ where $\Sigma = \{\sigma_1, \dots, \sigma_k\}$ such that each σ_i is a state-action trace over (S, Act) , and $L : S \rightarrow 2^{AP}$ is a function assigning each state the set of atomic propositions that are true in that state.

We assume that for a given SATM M , AP contains the variables representing the events that we wish to reason about. We also assume that for a given state-action trace σ , $L(s_i)$ contains the event variable corresponding to the action a_{i-1} . Notice that we consider event instances, not types. In other words, the n -th occurrence of some event of type E will be distinct in AP from the $n+1$ st occurrence of this event type. We next define an event order logic allowing us to reason about boolean conditions on the occurrence of events. The logic is using a set \mathcal{A} of event variables as well as the boolean connectives \wedge , \vee and \neg . To express the ordering of events we introduce the ordered conjunction operator \triangleleft . The formula $A \triangleleft B$ is satisfied if and only if events A and B occur in a trace and A occurs before B . The formal semantics of this logic is defined on SATMs. Notice that the \triangleleft operator is a temporal logic operator and that the SATM model is akin to a linearly ordered Kripke structure.

Definition 2. *Semantics of event order logic.* Let $M = (S, Act, AP, L, \Sigma)$ a SATM, ϕ and ψ formulae of the event order logic, and let \mathcal{A} a set of event variables, with $A \in \mathcal{A}$, over which ϕ and ψ are built. Let $\sigma = s_0, a_0, s_1, a_1, \dots, a_{n-1}, s_n$ a state-action trace over (S, Act) . We define that a formula is satisfied in state s_i of σ as follows:

- $s_i \models A$ iff $A \in L(s_i)$.
- $s_i \models \neg\phi$ iff not $s_i \models \phi$.
- $s_i \models \phi \wedge \psi$ iff $\exists j, k : i \leq j, k \leq n . s_j \models \phi$ and $s_k \models \psi$.
- $s_i \models \phi \vee \psi$ iff $\exists j, k : i \leq j, k \leq n . s_j \models \phi$ or $s_k \models \psi$.
- $s_i \models \phi \triangleleft \psi$ iff $\exists j, k : i \leq j \leq k \leq n . s_j \models \phi$ and $s_k \models \psi$.

We define that a sequence σ satisfies a formula ϕ , written as $\sigma \models \phi$ iff $\exists i . s_i \models \phi$. We define that the SATM M satisfies the formula ϕ , written as $M \models \phi$, iff $\exists \sigma \in \Sigma . \sigma \models \phi$.

In order to perform comparison operations between paths we define a number of path comparison operators.

Definition 3. *Path Comparison Operators.* Let $M = (S, Act, AP, L, \Sigma)$ a SATM, and σ_1 and σ_2 state-action traces in M .

- $=$: $\sigma_1 = \sigma_2$ iff $\forall e \in Act . \sigma_1 \models e \equiv \sigma_2 \models e$.
- \doteq : $\sigma_1 \doteq \sigma_2$ iff $\forall e_1, e_2 \in Act . \sigma_1 \models e_1 \wedge e_2 \equiv \sigma_2 \models e_1 \wedge e_2$.
- \sqsubseteq : $\sigma_1 \sqsubseteq \sigma_2$ iff $\forall e \in Act . \sigma_1 \models e \Rightarrow \sigma_2 \models e$. Furthermore, $\sigma_1 \subset \sigma_2$ iff $\sigma_1 \sqsubseteq \sigma_2$ and not $\sigma_1 = \sigma_2$.
- $\dot{\sqsubseteq}$: $\sigma_1 \dot{\sqsubseteq} \sigma_2$ iff $\forall e_1, e_2 \in Act . \sigma_1 \models e_1 \wedge e_2 \Rightarrow \sigma_2 \models e_1 \wedge e_2$. Furthermore, $\sigma_1 \dot{\subset} \sigma_2$ iff $\sigma_1 \dot{\sqsubseteq} \sigma_2$ and not $\sigma_1 \doteq \sigma_2$.

We are now ready to adopt the SEM to event orders. We interpret the SEM equations over a given SATM M . Again, without providing a detailed formalization, we assume the existence of a function $order_1$ assigning an order to the occurrence of the events M in the actual world, as well as a function $order_2$ which assigns a potentially different order in the alternate world. An event order logic formula ψ is considered a cause for an event represented by the event order logic formula φ , if the following conditions are satisfied:

AC1: Both ψ and φ are true in the actual world, assuming the context defined by the variables in U , given a valuation $val_1(V)$ and an order $order_1(V)$.

AC2: The set of endogenous events V is partitioned into sets Z and W , where the events in Z are involved in the causal process and the events in W are not involved in the causal process. It is assumed that $X \subseteq Z$ and that there exist valuations $val_2(X)$ and $val_2(W)$ and orders $order_2(X)$ and $order_2(W)$ such that:

1. Changing the values of the variables in X and W from val_1 to val_2 and the order of the variables in X and W from $order_1$ to $order_2$ changes φ from true to false.
2. Setting the values of the variables in W from val_1 to val_2 and the order of the variables in W from $order_1$ to $order_2$ should have no effect on φ as long as the values of the variables in X are kept at the values defined by val_1 , and the order as defined by $order_1$, even if all the variables in an arbitrary subset of $Z \setminus X$ are set to their value according to val_1 and $order_1$.

AC3: The set of variables X is minimal: no subset of X satisfies conditions AC1 and AC2.

If a formula ψ meets the above described conditions, the occurrence of the events included in ψ is causal for φ . However, condition AC2 does not imply that the order of the occurring events is causal. We introduce the following condition to express that the order of the variables occurring in ψ , or an arbitrary subset of these variables, has an influence on the causality of φ :

OC1: Let $Y \subseteq X$. Changing the order $order_1(Y)$ of the variables in Y to an arbitrary order $order_2(Y)$, while keeping the variables in $X \setminus Y$ at $order_1$, changes φ from true to false.

If for a subset of X OC1 is not satisfied, the order of the events in this subset has no influence on the causality of φ .

Fault Tree Generation. In order to automatically synthesize a fault tree from a stochastic counterexample, the combinations of basic events causing the top level event in the fault tree have to be identified. Using a stochastic model

checker we compute a counterexample which contains all paths leading to a state corresponding to the occurrence of some top level event T . This is achieved by computing the counterexample for the CSL formula $P =?(true U t)$, where t is a state formula representing the top level event T . We interpret counterexamples in the context of an SATM $M = (S, Act, AP, L, \Sigma)$. We assume that Σ is partitioned in disjoint sets Σ_G and Σ_C , where Σ_C contains all traces belonging to the counterexample, whereas Σ_G contains all maximal simple system traces that do not belong to the counterexample. The disjointness of Σ_C and Σ_G implies that M is deterministic with respect to the causality of T . Furthermore, we define $M_C = (S, Act, AP, L, \Sigma_C)$ as the restriction of M to only the counterexample traces, and refer to it as a counterexample model. W.l.o.g. we assume that every trace $\sigma \in M_C$ contains a last transition executing the top level event T , so that its last state $s_n \models T$, which implies that $M_C \models T$. In our interpretation of the SEM, actual world models will be derived from Σ_C , whereas alternate world models are part of Σ_G . Notice that in order to compute the full model probability of reaching T it is necessary to perform a complete state space exploration of the model that we analyze. We hence obtain M_C at no additional cost. We next define the candidate set of paths that we consider to be causal for T . We define this set in such a way that it includes all minimal paths. Paths are minimal if they do not contain a subpath according to the \subseteq operator that is also a member of the candidate set.

Definition 4 (Candidate Set). *Let $M_C = (S, Act, AP, L, \Sigma_C)$ a counterexample model, and T a top level event in M_C . We define the candidate set of paths belonging to the fault tree of T as $CFT(T)$:*

$$CFT(T) = \{\sigma \in \Sigma_C | \forall \sigma' \in \Sigma_C . \sigma' \subseteq \sigma \Rightarrow \sigma' = \sigma\}. \quad (1)$$

Notice that the candidate set is minimal in the sense that removing an event from some path in the candidate set means that the resulting path is no longer in the counterexample Σ_C . Given a counterexample model M_C , we state the following observations regarding the paths included in Σ_C :

- Each $\sigma \in \Sigma_C$ can be viewed as an ordered conjunction $A_1 \wedge \dots \wedge A_{n-1} \wedge T$ of events, where T is the top level event that we consider.
- On each path in the counterexample, there has to be at least one event causing the top level event. If that was not the case, the top level event would not have occurred on that path and as a result the path would not be in the counterexample.

The algorithm that we propose to compute fault trees is an over-approximation of the computation of the causal events X since computing the set X precisely is not efficiently possible [11]. Instead of computing the set X of events that are causal for some φ , we compute the set Z of events, which consists of all events that are part of the causal process of φ . Z will then be represented by ψ . Since X is a subset of Z we can assure that no event that is causal is omitted from

the fault tree. It is, however, possible that due to our overapproximation events that are not in X are added to the fault tree. This applies in particular to those events that are part of the causal process, and hence mediate between X and φ . However, as we show in Section 4, adding such events can be helpful to illustrate how the root cause is indirectly propagating by non-causal events to finally cause the top level event. We do not account for exogenous variables, since we believe them to be less relevant in the analysis of models of computational systems since the facts represented in those models all seem to be endogenous facts of the system. However, should one wish to consider exogenous variables, those can easily be retrofitted. We now define tests that will identify the set Z of ordered conjunctions of events that satisfy the conditions AC1 to AC3 and OC1, and which hence can be viewed as part of the causal process of the top level event. The starting point for this computation is the candidate set of Definition 4.

Test for AC1: The actual causal set Z that our algorithm computes is a subset of the events included in the candidate set $CFT(T)$ for some given top level event T . Since we assume that every path includes at least one causal event, Z is not empty. We may hence conclude that $CFT(T) \models \psi$ and $CFT(T) \models \varphi$.

Test for AC2(1): We check for each path $\sigma \in CFT(T)$ whether the ordered conjunctions of events that it is representing fulfills the condition AC2(1). We assume that the set of events Z is equal to the events occurring on the path σ . We further assume that $W = V \setminus Z$ and that $V = Act$. W hence contains all events that are possible, minus the ones that occur on path σ . More formally, for a given σ , $Z = \{e \in V \mid \sigma \models e\}$. This corresponds to using the valuation val_1 to assign true to all variables in Z and false to all variables in W in the formulation of AC2(1). Changing the valuation of the variables in Z to move from val_1 to some val_2 can imply removing variables from Z . Due to the minimality of σ this implies that the resulting trace σ' is no longer in Σ_C . Testing of this condition is hence implicit and implied by the minimality of the candidate set.

Test for AC2(2): We need to check that moving W from val_1 to val_2 and from $order_1$ to $order_2$ has no effect on the outcome of φ as long as the values of X are kept at the values defined by val_1 and the order defined by $order_1$. Recall that W denotes all events that are not currently considered to be causal, and that we compute Z as an overapproximation of X . For a path $\sigma \in CFT(T)$ changing W from val_1 to val_2 and from $order_1$ to $order_2$ implies that events are added to σ . Thus, we check for each path $\sigma \in CFT(T)$ whether there exists some path $\sigma' \in \Sigma_C$ for which $\sigma \dot{c} \sigma'$ holds. If there is no such path, there are no val_2 and $order_2$ of W that change the outcome of φ , and as a consequence AC2(2) is fulfilled by σ . If we do find such a path σ' , it contains all variables in Z with val_1 and $order_1$ and some events W with val_2 and $order_2$ that change the outcome of φ . In other words, the non-occurrence of the events in W on σ was causal for φ . In order to identify those events, we search for the minimal paths $R = \{\sigma' \in \Sigma_C \mid \sigma \dot{c} \sigma'\}$. For each path in R we negate the events that are in σ but not in σ' and add them to the candidate set. Subsequently, we remove σ from the candidate set. Consider the case $Z = G \wedge O \wedge T \wedge C$ in our rail road crossing model. It is necessary that no event G occurs between O and T for

this ordered conjunction of events to lead to a hazard. If the system execution $G \wedge O \wedge G \wedge T \wedge C$ is possible, which means that there is a path representing this execution in the set $NCX(A)$ for top level event A , we hence have to replace Z by $Z' = G \wedge O \wedge \neg G \wedge T \wedge C$.

Test for AC3: Due to the minimality property of the candidate set, no test for AC3 is necessary.

Test for OC1: We need to decide whether for all ordered conjunctions in $CFT(T)$ the order of the events is relevant to cause T . For each path $\sigma \in CFT(T)$, we check whether the order of the events to occur is important or not. If the order of events in σ is not important, then there has to be at least one path $\sigma' \in CFT(T)$ for which $\sigma = \sigma'$ and not $\sigma \doteq \sigma'$. For each event e_i in σ we check for all other events e_j with $i < j$ whether $\sigma' \models e_i \wedge e_j$ for all $\sigma' \in CFT(T)$. If $\sigma' \models e_i \wedge e_j$ for all $\sigma' \in CFT(T)$, we mark this pair of events as having an order which is important for causality. If we can not find such a σ' , we mark the whole path σ as having an order which is important for causality.

Adding Probabilities. In order to properly compute the probability mass that is to be attributed to the TLE T in the fault tree it is necessary to account for all paths that can cause T . If there are two paths $\sigma_1, \sigma_2 \in \Sigma_C$ which, when combined, deliver a path $\sigma_{12} \in \Sigma_C$, then the probability mass of all three paths needs to be taken into account when calculating the probability for reaching T . To illustrate this point, consider an extension of the railroad example introduced above. We add a traffic light indicating to the car driver that a train is approaching. Event R indicates that the traffic light on the crossing is red, while the red light being off is denoted by event L . The top level event A denoting the hazard is expressed as a state proposition applicable to the underlying stochastic model that states that neither the red light is on nor the gate is closed, and that the train approaches and the car is in the crossing. Assume that the above described test would identify the following ordered conjunctions of events to be causal: $\neg G \wedge T \wedge C$ and $\neg L \wedge T \wedge C$. Due to the minimality property of $CFT(A)$ the ordered conjunction $\neg G \wedge \neg L \wedge T \wedge C$ would be missing. We would hence lose the probability mass associated with the corresponding trace in the counterexample, as well as the qualitative information that the simultaneous failure of the red light and the gate also leads to a hazardous state. To account for this situation we introduce the combination condition CC1. CC1: Let $\sigma_1, \sigma_2, \dots, \sigma_k \in CFT(L)$ paths and $\psi_1, \psi_2, \dots, \psi_k$ the event conjunctions representing them. A path σ is added to $CFT(L)$ if for $k \geq 2$ paths in $CFT(L)$ it holds that $\sigma \models \psi_1 \wedge \sigma \models \psi_2 \wedge \dots \wedge \sigma \models \psi_k$. We can now assign each path in the candidate set the sum of the probability masses of the paths that it represents. This is done as follows: The probability of a path σ_1 in $CFT(L)$ is the probability sum of all paths σ' for which σ_1 is the only subset in $CFT(L)$. The last condition is necessary in order to correctly assign the probabilities to paths which were added to the fault tree by test CC1. All paths still in the candidate set are part of the fault tree and have now to be included in the fault tree representation. The fault trees generated by our approach all have a normal form, that is they start with an *intermediate*-gate representing the top level event, that is connected to an *OR*-gate. The paths in the candidate set

$CFT(L)$ will then be added as child nodes to the OR -gate as follows: Paths with a length of one and hence consisting of only one basic event are represented by the respective basic event. A path with length greater than one that has no subset of labels marked as ordered is represented by an AND -gate. This AND -gate connects the basic events belonging to that path. If a (subset of a) path is marked as ordered it is represented by a $PAND$ -gate that connects the basic events in addition with an *Order Condition* connected to the $PAND$ -gate constraining the order of the elements. The probability values of the AND -gates are the corresponding probabilities of the paths that they represent. In order to display the probabilities in the graphical representation of the fault tree, we add an intermediate event as parent node for each AND -gate. The resulting intermediate events are then connected by an OR -gate that leads to the top event, representing the hazard. Since the path probabilities are calculated for a path starting from an initial state to the hazard state, the probability of the OR -gate is the sum of the probability of all child elements.

Scalability and Complexity. As we show in detail in [15], the complexity of our algorithm is cubic in the size of the counterexample. The case studies presented in Section 4 show that the fault tree computation finishes in several seconds, while the computation of the counterexample took several minutes. Hence, the limiting factor of our approach is the time needed for the computation of the counterexample.

4 Case Study

We now briefly discuss a case study illustrating the applicability of our approach. Notice that for reasons of limited space we have to refer the reader to [15] for more detail, and also for more case studies. This case study is taken from [1] and models an industrial size airbag system. The airbag system architecture that we consider consists of two acceleration sensors whose task it is to detect front or rear crashes, one microcontroller to perform the crash evaluation, and an actuator that controls the deployment of the airbag. Although airbags save lives in crash situations, they may cause fatal behavior if they are inadvertently deployed. This is because the driver may lose control of the car when this deployment occurs. It is therefore a pivotal safety requirement that an airbag is never deployed if there is no crash situation. We are interested in generating the fault tree for an inadvertent ignition of the airbag. In CSL, this property can be expressed using the formula $\mathcal{P}_{=?}(noCrash \ U^{\leq T} \ AirbagIgnited)$. Notice that we assume the PRISM models in practical usage scenarios to be automatically synthesized from higher-level design models, such as for instance by our QuantUM tool [16]. However, the case study presented in this paper was directly modeled in the PRISM language. We computed the counterexamples using our counterexample generation tool DiPro [2], which in turn uses the PRISM model checker. Figure 1 shows the fault tree generated by CX2FTA. For better readability we have omitted the order constraints of the $PAND$ -gates. While the counterexample consists of 738 paths, the fault tree comprises only 5 paths. It is easy to see




Fig. 1. Fault Tree of the Airbag System

by which basic events, and with which probabilities, an inadvertent deployment of the airbag is caused. For instance, there is only one single fault that can lead to an inadvertent deployment, namely *FASICShortage*. It is also easy to see that the combination of the basic events *FETStuckHigh* and *FASICStuckHigh* only lead to an inadvertent deployment of the airbag if the basic event *FETStuckHigh* occurs prior to the basic event *FASICStuckHigh*. The case study shows that the fault tree is a compact and concise visualization of the counterexample which allows for an easy identification of the basic events that cause the inadvertent deployment of the airbag, along with the corresponding probabilities. If the order of the events is important, this can be seen in the fault tree by the *PAND*-gate. In the counterexample computed by DiPro one would have to manually compare the order of the events in all 738 paths, which is a tedious and time consuming task. Figure 2 shows the memory and run time consumption of the fault counterexample and fault tree computation⁴. The computational effort is dominated by the counterexample computation. Increasing the parameter t (mission time) in the process model has only a marginal influence on the computational effort needed.

5 Related Work

Work described in [7, 20] interprets fault trees in terms of temporal logic. This is the opposite direction of what we aim to accomplish, namely to derive fault trees

⁴ Experiments were performed on a PC with an Intel Core2Duo CPU (3.06 Ghz) and 8 GBs RAM.

t	Runtime CX (sec.)	Paths in CX	Mem. CX	Runtime FT	Paths in FT	Mem. FT
10	1 147 (\approx 19.12 min.)	738	29.17 MB	1.3 (sec.)	5	27 MB
100	1 148 (\approx 19.13 min.)	738	29.20 MB	1.3 (sec.)	5	27 MB
1000	1 263 (\approx 21.05 min.)	738	29.49 MB	1.8 (sec.)	5	27 MB

Fig. 2. Experiment results for T=10, T=100 and T=1000.

from system execution models. Various approaches to derive fault trees semi-automatically or automatically from various semi-formal or formal models have been proposed, e.g. [19, 8, 18]. Contrary to our method, none of these methods uses sets of system execution sequences as the basis of the fault tree derivation, or provides an automatic probabilistic assessment of the synthesized fault tree nodes. These approaches also lack a justification of the causality model used. Our work extends and improves on the approach of [6] in the following ways: We use a single set of system modeling and specification languages, namely PRISM and CSL. Whereas in the approach of [6] only minimal cut sets are generated, we generate complete fault trees. Contrary to [6], we support PAND-gates and provide a justification for the causality model used. Work documented in [5] uses the Halpern and Pearl approach to determine causality for counterexamples in functional CTL model checking. However, this approach considers only functional counterexamples that consist of single execution sequences. [11] contains a careful analysis of the complexity of computing causality in the SEM. Most notable is the result that even for an SEM with only binary variables computing causal relationships between variables is NP-complete.

6 Conclusion

We presented a method and tool that automatically generates a fault tree from a probabilistic counterexample. We demonstrated that our approach improves and facilitates the analysis of safety critical systems. The resulting fault trees were significantly smaller and hence easier to understand than the corresponding stochastic counterexample, but still contain all information to discern the causes for the occurrence of a hazard. The justification for the causalities determined by our method are based on an adoption of the Structural Equation Model of Halpern and Pearl. We illustrated how to use this model in the analysis of computing systems and extended it to account for event orderings as causal factors. We presented an over-approximating implementation of the causality tests derived from this model. To the best of our knowledge this is the first attempt at using the structural equation model in this fashion. In future work, we plan to further extend our approach, in particular to support the generation of dynamic fault-trees [10]. We are also interested in incorporating causality analysis directly into model checking algorithms.

Acknowledgments. The authors thank Mark Burgess for giving them access to the FaultCAT source code.

References

1. H. Aljazzar, M. Fischer, L. Grunske, M. Kuntz, F. Leitner-Fischer, and S. Leue. Safety Analysis of an Airbag System Using Probabilistic FMEA and Probabilistic Counterexamples. In *Proc. of QEST 2009*. IEEE Computer Society, 2009.
2. H. Aljazzar and S. Leue. Debugging of Dependability Models Using Interactive Visualization of Counterexamples. In *Proc. of QEST 2008*. IEEE Computer Society, 2008.
3. H. Aljazzar and S. Leue. Directed explicit state-space search in the generation of counterexamples for stochastic model checking. *IEEE Trans. Soft. Eng.*, 2009.
4. C. Baier, B. Haverkort, H. Hermanns, and J.-P. Katoen. Model-checking algorithms for continuous-time Markov chains. *IEEE Trans. Soft. Eng.*, 2003.
5. I. Beer, S. Ben-David, H. Chockler, A. Orni, and R. Trefer. Explaining counterexamples using causality. In *Proc. of CAV 2009*, LNCS. Springer, 2009.
6. E. Böde, T. Peikenkamp, J. Rakow, and S. Wischmeyer. Model Based Importance Analysis for Minimal Cut Sets. In *Proc. of ATVA 2008*, volume 5311 of LNCS. Springer, 2008.
7. M. Bozzano, A. Cimatti, and F. Tapparo. Symbolic Fault Tree Analysis for Reactive Systems. In *Proc. of ATVA 2007*, volume 4762 of LNCS. Springer, 2007.
8. B. Chen, G. Avrunin, L. Clarke, and L. Osterweil. Automatic Fault Tree Derivation From Little-Jil Process Definitions. In *Proc. of SPW/ProSim 2006*, volume 3966 of LNCS. Springer, 2006.
9. J. Collins, editor. *Causation and Counterfactuals*. MIT Press, 2004.
10. J. Dugan, S. Bavuso, and M. Boyd. Dynamic Fault Tree Models for Fault Tolerant Computer Systems. *IEEE Trans. Reliability*, 1992.
11. T. Eiter and T. Lukasiewicz. Complexity results for structure-based causality. *Artificial Intelligence*, 2002.
12. J. Halpern and J. Pearl. Causes and explanations: A structural-model approach. Part I: Causes. *The British Journal for the Philosophy of Science*, 2005.
13. T. Han, J.-P. Katoen, and B. Damman. Counterexample generation in probabilistic model checking. *IEEE Trans. Softw. Eng.*, 2009.
14. A. Hinton, M. Kwiatkowska, G. Norman, and D. Parker. PRISM: A Tool for Automatic Verification of Probabilistic Systems. In *Proc. of TACAS 2006*, volume 3966 of LNCS. Springer, 2006.
15. M. Kuntz, F. Leitner-Fischer, and S. Leue. From probabilistic counterexamples via causality to fault trees. Technical Report soft-11-02, Chair for Software Engineering, University of Konstanz, 2011. Available from <http://www.inf.uni-konstanz.de/soft/research/publications/pdf/soft-11-02.pdf>.
16. F. Leitner-Fischer and S. Leue. QuantUM: Quantitative safety analysis of UML models. In *In Proc. of QAPL 2011*, 2011.
17. D. Lewis. *Counterfactuals*. Wiley-Blackwell, 2001.
18. M. McKelvin Jr, G. Eirea, C. Pinello, S. Kanajan, and A. Sangiovanni-Vincentelli. A Formal Approach to Fault Tree Synthesis for the Analysis of Distributed Fault Tolerant Systems. In *Proc. of EMSOFT 2005*. ACM, 2005.
19. G. Pai and J. Dugan. Automatic synthesis of dynamic fault trees from UML system models. In *Proc. of ISSRE 2002*. IEEE Computer Society, 2002.
20. G. Schellhorn, A. Thums, and W. Reif. Formal fault tree semantics. In *Proc. IDPT 2002*. Society for Design and Process Science, 2002.
21. U.S. Nuclear Regulatory Commission. *Fault Tree Handbook*, 1981. NUREG-0492.
22. A. Zeller. *Why Programs Fail: A Guide to Systematic Debugging*. Elsevier, 2009.