

Significant Diagnostic Counterexamples in Probabilistic Model Checking

Miguel E. Andrés¹ *, Pedro D’Argenio² **, Peter van Rossum¹

¹Institute for Computing and Information Sciences, The Netherlands.
 {mandres,petervr}@cs.ru.nl

²FaMAF, Universidad Nacional de Córdoba, CONICET, Argentina.
 dargenio@famaf.unc.edu.ar

Abstract. This paper presents a novel technique for counterexample generation in probabilistic model checking of Markov Chains and Markov Decision Processes. (Finite) paths in counterexamples are grouped together in witnesses that are likely to provide similar debugging information to the user. We list five properties that witnesses should satisfy in order to be useful as debugging aid: similarity, accuracy, originality, significance, and finiteness. Our witnesses contain paths that behave similar outside strongly connected components. This paper shows how to compute these witnesses by reducing the problem of generating counterexamples for general properties over Markov Decision Processes, in several steps, to the easy problem of generating counterexamples for reachability properties over acyclic Markov Chains.

1 Introduction

Model checking is an automated technique that, given a finite-state model of a system and a property stated in an appropriate logical formalism, systematically checks the validity of this property. Model checking is a general approach and is applied in areas like hardware verification and software engineering.

Nowadays, the interaction geometry of distributed systems and network protocols calls for probabilistic, or more generally, quantitative estimates of, e.g., performance and cost measures. Randomized algorithms are increasingly utilized to achieve high performance at the cost of obtaining correct answers only with high probability. For all this, there is a wide range of models and applications in computer science requiring quantitative analysis. Probabilistic model checking allow us to check whether or not a probabilistic property is satisfied in a given model, e.g., “Is every message sent successfully received with probability greater or equal than 0.99?”.

A major strength of model checking is the possibility of generating diagnostic information in case the property is violated. This diagnostic information is provided through a *counterexample* showing an execution of the model that invalidates the property under verification. Apart from the immediate feedback in model checking, counterexamples are also used in abstraction-refinement techniques [CGJ⁺00], and provide the foundations for schedule derivation (see, e.g., [BLR05]).

Although counterexample generation was studied from the very beginning in most model checking techniques, this has not been the case for probabilistic model checking. Only recently attention was drawn to this subject [AHL05,AL06,HK07a,HK07b,AL07], fifteen years after the first studies on probabilistic model checking. Contrarily to other model checking techniques, counterexamples in this setting are *not* given by a single execution path. Instead, they are *sets of executions* of the system satisfying a certain undesired property whose probability mass is higher than a given bound. Since counterexamples are used as a diagnostic tool, previous works on counterexamples have presented them as sets of finite paths of large enough probability. We refer to these sets as

* Supported by NWO project 612.000.526

** Supported by the ANPCyT project PICT 26135 and CONICET project PIP 6391

representative counterexamples. Elements of representative counterexamples with high probability have been considered the most informative since they contribute mostly to the property refutation.

A challenge in counterexample generation for probabilistic model checking is that (1) representative counterexamples are very large (often infinite), (2) many of its elements have very low probability, and (3) that elements can be extremely similar to each other (consequently providing similar diagnostic information). Even worse, (4) sometimes the finite paths with highest probability do not indicate the most likely violation of the property under consideration.

For example, look at the Markov chain \mathcal{M} in Figure 1. The property $\mathcal{M} \models_{\leq 0.5} \diamond\psi$ stating that execution reaches a state satisfying ψ (i.e., reaches s_3 or s_4) with probability lower or equal than 0.5 is violated (since the probability of reaching ψ is 1). The left hand side of table in Figure 2 lists finite paths reaching ψ ranked according to their probability. Note that finite paths with highest probability take the left branch in the system, whereas the right branch in itself has higher probability, illustrating Problem 4. To adjust the model so that it does satisfy the property (bug fixing), it is not sufficient to modify the left hand side of the system alone; no matter how one changes the left hand side, the probability of reaching ψ remains at least 0.6. Furthermore, the first six finite paths provide similar diagnostic information: they just make extra loops in s_1 . This is an example of Problem 3. Also, the probability of every single finite path is far below the bound 0.5, making it unclear if a particular path is important; see Problem 2 above. Finally, the (unique) counterexample for the property $\mathcal{M} \models_{< 1} \diamond\psi$ consists of infinitely many finite paths (namely all finite paths of \mathcal{M}); see Problem 1. To overcome these problems, we partition a representative counterexample into sets of




Fig. 1: Markov chain

Rank	Single paths		Witnesses	
	F. Path	Prob	Witness	Mass
1	$s_0(s_1)^1s_3$	0.2	$[s_0s_2s_4]$	0.6
2	$s_0(s_1)^2s_3$	0.1	$[s_0s_1s_3]$	0.4
3	$s_0(s_1)^3s_3$	0.05		
4	$s_0(s_1)^4s_3$	0.025		
5	$s_0(s_1)^5s_3$	0.0125		
6	$s_0(s_1)^6s_3$	0.00625		
7	$s_0(s_2)^1s_4$	0.006		
8	$s_0(s_2)^2s_4$	0.0059		
9	$s_0(s_2)^3s_4$	0.0058		
\vdots	\vdots	\vdots		

Fig. 2: Comparison Table

finite paths that follow a similar pattern. We call these sets *witnesses*. To ensure that witnesses provide valuable diagnostic information, we desire that the set of witnesses that form a counterexample satisfies several properties: two different witnesses should provide different diagnostic information (solving Problem 3) and elements of a single witness should provide similar diagnostic information, as a consequence witnesses have a high probability mass (solving Problems 2 and 4), and the number of witnesses of a representative counterexample should be finite (solving Problem 1).

In our setting, witnesses consist of paths that behave the same outside strongly connected components. In the example of Figure 1, there are two witnesses: the set of all finite paths going right, represented by $[s_0s_2s_4]$ whose probability (mass) is 0.6, and the set of all finite paths going left, represented by $[s_0s_1s_3]$ with probability (mass) 0.4.

In this paper, we show how to obtain such sets of witnesses for bounded probabilistic LTL properties on Markov decision processes (MDP). In fact, we first show how to reduce this problem to finding witnesses for upper bounded probabilistic reachability

properties on discrete time Markov chains (MCs). The major technical matters lie on this last problem to which most of the paper is devoted.

In a nutshell, the process to find witnesses for the violation of $\mathcal{M} \models_{\leq p} \diamond\psi$, with \mathcal{M} being a MC, is as follows. We first eliminate from the original MC all the “uninteresting” parts. This proceeds as the first steps of the model checking process: make absorbing all state satisfying ψ , and all states that cannot reach ψ , obtaining a new MC \mathcal{M}_ψ . Next reduce this last MC to an acyclic MC $\text{Ac}(\mathcal{M}_\psi)$ in which all strongly connected components have been conveniently abstracted with a single probabilistic transition. The original and the acyclic MCs are related by a mapping that, to each finite path in $\text{Ac}(\mathcal{M}_\psi)$ (that we call rail), assigns a set of finite paths behaving similarly in \mathcal{M} (that we call torrent). This map preserves the probability of reaching ψ and hence relates counterexamples in $\text{Ac}(\mathcal{M}_\psi)$ to counterexamples in \mathcal{M} . Finally, counterexamples in $\text{Ac}(\mathcal{M}_\psi)$ are computed by reducing the problem to a k shortest path problem, as in [HK07a]. Because $\text{Ac}(\mathcal{M}_\psi)$ is acyclic, the complexity is lower than the corresponding problem in [HK07a].

It is worth to mention that our technique can also be applied to simple pCTL formulas without nested path quantifiers.

Organization of the paper. Section 2 presents the necessary background on Markov chains (MC), Markov Decision Processes (MDP), and Linear Temporal Logic (LTL). Section 3 presents the definition of counterexamples and discuss the reduction from general LTL formulas to upper bounded probabilistic reachability properties, and the extraction of the maximizing MC in a MDP. Section 4 discusses desire properties of counterexamples. In Sections 5 and 6, we introduce the fundamentals on rails and torrents, the reduction of the original MC to the acyclic one, and our notion of significant diagnostic counterexamples. Section 7 then present the techniques to actually compute counterexamples. In Section 8 we discuss related work and give final conclusions.

2 Preliminaries

2.1 Markov Decision Processes and Markov chains

Markov Decision Processes (MDPs) constitute a formalism that combines nondeterministic and probabilistic choices. They are the dominant model in corporate finance, supply chain optimization, system verification and optimization. There are many slightly different variants of this formalism such as action-labeled MDPs [Bel57,FV97], probabilistic automata [SL95,SdV04]; we work with the state-labeled MDPs from [BdA95].

Definition 2.1. Let S be a set. A *discrete probability distribution* on S is a function $p: S \rightarrow [0, 1]$ with countable or finite carrier and such that $\sum_{s \in S} p(s) = 1$. We denote the set of all discrete probability distributions on S by $\text{Distr}(S)$. Additionally, we define the *Dirac distribution* on an element $s \in S$ as 1_s , i.e., $1_s(s) = 1$ and $1_s(t) = 0$ for all $t \in S \setminus \{s\}$.

Definition 2.2. A *Markov Decision Process* (MDP) is a four-tuple $\mathcal{D} = (S, s_0, L, \tau)$, where

- S is the finite state space of the system;
- $s_0 \in S$ is the initial state;
- L is a labeling function that associates to each state $s \in S$ a set $L(s)$ of propositional variables that are *valid* in s ;
- $\tau: S \rightarrow \wp(\text{Distr}(S))$ is a function that associates to each $s \in S$ a non-empty and finite subset of $\text{Distr}(S)$ of probability distributions.

Definition 2.3. Let $\mathcal{D} = (S, s_0, \tau, L)$ be a MDP. We define a *successor* relation $\delta \subseteq S \times S$ by $\delta \triangleq \{(s, t) \mid \exists \pi \in \tau(s) . \pi(t) > 0\}$ and for each state $s \in S$ we define the sets

$$\begin{aligned} \text{Paths}(\mathcal{D}, s) &\triangleq \{s_0 s_1 s_2 \dots \in S^\omega \mid s_0 = s \wedge \forall n \in \mathbb{N} . \delta(s_n, s_{n+1})\} \text{ and} \\ \text{Paths}^*(\mathcal{D}, s) &\triangleq \{s_0 s_1 \dots s_n \in S^* \mid s_0 = s \wedge \forall 0 \leq i < n . \delta(s_i, s_{i+1})\} \end{aligned}$$

of paths and finite paths respectively beginning at s . We usually omit \mathcal{D} from the notation; we also abbreviate $\text{Paths}(\mathcal{D}, s_0)$ as $\text{Paths}(\mathcal{D})$ and $\text{Paths}^*(\mathcal{D}, s_0)$ as $\text{Paths}^*(\mathcal{D})$. For $\omega \in \text{Paths}(s)$, we write the $(n+1)$ -st state of ω as ω_n . As usual, we let $\mathcal{B}_s \subseteq \wp(\text{Paths}(s))$ be the Borel σ -algebra on the cones $\langle s_0 \dots s_n \rangle \triangleq \{\omega \in \text{Paths}(s) \mid \omega_0 = s_0 \wedge \dots \wedge \omega_n = s_n\}$. Additionally, for a set of finite paths $A \subseteq \text{Paths}^*(s)$, we define $\langle A \rangle \triangleq \bigcup_{\sigma \in A} \langle \sigma \rangle$.




Fig. 3: Markov Decision Process

Figure 3 shows a MDP. Absorbing states (i.e., states s with $\tau(s) = \{1_s\}$) are represented by double lines. This MDP features a single nondeterministic decision, to be made in state s_0 , namely π_1 and π_2 .

Definition 2.4. Let $\mathcal{D} = (S, s_0, \tau, L)$ be a MDP and $\mathcal{A} \subseteq S$. We define the sets of paths and finite paths reaching \mathcal{A} as

$$\text{Reach}(\mathcal{D}, s, \mathcal{A}) \triangleq \{\omega \in \text{Paths}(\mathcal{D}, s) \mid \exists_{i \geq 0} \omega_i \in \mathcal{A}\} \text{ and}$$

$$\text{Reach}^*(\mathcal{D}, s, \mathcal{A}) \triangleq \{\sigma \in \text{Paths}^*(\mathcal{D}, s) \mid \text{last}(\sigma) \in \mathcal{A} \wedge \forall_{i \leq |\sigma|-1} \sigma_i \notin \mathcal{A}\}$$

respectively. Note that $\text{Reach}^*(\mathcal{D}, s, \mathcal{A})$ consists of those finite paths σ reaching \mathcal{A} exactly once, at the end of the execution. It is easy to check that these sets are *prefix free*, i.e. contain finite paths such that none of them is a prefix of another one.

2.2 Schedulers

Schedulers (also called strategies, adversaries, or policies) resolve the nondeterministic choices in a MDP [PZ93, Var85, BdA95].

Definition 2.5. Let $\mathcal{D} = (S, s_0, \tau, L)$ be a MDP. A *scheduler* η on \mathcal{D} is a function from $\text{Paths}^*(\mathcal{D})$ to $\text{Distr}(\wp(\text{Distr}(S)))$ such that for all $\sigma \in \text{Paths}^*(\mathcal{D})$ we have $\eta(\sigma) \in \text{Distr}(\tau(\text{last}(\sigma)))$. We denote the set of all schedulers on \mathcal{D} by $\text{Sch}(\mathcal{D})$.

Note that our schedulers are randomized, i.e., in a finite path σ a scheduler chooses an element of $\tau(\text{last}(\sigma))$ probabilistically. Under a scheduler η , the probability that the next state reached after the path σ is t , equals $\sum_{\pi \in \tau(\text{last}(\sigma))} \eta(\sigma)(\pi) \cdot \pi(t)$. In this way, a scheduler induces a probability measure on \mathcal{B}_s as usual.

Definition 2.6. Let \mathcal{D} be a MDP, $s \in S$, and η an s -scheduler on \mathcal{D} . We define the probability measure $\mu_{s, \eta}$ as the unique measure on \mathcal{B}_s such that for all $s_0 s_1 \dots s_n \in \text{Paths}^*(s)$

$$\Pr_{s, \eta}(\langle s_0 s_1 \dots s_n \rangle) = \prod_{i=0}^{n-1} \sum_{\pi \in \tau(s_i)} \eta(s_0 s_1 \dots s_i)(\pi) \cdot \pi(s_{i+1}).$$

We now recall the notions of deterministic and memoryless schedulers.

Definition 2.7. Let \mathcal{D} be a MDP, $s \in S$, and η an scheduler of \mathcal{D} . We say that η is *deterministic* if $\eta(\sigma)(\pi_i)$ is either 0 or 1 for all $\pi_i \in \tau(\text{last}(\sigma))$ and all $\sigma \in \text{Paths}^*(\mathcal{D})$. We say that a scheduler is *memoryless* if for all finite paths σ_1, σ_2 of \mathcal{D} with $\text{last}(\sigma_1) = \text{last}(\sigma_2)$ we have $\eta(\sigma_1) = \eta(\sigma_2)$

Definition 2.8. Let \mathcal{D} be a MDP, $s \in S$, and $\Delta \in \mathcal{B}_s$. Then the *maximal and minimal probabilities of Δ* , $\mathbf{Pr}_s^+(\Delta), \mathbf{Pr}_s^-(\Delta)$, are defined by

$$\mathbf{Pr}_s^+(\Delta) \triangleq \sup_{\eta \in \text{Sch}_s(\mathcal{D})} \mathbf{Pr}_{s,\eta}(\Delta) \quad \text{and} \quad \mathbf{Pr}_s^-(\Delta) \triangleq \inf_{\eta \in \text{Sch}_s(\mathcal{D})} \mathbf{Pr}_{s,\eta}(\Delta).$$

A scheduler that attains $\mathbf{Pr}_s^+(\Delta)$ or $\mathbf{Pr}_s^-(\Delta)$ is called a *maximizing* or *minimizing* scheduler respectively.

A *Markov chain* (MC) is a MDP associating exactly one probability distribution to each state. In this way nondeterministic choices are not longer allowed.

Definition 2.9 (Markov chain). Let $\mathcal{D} = (S, s_0, \tau, L)$ be a MDP. If $|\tau(s)| = 1$ for all $s \in S$, then we say that \mathcal{D} is a *Markov chain* (MC).

2.3 Linear Temporal Logic

Linear temporal logic (LTL) [MP91] is a modal temporal logic with modalities referring to time. In LTL is possible to encode formulas about the future of paths: a condition will eventually be true, a condition will be true until another fact becomes true, etc.

Definition 2.10. LTL is built up from the set of propositional variables \mathcal{V} , the logical connectives \neg, \wedge , and a temporal modal operator by the following grammar:

$$\phi ::= \mathcal{V} \mid \neg\phi \mid \phi \wedge \phi \mid \phi \mathcal{U} \phi.$$

Using these operators we define $\vee, \rightarrow, \diamond,$ and \square in the standard way.

Definition 2.11. Let $\mathcal{D} = (S, s_0, \tau, L)$ be a MDP. We define satisfiability for paths ω in \mathcal{D} and LTL formulas ϕ, ψ inductively by

$$\begin{aligned} \omega \models_{\mathcal{D}} v &\Leftrightarrow v \in L(\omega_0) \\ \omega \models_{\mathcal{D}} \neg\phi &\Leftrightarrow \text{not}(\omega \models_{\mathcal{D}} \phi) \\ \omega \models_{\mathcal{D}} \phi \wedge \psi &\Leftrightarrow \omega \models_{\mathcal{D}} \phi \text{ and } \omega \models_{\mathcal{D}} \psi \\ \omega \models_{\mathcal{D}} \phi \mathcal{U} \psi &\Leftrightarrow \exists_{i \geq 0} \omega_{\downarrow i} \models_{\mathcal{D}} \psi \text{ and } \forall_{0 \leq j < i} \omega_{\downarrow j} \models_{\mathcal{D}} \phi \end{aligned}$$

where $\omega_{\downarrow i}$ is the i -th suffix of ω . When confusion is unlikely, we omit the subscript \mathcal{D} on the satisfiability relation.

Definition 2.12. Let \mathcal{D} be a MDP. We define the language $\text{Sat}_{\mathcal{D}}(\phi)$ associated to an LTL formula ϕ as the set of paths satisfying ϕ , i.e. $\text{Sat}_{\mathcal{D}}(\phi) \triangleq \{\omega \in \text{Paths}(\mathcal{D}) \mid \omega \models \phi\}$. Here we also generally omit the subscript \mathcal{D} .

We now define satisfiability of an LTL formula ϕ on a MDP \mathcal{D} . We say that \mathcal{D} satisfies ϕ with probability at most p ($\mathcal{D} \models_{\leq p} \phi$) if the probability of getting an execution satisfying ϕ is at most p .

Definition 2.13. Let \mathcal{D} be a MDP, ϕ an LTL formula and $p \in [0, 1]$. We define $\models_{\leq p}$ and $\models_{\geq p}$ by

$$\begin{aligned} \mathcal{D} \models_{\leq p} \phi &\Leftrightarrow \mathbf{Pr}_{\mathcal{D}}^+(\text{Sat}(\phi)) \leq p, \\ \mathcal{D} \models_{\geq p} \phi &\Leftrightarrow \mathbf{Pr}_{\mathcal{D}}^-(\text{Sat}(\phi)) \geq p. \end{aligned}$$

We define $\mathcal{D} \models_{< p} \phi$ and $\mathcal{D} \models_{> p} \phi$ in a similar way.

In case the MDP is fully probabilistic, i.e., a MC, the satisfiability problem is reduced to $\mathcal{D} \models_{\bowtie p} \phi \Leftrightarrow \mathbf{Pr}_{\mathcal{D}}(\text{Sat}(\phi)) \bowtie p$, where $\bowtie \in \{<, \leq, >, \geq\}$.

3 Counterexamples

In this section, we define what counterexamples are and how the problem of finding counterexamples for a general LTL property over Markov Decision Processes reduces to finding counterexamples to reachability problems over Markov chains.

Definition 3.1 (Counterexamples). Let \mathcal{D} be a MDP and ϕ an LTL formula. A *counterexample* to $\mathcal{D} \models_{>p} \phi$ is a measurable set $\mathcal{C} \subseteq \text{Sat}(\phi)$ such that $\Pr_{\mathcal{D}}^+(\mathcal{C}) > p$. Counterexamples to $\mathcal{D} \models_{\leq p} \phi$ are defined similarly.

Counterexamples to $\mathcal{D} \models_{>p} \phi$ and $\mathcal{D} \models_{\geq p} \phi$ cannot be defined straightforwardly as it is always possible to find a set $\mathcal{C} \subseteq \text{Sat}(\phi)$ such that $\Pr_{\mathcal{D}}^-(\mathcal{C}) \leq p$ or $\Pr_{\mathcal{D}}^-(\mathcal{C}) < p$, note that the empty set trivially satisfies it. Therefore, the best way to find counterexamples to lower bounded probabilities is to find counterexamples to the dual properties $\mathcal{D} \models_{<1-p} \neg\phi$ and $\mathcal{D} \models_{\leq 1-p} \neg\phi$. That is, while for upper bounded probabilities, a counterexample is a set of paths satisfying the property beyond the bound, for lower bounded probabilities the counterexample is a set of paths that *does not* satisfy the property with sufficient probability.

Example 1. Consider the MDP \mathcal{D} of Figure 4 and the LTL formula $\diamond v$. it is easy to check that $\mathcal{D} \not\models_{<1} \diamond v$. The set $\mathcal{C} = \text{Sat}(\diamond v) = \{\gamma \in \text{Paths}(s_0) \mid \exists i \geq 0. \gamma = s_0(s_1)^i(s_4)^\omega\} \cup \{\gamma \in \text{Paths}(s_0) \mid \exists i \geq 0. \gamma = s_0(s_3)^i(s_5)^\omega\}$ is a counterexample. Note that $\Pr_{s_0, \eta}^+(\mathcal{C}) = 1$ where η is any deterministic scheduler of \mathcal{M} satisfying $\eta(s_0) = \pi_1$.

LTL formulas are actually checked by reducing the model checking problem to a reachability problem [dAKM97]. For checking upper bounded probabilities, the LTL formula is translated into an equivalent deterministic Rabin automaton and composed with the MDP under verification. On the obtained MDP, the set of states forming accepting end components (maximal components that traps accepting conditions with probability 1) are identified. The maximum probability of the LTL property on the original MDP is the same as the maximum probability of reaching a state of an accepting end component in the final MDP. Hence, from now on we will focus on counterexamples to properties of the form $\mathcal{D} \models_{\leq p} \diamond\psi$ or $\mathcal{D} \models_{<p} \diamond\psi$, where ψ is a propositional formula, i.e., a formula without temporal operators.

In the following, it will be useful to identify the set of states in which a propositional property is valid.

Definition 3.2. Let \mathcal{D} be a MDP. We define the state language $\text{Sat}_{\mathcal{D}}(\psi)$ associated to a propositional formula ψ as the set of states satisfying ψ , i.e., $\text{Sat}_{\mathcal{D}}(\psi) \triangleq \{s \in S \mid s \models \psi\}$, where \models has the obvious satisfaction meaning for states. As usual, we generally omit the subscript \mathcal{D} .

To find a counterexample to a property in a MDP with respect to an upper bound, it suffices to find a counterexample for the maximizing scheduler. A scheduler defines a Markov chain, and hence finding a counterexample on a MDP amounts to finding a counterexample in the Markov chain induced by the maximizing scheduler. The maximizing scheduler turns out to be deterministic and memoryless [BdA95]; consequently the induced Markov chain can be easily extracted from the MDP as follows.

Definition 3.3. Let $\mathcal{D} = (S, s_0, \tau, L)$ be a MDP and η a deterministic memoryless scheduler. Then we define the MC η -associated to \mathcal{D} as $\mathcal{D}_{\eta} = (S, s_0, \mathcal{P}_{\eta}, L)$ where $\mathcal{P}_{\eta}(s, t) = (\eta(s))(t)$ for all $s, t \in S$.




Fig. 4:

Now we state that finding counterexamples for upper bounded probabilistic reachability LTL properties on MDPs can be reduced to finding counterexamples for upper bounded probabilistic reachability LTL properties on MCs.

Theorem 3.4. *Let \mathcal{D} be a MDP, ψ a propositional formula and $p \in [0, 1]$. Then, there is a maximizing (deterministic memoryless) scheduler η such that $\mathcal{D} \models_{\leq p} \diamond\psi \Leftrightarrow \mathcal{D}_\eta \models_{\leq p} \diamond\psi$. Moreover, \mathcal{C} is a counterexample to $\mathcal{D}_\eta \models_{\leq p} \diamond\psi$ if and only if $\mathcal{C}^{\leq p}$ is also a counterexample to $\mathcal{D} \models_{\leq p} \diamond\psi$.*

4 Representative Counterexamples, Partitions and Witnesses

The notion of counterexample from Definition 3.1 is very broad: just an arbitrary (measurable) set of paths with high enough probability. To be useful as a debugging tool (and in fact to be able to present the counterexample to a user), we need counterexamples with specific properties. We will partition counterexamples (or rather, representative counterexamples) in witnesses and list five properties that witnesses should satisfy.

The first point to stress is that for reachability properties it is sufficient to consider counterexamples that consist of finite paths.

Definition 4.1 (Representative counterexamples). Let \mathcal{D} be a MDP, ψ a propositional formula and $p \in [0, 1]$. A *representative counterexample* to $\mathcal{D} \models_{\leq p} \diamond\psi$ is a set $\mathcal{C} \subseteq \text{Reach}^*(\mathcal{D}, \text{Sat}(\psi))$ such that $\Pr_{\mathcal{D}}^+(\langle\mathcal{C}\rangle) > p$. We denote the set of $\leq p$ all representative counterexamples to $\mathcal{M} \models_{\leq p} \diamond\psi$ by $\mathcal{R}(\mathcal{M}, p, \psi)$.

Theorem 4.2. *Let \mathcal{D} be a MDP, ψ a propositional formula and $p \in [0, 1]$. If \mathcal{C} is a representative counterexample to $\mathcal{D} \models_{\leq p} \diamond\psi$, then $\langle\mathcal{C}\rangle$ is a counterexample to $\mathcal{D} \models_{\leq p} \diamond\psi$. Furthermore, there exists a counterexample to $\mathcal{D} \models_{\leq p} \diamond\psi$ if and only if there exists a representative counterexample to $\mathcal{D} \models_{\leq p} \psi$.*

Following [HK07a], we present the notions of *minimum counterexample*, *strongest evidence* and *most indicative counterexamples*.

Definition 4.3 (Minimum counterexample). Let \mathcal{M} be a MC, ψ a propositional formula and $p \in [0, 1]$. We say that $\mathcal{C} \in \mathcal{R}(\mathcal{M}, p, \psi)$ is a *minimum counterexample* if $|\mathcal{C}| \leq |\mathcal{C}'|$, for all $\mathcal{C}' \in \mathcal{R}(\mathcal{M}, p, \psi)$.

Definition 4.4 (Strongest evidence). Let \mathcal{M} be a MC, ψ a propositional formula and $p \in [0, 1]$. A *strongest evidence* to $\mathcal{M} \not\models_{\leq p} \diamond\psi$ is a finite path $\sigma \in \text{Reach}^*(\mathcal{M}, \text{Sat}(\psi))$ such that $\Pr_{\mathcal{M}}(\langle\sigma\rangle) \geq \Pr_{\mathcal{M}}(\langle\rho\rangle)$, for all $\rho \in \text{Reach}^*(\mathcal{M}, \text{Sat}(\psi))$.

Definition 4.5 (Most indicative counterexample). Let \mathcal{M} be a MC, ψ a propositional formula and $p \in [0, 1]$. We call $\mathcal{C} \in \mathcal{R}(\mathcal{M}, p, \psi)$ a *most indicative counterexample* if it is minimum and $\Pr(\langle\mathcal{C}\rangle) \geq \Pr(\langle\mathcal{C}'\rangle)$, for all minimum counterexamples $\mathcal{C}' \in \mathcal{R}(\mathcal{M}, p, \psi)$.

Unfortunately, very often most indicative counterexamples are very large (even infinite), many of its elements have insignificant measure and elements can be extremely similar to each other (consequently providing the same diagnostic information). Even worse, sometimes the finite paths with highest probability do not exhibit the way in which the system accumulates higher probability to reach the undesired property (and consequently where an error occurs with higher probability). For these reasons, we are of the opinion that representative counterexamples are still too general in order to be useful as feedback information. We approach this problem by splitting out the representative counterexample into sets of finite paths following a “similarity” criteria (introduced in Section 5). These sets are called *witnesses of the counterexample*.

Recall that a set Y of nonempty sets is a partition of X if the elements of Y cover X and the elements of Y are pairwise disjoint. We define counterexample partitions in the following way.

Definition 4.6 (Counterexample partitions and witnesses). Let \mathcal{D} be a MDP, ψ a propositional formula, $p \in [0, 1]$, and \mathcal{C} a representative counterexample to $\mathcal{D} \models_{\leq p} \diamond\psi$. A *counterexample partition* $W_{\mathcal{C}}$ is a partition of \mathcal{C} . We call the elements of $W_{\mathcal{C}}$ *witnesses*.

Since not every partition generates useful witnesses (from the debugging perspective), we now state properties that witnesses must satisfy in order to be valuable as diagnostic information. In Section 7 we show how to partition the detailed counterexample in order to obtain useful witnesses.

Similarity: Elements of a witness should provide similar debugging information.

Accuracy: Witnesses with higher probability should show evolution of the system with higher probability of containing errors.

Originality: Different witnesses should provide different debugging information.

Significance: The probability of a witnesses should be close to the probability bound p .

Finiteness: The number of witnesses of a counterexamples partition should be finite.

5 Rails and Torrents

As argued before we consider that representative counterexamples are excessively general to be useful as feedback information. Therefore, we group finite paths of a representative counterexample in witnesses if they are “similar enough”. We will consider finite paths that behave the same outside SCCs of the system as providing similar feedback information.

In order to formalize this idea, we first reduce the original Markov chain to an acyclic one that preserves reachability probabilities. We do so by removing all SCCs K of \mathcal{M} keeping just *input states* of K . In this way, we get a new acyclic MC denoted by $\text{Ac}(\mathcal{M})$. The probability matrix of the Markov chain relates input states of each SCC with its *output states* with the reachability probability between these states in \mathcal{M} . Secondly, we establish a map between finite paths σ in $\text{Ac}(\mathcal{M})$ (*rails*) and sets of finite paths W_{σ} in \mathcal{M} (*torrents*). Each torrent contains finite paths that are similar, i.e., behave the same outside SCCs. Additionally we show that the probability of σ is equal to the probability of W_{σ} .

Reduction to Acyclic Markov Chains


Consider a MC $\mathcal{M} = (S, s_0, \mathcal{P}, L)$. Recall that a subset $K \subseteq S$ is called *strongly connected* if for every $s, t \in K$ there is a finite path from s to t . Additionally K is called a *strongly connected component* (SCC) if it is a maximally (with respect to \subseteq) strongly connected subset of S .

Note that every state is a member of exactly one SCC of \mathcal{M} (even those states that are not involved in cycles, since the trivial finite path s connects s to itself). From now on we let SCC^* be the set of non trivial strongly connected components of a MC, i.e., those composed of more than one state.

A Markov chain is called *acyclic* if it does not have non trivial SCCs. Note that an acyclic Markov chain still has absorbing states.

Definition 5.1. Let $\mathcal{M} = (S, s_0, \mathcal{P}, L)$ be a MC. Then, for each $\text{SCC}^* K$ of \mathcal{M} , we define the sets $\text{Inp}_K \subseteq S$ of all states in K that have an incoming transition from a state outside of K and $\text{Out}_K \subseteq S$ of all states outside of K that have an incoming transition from a state of K in the following way

$$\begin{aligned} \text{Inp}_K &\triangleq \{u \in K \mid \exists s \in S \setminus K. \mathcal{P}(s, u) > 0\}, \\ \text{Out}_K &\triangleq \{s \in S \setminus K \mid \exists u \in K. \mathcal{P}(u, s) > 0\}. \end{aligned}$$



We also define for each SCC^* K a MC related to K as $\mathcal{M}_K \triangleq (K \cup \text{Out}_K, s_K, \mathcal{P}_K, L_K)$ where s_K is any state in Inp_K , $L_K(s) \triangleq L(s)$, and $\mathcal{P}_K(s, t)$ is equal to $\mathcal{P}(s, t)$ if $s \in K$ and equal to 1_s otherwise. Additionally, for every state s involved in non trivial SCCs we define SCC_s^+ as \mathcal{M}_K , where K is the SCC^* of \mathcal{M} such that $s \in K$.

Now we are able to define an acyclic MC $\text{Ac}(\mathcal{M})$ related to \mathcal{M} .

Definition 5.2. Let $\mathcal{M} = (S, s_0, \mathcal{P}, L)$ be a MC. We define $\text{Ac}(\mathcal{M}) \triangleq (S', s_0, \mathcal{P}', L')$ where

- $S' \triangleq S \setminus \bigcup_{K \in \text{SCC}^*} K \cup \bigcup_{K \in \text{SCC}^*} \text{Inp}_K$
- $L' \triangleq L|_{S'}$,
- $\mathcal{P}'(s, t) \triangleq \begin{cases} \mathcal{P}(s, t) & \text{if } s \in S_{\text{com}}, \\ \mathbf{Pr}_{\mathcal{M}, s}(\text{Reach}(\text{SCC}_s^+, s, \{t\})) & \text{if } s \in S_{\text{inp}} \wedge t \in \text{Out}_{\text{SCC}_s^+}, \\ 1_s & \text{if } s \in S_{\text{inp}} \wedge \text{Out}_{\text{SCC}_s^+} = \emptyset, \\ 0 & \text{otherwise.} \end{cases}$

Note that $\text{Ac}(\mathcal{M})$ is indeed acyclic.

Example 2. Consider the MC \mathcal{M} of Figure 5(a). The strongly connected components of \mathcal{M} are $K_1 \triangleq \{s_1, s_3, s_4, s_7\}$, $K_2 \triangleq \{s_5, s_6, s_8\}$ and the singletons $\{s_0\}$, $\{s_2\}$, $\{s_9\}$, $\{s_{10}\}$, $\{s_{11}\}$, $\{s_{12}\}$, $\{s_{13}\}$, and $\{s_{14}\}$. The input states of K_1 are $\text{Inp}_{K_1} = \{s_1\}$ and its output states are $\text{Out}_{K_1} = \{s_9, s_{10}\}$. For K_2 , $\text{Inp}_{K_2} = \{s_5, s_6\}$ and $\text{Out}_{K_2} = \{s_{11}, s_{14}\}$. The reduced acyclic MC of \mathcal{M} is shown in Figure 5(b).




Fig. 5:

Rails and Torrents

We now relate (finite) paths in $\text{Ac}(\mathcal{M})$ (rails) to sets of (finite) paths in \mathcal{M} (torrents).

Definition 5.3 (Rails). Let \mathcal{M} be a MC. A finite path $\sigma \in \text{Paths}^*(\text{Ac}(\mathcal{M}))$ will be called a *rail* of \mathcal{M} .

Consider a rail σ , i.e., a finite path of $\text{Ac}(\mathcal{M})$. We will use σ to represent those paths ω of \mathcal{M} that behave “similar to” σ outside SCCs of \mathcal{M} . Naively, this means that σ is a subsequence of ω . There are two technical subtleties to deal with: every input state in σ must be the first state in its SCC in ω (freshness) and every SCC visited by ω must be also visited by σ (inertia) (see Definition 5.5). We need these extra conditions to make sure that no path ω behaves “similar to” two distinct rails (see Lemma 5.7).

Recall that given a finite sequence σ and a (possible infinite) sequence ω , we say that σ is a *subsequence* of ω , denoted by $\sigma \sqsubseteq \omega$, if and only if there exists a strictly increasing function $f : \{0, 1, \dots, |\sigma| - 1\} \rightarrow \{0, 1, \dots, |\omega| - 1\}$ such that $\forall_{0 \leq i < |\sigma|} \sigma_i = \omega_{f(i)}$. If ω is an infinite sequence, we interpret the codomain of f as \mathbb{N} . In case f is such a function we write $\sigma \sqsubseteq_f \omega$. Note that finite paths and paths are sequences.

Definition 5.4. Let $\mathcal{M} = (S, s_0, \mathcal{P}, L)$ be a MC. On S we consider the equivalence relation $\approx_{\mathcal{M}}$ satisfying $s \approx_{\mathcal{M}} t$ if and only if s and t are in the same strongly connected component. Again, we usually omit the subscript \mathcal{M} from the notation.

The following definition refines the notion of subsequence, taking care of the two technical subtleties noted above.

Definition 5.5. Let $\mathcal{M} = (S, s_0, \mathcal{P}, L)$ be a MC, ω a (finite) path of \mathcal{M} , and $\sigma \in \text{Paths}^*(\text{Ac}(\mathcal{M}))$ a finite path of $\text{Ac}(\mathcal{M})$. Then we write $\sigma \preceq \omega$ if there exists $f : \{0, 1, \dots, |\sigma| - 1\} \rightarrow \mathbb{N}$ such that $\sigma \sqsubseteq_f \omega$ and for all $0 \leq i < |\sigma|$ we have

$$\begin{aligned} \forall_{0 \leq j < f(i)} : \omega_{f(i)} \not\sim \omega_j; & \text{ for all } i = 0, 1, \dots, |\sigma| - 1, & [\textit{Freshness property}] \\ \forall_{f(i) < j < f(i+1)} : \omega_{f(i)} \sim \omega_j; & \text{ for all } i = 0, 1, \dots, |\sigma| - 2. & [\textit{Inertia property}] \end{aligned}$$

In case f is such a function we write $\sigma \preceq_f \omega$.

Example 3. Let $\mathcal{M} = (S, s_0, \mathcal{P}, L)$ be the MC of Figure 5(a) and take $\sigma = s_0 s_2 s_6 s_{14}$. Then for all $i \in \mathbb{N}$ we have $\sigma \preceq_{f_i} \omega_i$ where $\omega_i = s_0 s_2 s_6 (s_5 s_8 s_6)^i s_{14}$ and $f_i(0) \triangleq 0$, $f_i(1) \triangleq 1$, $f_i(2) \triangleq 2$, and $f_i(3) \triangleq 3 + 3i$. Additionally, $\sigma \not\preceq s_0 s_2 s_5 s_8 s_6 s_{14}$ since for all f satisfying $\sigma \sqsubseteq_f s_0 s_2 s_5 s_8 s_6 s_{14}$ we must have $f(2) = 5$; this implies that f does not satisfy the freshness property. Finally, note that $\sigma \not\preceq s_0 s_2 s_6 s_{11} s_{14}$ since for all f satisfying $\sigma \sqsubseteq_f s_0 s_2 s_6 s_{11} s_{14}$ we must have $f(2) = 2$; this implies that f does not satisfy the inertia property.

We now give the formal definition of torrents.

Definition 5.6 (Torrents). Let $\mathcal{M} = (S, s_0, \mathcal{P}, L)$ be a MC and σ a sequence of states in S . We define the function Torr by

$$\text{Torr}(\mathcal{M}, \sigma) \triangleq \{\omega \in \text{Paths}(\mathcal{M}) \mid \sigma \preceq \omega\}.$$

We call $\text{Torr}(\mathcal{M}, \sigma)$ the *torrent* associated to σ .

We now show that torrents are disjoint (Lemma 5.7) and that the probability of a rail is equal to the probability of its associated torrent (Theorem 5.10). For this last result, we first show that torrents can be represented as the disjoint union of cones of finite paths. We call these finite paths *generators* of the torrent (Definition 5.8).

Lemma 5.7. *Let \mathcal{M} be a MC. For every $\sigma, \rho \in \text{Paths}^*(\text{Ac}(\mathcal{M}))$ we have*

$$\sigma \neq \rho \Rightarrow \text{Torr}(\mathcal{M}, \sigma) \cap \text{Torr}(\mathcal{M}, \rho) = \emptyset$$

Definition 5.8 (Torrent Generators). Let \mathcal{M} be a MC. Then we define for every rail $\sigma \in \text{Paths}^*(\text{Ac}(\mathcal{M}))$ the set

$$\text{GenTorr}(\mathcal{M}, \sigma) \triangleq \{\rho \in \text{Paths}^*(\mathcal{M}) \mid \exists f : \sigma \preceq_f \rho \wedge f(|\sigma| - 1) = |\rho| - 1\}.$$

In the example from the Introduction (see Figure 1), $s_0 s_1 s_3$ and $s_0 s_2 s_4$ are rails. The associated torrents are, respectively, $\{s_0 s_1^n s_3^{\omega} \mid n \in \mathbb{N}^*\}$ and $\{s_0 s_2^n s_4^{\omega} \mid n \in \mathbb{N}^*\}$ (note that s_3 and s_4 are absorbing states), i.e. the paths going left and the paths going right. The generators of the first torrent are $\{s_0 s_1^n s_3 \mid n \in \mathbb{N}^*\}$ and similarly for the second torrent.

Lemma 5.9. *Let \mathcal{M} be a MC and $\sigma \in \text{Paths}^*(\text{Ac}(\mathcal{M}))$ a rail of \mathcal{M} . Then we have*

$$\text{Torr}(\mathcal{M}, \sigma) = \bigsqcup_{\rho \in \text{GenTorr}(\mathcal{M}, \sigma)} \langle \rho \rangle.$$

Theorem 5.10. *Let \mathcal{M} be a MC. Then for every rail $\sigma \in \text{Paths}^*(\text{Ac}(\mathcal{M}))$ we have*

$$\mathbf{Pr}_{\text{Ac}(\mathcal{M})}(\langle \sigma \rangle) = \mathbf{Pr}_{\mathcal{M}}(\text{Torr}(\mathcal{M}, \sigma)).$$

6 Significant Diagnostic Counterexamples

So far we have formalized the notion of paths behaving similarly (i.e., behaving the same outside SCCs) in a MC \mathcal{M} by removing all SCC of \mathcal{M} , obtaining $\text{Ac}(\mathcal{M})$. A representative counterexample to $\text{Ac}(\mathcal{M}) \models_{\leq p} \diamond\psi$ will give rise to a representative counterexample to $\mathcal{M} \models_{\leq p} \diamond\psi$. For every finite path σ in the counterexample to $\text{Ac}(\mathcal{M}) \models_{\leq p} \diamond\psi$, the set $\text{GenTorr}(\mathcal{M}, \sigma)$ will be a witness. The union of these is the representative counterexample to $\mathcal{M} \models_{\leq p} \diamond\psi$.

Before giving a formal definition, there is still one technical issue to resolve: we need to be sure that by removing SCCs we are not discarding useful information. Because torrents are built from rails, we need to make sure that when we discard SCCs, we do not discard rails that reach ψ .

We achieve this by first making states satisfying ψ absorbing. Additionally, we make absorbing states from which it is not possible to reach ψ . Note that this does not affect counterexamples.

Definition 6.1. Let $\mathcal{M} = (S, s_0, \mathcal{P}, L)$ be a MC and ψ a propositional formula. We define the MC $\mathcal{M}_\psi \triangleq (S, s_0, \mathcal{P}_\psi, L)$, with

$$\mathcal{P}_\psi(s, t) \triangleq \begin{cases} 1 & \text{if } s \notin \text{Sat}_\diamond(\psi) \wedge s = t, \\ 1 & \text{if } s \in \text{Sat}(\psi) \wedge s = t, \\ \mathcal{P}(s, t) & \text{if } s \in \text{Sat}_\diamond(\psi) - \text{Sat}(\psi), \\ 0 & \text{otherwise,} \end{cases}$$

where $\text{Sat}_\diamond(\psi) \triangleq \{s \in S \mid \mathbf{Pr}_s(\text{Reach}(\mathcal{M}, s, \text{Sat}(\psi))) > 0\}$ is the set of states reaching ψ in \mathcal{M} .

The following theorem shows the relation between paths, finite paths, and probabilities of \mathcal{M} , \mathcal{M}_ψ , and $\text{Ac}(\mathcal{M}_\psi)$. Most importantly, the probability of a rail σ (in $\text{Ac}(\mathcal{M}_\psi)$) is equal to the probability of its associated torrent (in \mathcal{M}) (item 5 below) and the probability of $\diamond\psi$ is not affected by reducing \mathcal{M} to $\text{Ac}(\mathcal{M}_\psi)$ (item 6 below).

Note that a rail σ is always a finite path in $\text{Ac}(\mathcal{M}_\psi)$, but that we can talk about its associated torrent $\text{Torr}(\mathcal{M}_\psi, \sigma)$ in \mathcal{M}_ψ and about its associated torrent $\text{Torr}(\mathcal{M}, \sigma)$ in \mathcal{M} . The former exists for technical convenience; it is the latter that we are ultimately interested in. The following theorem also shows that for our purposes, viz. the definition of the generators of the torrent and the probability of the torrent, there is no difference (items 3 and 4 below).

Theorem 6.2. Let $\mathcal{M} = (S, s_0, \mathcal{P}, L)$ be a MC and ψ a propositional formula. Then for every $\sigma \in \text{Paths}^*(\mathcal{M}_\psi)$

1. $\text{Reach}^*(\mathcal{M}_\psi, s_0, \text{Sat}(\psi)) = \text{Reach}^*(\mathcal{M}, s_0, \text{Sat}(\psi))$,
2. $\mathbf{Pr}_{\mathcal{M}_\psi}(\langle\sigma\rangle) = \mathbf{Pr}_{\mathcal{M}}(\langle\sigma\rangle)$,
3. $\text{GenTorr}(\mathcal{M}_\psi, \sigma) = \text{GenTorr}(\mathcal{M}, \sigma)$,
4. $\mathbf{Pr}_{\mathcal{M}_\psi}(\text{Torr}(\mathcal{M}_\psi, \sigma)) = \mathbf{Pr}_{\mathcal{M}}(\text{Torr}(\mathcal{M}, \sigma))$,
5. $\mathbf{Pr}_{\text{Ac}(\mathcal{M}_\psi)}(\langle\sigma\rangle) = \mathbf{Pr}_{\mathcal{M}}(\text{Torr}(\mathcal{M}, \sigma))$,
6. $\text{Ac}(\mathcal{M}_\psi) \models_{\leq p} \diamond\psi$ if and only if $\mathcal{M} \models_{\leq p} \diamond\psi$, for any $p \in [0, 1]$.

Definition 6.3 (Torrent-Counterexamples). Let $\mathcal{M} = (S, s_0, \mathcal{P}, L)$ be a MC, ψ a propositional formula, and $p \in [0, 1]$ such that $\mathcal{M} \not\models_{\leq p} \diamond\psi$. Let \mathcal{C} be a representative counterexample to $\text{Ac}(\mathcal{M}_\psi) \models_{\leq p} \diamond\psi$. We define the set

$$\text{TorRepCount}(\mathcal{C}) \triangleq \{\text{GenTorr}(\mathcal{M}, \sigma) \mid \sigma \in \mathcal{C}\}.$$

We call the set $\text{TorRepCount}(\mathcal{C})$ a *torrent-counterexample* of \mathcal{C} . Note that this set is a partition of a counterexample to $\mathcal{M} \models_{\leq p} \diamond\psi$. Additionally, we denote by $\mathcal{R}_t(\mathcal{M}, p, \psi)$ to the set of all torrent-counterexamples to $\mathcal{M} \models_{\leq p} \diamond\psi$, i.e., $\{\text{TorRepCount}(\mathcal{C}) \mid \mathcal{C} \in \mathcal{R}(\mathcal{M}, p, \psi)\}$.

Theorem 6.4. *Let $\mathcal{M} = (S, s_0, \mathcal{P}, L)$ be a MC, ψ a propositional formula, and $p \in [0, 1]$ such that $\mathcal{M} \not\models_{\leq p} \diamond\psi$. Take \mathcal{C} a representative counterexample to $\text{Ac}(\mathcal{M}_\psi) \models_{\leq p} \diamond\psi$. Then the set of finite paths $\biguplus_{W \in \text{TorRepCount}(\mathcal{C})} W$ is a representative counterexample to $\mathcal{M} \models_{\leq p} \diamond\psi$.*

Note that for each $\sigma \in \mathcal{C}$ we get a witness $\text{GenTorr}(\mathcal{M}, \sigma)$. Also note that the number of rails is finite, so there are also only finitely many witnesses.

Following [HK07a], we extend the notions of *minimum counterexamples*, *strongest evidence* and *smallest counterexample* to torrents.

Definition 6.5 (Minimum torrent-counterexample). Let \mathcal{M} be a MC, ψ a propositional formula and $p \in [0, 1]$. We say that $\mathcal{C}_t \in \mathcal{R}_t(\mathcal{M}, p, \psi)$ is a *minimum torrent-counterexample* if $|\mathcal{C}_t| \leq |\mathcal{C}'_t|$, for all $\mathcal{C}'_t \in \mathcal{R}_t(\mathcal{M}, p, \psi)$.

Definition 6.6 (Strongest torrent-evidence). Let \mathcal{M} be a MC, ψ a propositional formula and $p \in [0, 1]$. A *strongest torrent-evidence* to $\mathcal{M} \not\models_{\leq p} \diamond\psi$ is a torrent $W_\sigma \in \text{Torr}(\mathcal{M}, \text{Sat}(\psi))$ such that $\Pr_{\mathcal{M}}(W_\sigma) \geq \Pr_{\mathcal{M}}(W_\rho)$ for all $W_\rho \in \text{Torr}(\mathcal{M}, \text{Sat}(\psi))$.

Now we define our notion of significant diagnostic counterexamples. It is the generalization of most indicative counterexample from [HK07a] to our setting.

Definition 6.7 (Most indicative torrent-counterexample). Let \mathcal{M} be a MC, ψ a propositional formula and $p \in [0, 1]$. We call $\mathcal{C}_t \in \mathcal{R}_t(\mathcal{M}, p, \psi)$ a *most indicative torrent-counterexample* if it is a minimum torrent-counterexample and $\Pr(\bigcup_{W \in \mathcal{C}_t} \langle W \rangle) \geq \Pr(\bigcup_{W \in \mathcal{C}'_t} \langle W \rangle)$ for all minimum torrent counterexamples $\mathcal{C}'_t \in \mathcal{R}_t(\mathcal{M}, p, \psi)$.

By Theorem 6.4 it is possible to obtain strongest torrent-evidence and most indicative torrent-counterexamples of a MC \mathcal{M} by obtaining strongest evidence and most indicative counterexamples of $\text{Ac}(\mathcal{M}_\psi)$ respectively.

7 Computing Counterexamples

In this section we show how to compute most indicative torrent-counterexamples. We also discuss what information to present to the user: how to present witnesses and how to deal with overly large strongly connected components.

7.1 Maximizing Schedulers

The calculation of a maximal probability on a reachability problem can be performed by solving a linear minimization problem [BdA95,dA97]. This minimization problem is defined on a system of inequalities that has a variable x_i for each different state s_i and an inequality $\sum_j \pi(s_j) \cdot x_j \leq x_i$ for each distribution $\pi \in \tau(s_i)$. The maximizing (deterministic memoryless) scheduler η can be easily extracted out of such system of inequalities after obtaining the solution. If p_0, \dots, p_n are the values that minimize $\sum_i x_i$ in the previous system, then η is such that, for all s_i , $\eta(s_i) = \pi$ whenever $\sum_j \pi(s_j) \cdot p_j = p_i$. In the following we denote $\mathbf{P}_{s_i}[\diamond\psi] \triangleq x_i$.

7.2 Computing most indicative torrent-counterexamples

We divide the computation of most indicative torrent-counterexamples to $\mathcal{D} \models_{\leq p} \diamond\psi$ in three stages: *pre-processing*, *SCC analysis*, and *searching*.

Pre-processing stage. We first modify the original MC \mathcal{M} by making all states in $\text{Sat}(\psi) \cup S \setminus \text{Sat}_\diamond(\psi)$ absorbing. In this way we obtain the MC \mathcal{M}_ψ from Definition 6.1. Note that we do not have to spend additional computational resources to compute this set, since $\text{Sat}_\diamond(\psi) = \{s \in S \mid \mathbf{P}_s[\psi] > 0\}$ and hence all required data is already available from the LTL model checking phase.

SCC analysis stage. We remove all SCCs K of \mathcal{M}_ψ keeping just *input states* of K , getting the acyclic MC $\text{Ac}(\mathcal{M}_\psi)$ according to Definition 5.2.

To compute this, we first need to find the SCCs of \mathcal{M}_ψ . There exists well known algorithms to achieve this: Kosaraju’s, Tarjan’s, Gabow’s algorithms (among others). We also have to compute the reachability probability from input states to output states of every SCC. This can be done by using steady state analysis techniques [Cas93].

Searching stage. To find most indicative torrent-counterexamples in \mathcal{M} , we find most indicative counterexamples in $\text{Ac}(\mathcal{M}_\psi)$. For this we use the same approach as [HK07a], turning the MC into a weighted digraph to exchange the problem of finding the finite path with highest probability by a shortest path problem. The nodes of the digraph are the states of the MC and there is an edge between s and t if $\mathcal{P}(s, t) > 0$. The weight of such an edge is $-\log \mathcal{P}(s, t)$.

Finding the most indicative counterexample in $\text{Ac}(\mathcal{M}_\psi)$ is now reduced to finding k shortest paths. As explained in [HK07a], our algorithm has to compute k on the fly. Eppstein’s algorithm [Epp98] produces the k shortest paths in general in $O(m + n \log n + k)$, where m is the number of nodes and n the number of edges. In our case, since $\text{Ac}(\mathcal{M}_\psi)$ is acyclic, the complexity decreases to $O(m + k)$.

7.3 Debugging issues

Representative finite paths. What we have computed so far is a most indicative counterexample to $\text{Ac}(\mathcal{M}_\psi) \models_{\leq p} \diamond\psi$. This is a finite set of rails, i.e., a finite set of paths in $\text{Ac}(\mathcal{M}_\psi)$. Each of these paths σ represents a witness $\text{GenTorr}(\mathcal{M}, \sigma)$. Note that this witness itself has usually infinitely many elements.

In practice, one somehow has to display a witness to the user. The obvious way would be to show the user the rail σ . This, however, may be confusing to the user as σ is not a finite path of the original Markov Decision Process. Instead of presenting the user with σ , we therefore show the user the element of $\text{GenTorr}(\mathcal{M}, \sigma)$ with highest probability.

Definition 7.1. Let \mathcal{M} be a MC, and $\sigma \in \text{Paths}^*(\text{Ac}(\mathcal{M}_\psi))$ a rail of \mathcal{M} . We define the *representant* of $\text{Torr}(\mathcal{M}, \sigma)$ as

$$\text{repTorr}(\mathcal{M}, \sigma) = \text{repTorr} \left(\biguplus_{\rho \in \text{GenTorr}(\mathcal{M}, \sigma)} \langle \rho \rangle \right) \triangleq \arg \max_{\rho \in \text{GenTorr}(\mathcal{M}, \sigma)} \Pr(\langle \rho \rangle)$$

Note that given $\text{repTorr}(\mathcal{M}, \sigma)$, one can easily recover σ . Therefore, no information is lost by presenting torrents as a single element of the torrent instead of as a rail.

Expanding SCC. It is possible that the system contains some very large strongly connected components. In that case, a single witness could have a very large probability mass and one could argue that the information presented to the user is not detailed enough. For instance, consider the Markov chain of Figure 6 in which there is a single large SCC with input state t and output state u .

The most-indicative torrent counterexample to the property $\mathcal{M} \models_{\leq 0.9} \diamond\psi$ is simply $\{\text{GenTorr}(stu)\}$, i.e., a single witness with probability mass 1 associated to the rail stu . Although this may seem uninformative, we argue that it is more informative than listing several paths of the form $st \cdots u$ with probability summing up to, say, 0.91. Our single witness counterexample suggests that the outgoing edge to a state not reaching ψ was simply forgotten; the listing of paths still allows the possibility that one of the probabilities in the whole system is simply wrong.




Fig. 6:

Nevertheless, if the user needs more information to tackle bugs inside strongly connected components, note that there is more information available at this point. In particular, for every strongly connected component K , every input state s of K (even for every state in K), and every output state t of K , the probability of reaching t from s is already available from the computation of $\text{Ac}(\mathcal{M}_\psi)$ during the SCC analysis stage of Section 7.2.

8 Final Discussion

We have presented a novel technique for representing and computing counterexamples for nondeterministic and probabilistic systems. We partition a counterexample in witnesses and state five properties that we believe good witnesses should satisfy in order to be useful as debugging tool: (similarity) elements of a witness should provide similar debugging information; (originality) different witnesses should provide different debugging information; (accuracy) witnesses with higher probability should indicate system behavior more likely to contain errors; (significance) probability of a witness should be relatively high; (finiteness) there should be finitely many witnesses. We achieve this by grouping finite paths in a counterexample together in a witness if they behave the same outside the strongly connected components.

Presently, some work has been done on counterexample generation techniques for different variants of probabilistic models (Discrete Markov chains and Continuous Markov chains) [AHL05,AL06,HK07a,HK07b]. In our terminology, these works consider witnesses consisting of a *single* finite path. We have already discussed in the Introduction that the single path approach does not meet the properties of accuracy, originality, significance, and finiteness.

Instead, our witness/torrent approach provides a high level of abstraction of a counterexample. By grouping together finite paths that behave the same outside strongly connected components in a single witness, we can achieve these properties to a higher extent. Behaving the same outside strongly connected components is a reasonable way of formalizing the concept of providing *similar* debugging information. This grouping also makes witnesses significantly different from each other: each witness comes from a different rail and each rail provides a different way to reach the undesired property. Then each witness provides *original* information. Of course, our witnesses are more *significant* than single finite paths, because they are sets of finite paths. This also gives us more *accuracy* than the approach with single finite paths, as a collection of finite paths behaving the same and reaching an undesired condition with high probability is more likely to show how the system reaches this condition than just a single path. Finally, because there is a finite number of rails, there is also a *finite* number of witnesses.

Another key difference of our work to previous ones is that our technique allows us to generate counterexamples for probabilistic systems *with* nondeterminism. However, a recent report [AL07] also considers counterexample generation for MDPs. This work is limited to upper bounded pCTL formulas without nested temporal operators. Besides, their technique significantly differs from ours.

Finally, among the related work, we would like to stress the result of [HK07a], which provides a systematic characterization of counterexample generation in terms of shortest paths problems. We use this result to generate counterexamples for the acyclic Markov Chains.

In the future we intend to implement a tool to generate our significant diagnostic counterexamples; a very preliminary version has already been implemented. There is still work to be done on improving the visualization of the witnesses, in particular, when a witness captures a large strongly connected component. Another direction is to investigate how this work can be extended to timed systems, either modeled with continuous time Markov chains or with probabilistic timed automata.

References

- AHL05. Husain Aljazzar, Holger Hermanns, and Stefan Leue. Counterexamples for timed probabilistic reachability. In *Formal Modeling and Analysis of Timed Systems (FORMATS '05)*, volume 3829, pages 177–195, 2005.
- AL06. Husain Aljazzar and Stefan Leue. Extended directed search for probabilistic timed reachability. In *Formal Modeling and Analysis of Timed Systems (FORMATS '06)*, pages 33–51, 2006.
- AL07. Husain Aljazzar and Stefan Leue. Counterexamples for model checking of markov decision processes. Computer Science Technical Report soft-08-01, University of Konstanz, December 2007.
- BdA95. Andrea Bianco and Luca de Alfaro. Model checking of probabilistic and non-deterministic systems. In G. Goos, J. Hartmanis, and J. van Leeuwen, editors, *Foundations of Software Technology and Theoretical Computer Science (FSTTCS '95)*, volume 1026, pages 499–513, 1995.
- Bel57. Richard E. Bellman. A Markovian decision process. *J. Math. Mech.*, 6:679–684, 1957.
- BLR05. Gerd Behrmann, Kim G. Larsen, and Jacob I. Rasmussen. Optimal scheduling using priced timed automata. *SIGMETRICS Perform. Eval. Rev.*, 32(4):34–40, 2005.
- Cas93. Christos G. Cassandras. *Discrete Event Systems: Modeling and Performance Analysis*. Richard D. Irwin, Inc., and Aksen Associates, Inc., 1993.
- CGJ⁺00. Edmund M. Clarke, Orna Grumberg, Somesh Jha, Yuan Lu, and Helmut Veith. Counterexample-guided abstraction refinement. In *Computer Aided Verification*, pages 154–169, 2000.
- dA97. Luca de Alfaro. *Formal Verification of Probabilistic Systems*. PhD thesis, Stanford University, 1997.
- dAKM97. Luca de Alfaro, Arjun Kapur, and Zohar Manna. Hybrid diagrams: A deductive-algorithmic approach to hybrid system verification. In *Symposium on Theoretical Aspects of Computer Science*, pages 153–164, 1997.
- Epp98. David Eppstein. Finding the k shortest paths. In *SIAM Journal of Computing*, pages 652–673, 1998.
- FV97. J. Filar and K. Vrieze. *Competitive Markov Decision Processes*. 1997.
- HK07a. Tingting Han and Joost-Pieter Katoen. Counterexamples in probabilistic model checking. In *Tools and Algorithms for the Construction and Analysis of Systems: 13th International Conference (TACAS '07)*, volume 4424, pages 60–75, 2007.
- HK07b. Tingting Han and Joost-Pieter Katoen. Providing evidence of likely being on time counterexample generation for ctmc model checking. In *International Symposium on Automated Technology for Verification and Analysis (ATVA '07)*, volume 4762, pages 331–346, 2007.
- MP91. Z. Manna and A. Pnueli. *The Temporal Logic of Reactive and Concurrent Systems: Specification*. Springer, 1991.
- PZ93. Amir Pnueli and Lenore D. Zuck. Probabilistic verification. *Information and Computation*, 103(1):1–29, 1993.
- SdV04. Ana Sokolova and Erik P. de Vink. Probabilistic automata: System types, parallel composition and comparison. In Christel Baier, Boudewijn R. Haverkort, Holger Hermanns, Joost-Pieter Katoen, and Markus Siegle, editors, *Validation of Stochastic Systems: A Guide to Current Research*, volume 2925, pages 1–43. 2004.
- SL95. Roberto Segala and Nancy Lynch. Probabilistic simulations for probabilistic processes. *Nordic Journal of Computing*, 2(2):250–273, 1995.
- Var85. M.Y. Vardi. Automatic verification of probabilistic concurrent finite-state systems. In *Proc. 26th IEEE Symp. Found. Comp. Sci.*, pages 327–338, 1985.

Appendix: Proofs

In this appendix we give proofs of the results that were omitted from the paper for space reasons.

Observation 8.1. *Let \mathcal{M} be a MC. Since $\text{Ac}(\mathcal{M})$ is acyclic we have $\sigma_i \not\sim \sigma_j$ for every $\sigma \in \text{Paths}^*(\text{Ac}(\mathcal{M}))$ and $i \neq j$ (with the exception of absorbing states).*

Observation 8.2. *Let σ, ω and f be such that $\sigma \preceq_f \omega$. Then $\forall i : \exists j : \omega_i \sim \sigma_j$. This follows from $\sigma \sqsubseteq_f \omega$ and the inertia property.*

Lemma 8.3. *Let \mathcal{M} be a MC, and $\sigma ts \in \text{Paths}^*(\text{Ac}(\mathcal{M}))$. Additionally let $\Delta_{\sigma ts} \triangleq \{\rho \text{tail}(\pi) \mid \rho \in \text{GenTorr}(\sigma t), \pi \in \text{Paths}^*(\text{SCC}_t^+, t, \{s\})\}$. Then $\Delta_{\sigma ts} = \text{GenTorr}(\sigma ts)$.*

Proof. .

(\supseteq) Let $\rho_0 \rho_1 \cdots \rho_k \in \text{GenTorr}(\sigma ts)$ and n_t the lowest subindex of ρ such that $\rho_{n_t} = t$. Take $\rho \triangleq \rho_0 \rho_1 \cdots \rho_{n_t}$ and $\pi \triangleq \rho_{n_t} \cdots \rho_k$ (Note that $\rho_0 \rho_1 \cdots \rho_k = \rho \text{tail}(\pi)$). In order to prove that $\rho_0 \rho_1 \cdots \rho_k \in \Delta_{\sigma ts}$ we need to prove that

- (1) $\rho \in \text{GenTorr}(\sigma t)$, and
- (2) $\pi \in \text{Paths}^*(\text{SCC}_t^+, t, \{s\})$.

(1) Let f be such that $\sigma ts \preceq_f \rho_0 \rho_1 \cdots \rho_k$ and $f(|\sigma ts| - 1) = k$. Take $g : \{0, 1, \dots, |\sigma t| - 1\} \rightarrow \mathbb{N}$ be the restriction of f . It is easy to check that $\sigma t \preceq_g \rho$. Additionally $f(|\sigma t| - 1) = n_t$ (otherwise f would not satisfy the freshness property for $i = |\sigma t| - 1$). Then, by definition of g , we have $g(|\sigma t| - 1) = n_t$.

(2) It is clear that π is a path from t to s . Therefore we only have to show that every state of π is in SCC_t^+ . By definition of SCC_t^+ , $\pi_0 = t \in \text{SCC}_t^+$ and $s \in \text{SCC}_t^+$ since $s \in \text{Out}_{\text{SCC}_t^+}$. Additionally, since f satisfies inertia property we have that $\forall_{f(|\sigma t|-1) < j < f(|\sigma ts|-1)} : \rho_{f(|\sigma t|-1)} \sim \rho_j$, since $f(|\sigma t| - 1) = n_t$ and $\pi \triangleq \rho_{n_t} \cdots \rho_k$ we have $\forall_{0 < j < |\pi| - 1} : t \sim \pi_j$ proving that $\pi_j \in \text{SCC}_t^+$ for $j \in \{1, \dots, |\pi| - 2\}$.

(\subseteq) Take $\rho \in \text{GenTorr}(\sigma t)$ and $\text{tail}(\pi) \in \text{Paths}^*(\text{SCC}_t^+, t, \{s\})$. In order to prove that $\rho \text{tail}(\pi) \in \text{GenTorr}(\sigma ts)$ we need to show that there exists a function g such that:

- (1) $\sigma ts \preceq_g \rho \text{tail}(\pi)$,
- (2) $g(|\sigma ts| - 1) = |\rho \text{tail}(\pi)| - 1$.

Since $\rho \in \text{GenTorr}(\sigma t)$ we know that there exists f be such that $\sigma t \preceq_f \rho$ and $f(|\sigma t| - 1) = |\rho| - 1$. We define $g : \{0, 1, \dots, |\sigma ts| - 1\} \rightarrow \{0, 1, \dots, |\rho \text{tail}(\pi)| - 1\}$ by

$$g(i) \triangleq \begin{cases} f(i) & \text{if } i < |\sigma ts| - 1, \\ |\rho \text{tail}(\pi)| - 1 & \text{if } i = |\sigma ts| - 1. \end{cases}$$

(1) It is easy to check that $\sigma ts \sqsubseteq_g \rho \text{tail}(\pi)$. Now we will show that g satisfies Freshness and Inertia properties.

Freshness property: We need to show that for all $0 \leq i < |\sigma ts|$ we have $\forall_{0 \leq j < g(i)} : \rho \text{tail}(\pi)_{g(i)} \not\sim \rho \text{tail}(\pi)_j$. For the cases $i \in \{0, \dots, |\sigma t| - 1\}$ this holds since $\sigma t \preceq_f \rho$ and definition of g .

Consider $i = |\sigma ts| - 1$, in this case we have to prove $\forall_{0 \leq j < |\rho \text{tail}(\pi)| - 1} : \rho \text{tail}(\pi)_{|\rho \text{tail}(\pi)| - 1} \not\sim \rho \text{tail}(\pi)_j$ or equivalently $\forall_{0 \leq j < |\rho \text{tail}(\pi)| - 1} : s \not\sim \rho \text{tail}(\pi)_j$.

Case $j \in \{|\rho|, \dots, |\rho \text{tail}(\pi)| - 1\}$

Since $\pi \in \text{Paths}^*(\text{SCC}_t^+, t, \{s\})$ and $s \in \text{Out}_{\text{SCC}_t^+}$ we have $\forall_{0 \leq j < |\text{tail}(\pi)| - 1} : s \not\sim \text{tail}(\pi)_j$

Case $j \in \{0, \dots, |\rho| - 1\}$

Since $\sigma ts \in \text{Paths}^*(\text{Ac}(\mathcal{M}))$ and Observation 8.1 we have $\forall_{0 \leq j < |\sigma t| - 1} : s \not\sim \sigma t_j$. Additionally, $\sigma t \preceq_f \rho$, def. g , and Observation 8.2 imply $\forall_{0 \leq j < |\rho|} : s \not\sim \rho_j$ or equivalently $\forall_{0 \leq j < |\rho|} : s \not\sim \rho \text{tail}(\pi)_j$.

Inertia property: Since $\pi \in \text{Paths}^*(\text{SCC}_t^+, t, \{s\})$ we know that $\forall_{0 \leq j < |\pi| - 1} : t \sim \pi_j$ which implies that $\forall_{|\rho| - 1 < j < |\rho \text{tail}(\pi)| - 1} : \rho \text{tail}(\pi)_{|\rho| - 1} \sim \rho \text{tail}(\pi)_j$ or equivalently $\forall_{g(|\sigma| - 1) < j < g(|\sigma s| - 1)} : \rho \text{tail}(\pi)_{g(|\rho| - 1)} \sim \rho \text{tail}(\pi)_j$ showing that g satisfies the inertia property.

(2) Follows from the definition of g .

□

Theorem 5.10. *Let $\mathcal{M} = (S, s_0, \mathcal{P}, L)$ be a MC. Then for every rail $\sigma \in \text{Paths}^*(\text{Ac}(\mathcal{M}))$ we have*

$$\mathbf{Pr}_{\text{Ac}(\mathcal{M})}(\langle \sigma \rangle) = \mathbf{Pr}_{\mathcal{M}}(\text{Torr}(\sigma)).$$

Proof. By induction on the structure of σ .

Base Case: $\mathbf{Pr}_{\text{Ac}(\mathcal{M})}(\langle s_0 \rangle) = \mathbf{Pr}_{\text{Ac}(\mathcal{M})}(\text{Paths}(\text{Ac}(\mathcal{M}), s_0)) = 1 = \mathbf{Pr}_{\mathcal{M}}(\text{Paths}(\mathcal{M}, s_0)) = \mathbf{Pr}_{\mathcal{M}}(\text{Torr}(s_0))$.

Inductive Step: Let t be such that $\text{last}(\sigma) = t$. Suppose that $t \in S_{\text{Com}}$. Then

$$\begin{aligned} & \mathbf{Pr}_{\text{Ac}(\mathcal{M})}(\langle \sigma s \rangle) \\ &= \mathbf{Pr}_{\text{Ac}(\mathcal{M})}(\langle \sigma \rangle) \cdot \text{Ac}(\mathcal{P})(t, s) \\ &= \mathbf{Pr}_{\mathcal{M}}(\text{Torr}(\sigma)) \cdot \mathcal{P}(t, s) \\ & \quad \text{\{Inductive Hypothesis and definition of } \mathcal{P}\text{\}} \\ &= \mathbf{Pr}_{\mathcal{M}}(\bigsqcup_{\rho \in \text{GenTorr}(\sigma)} \langle \rho \rangle) \cdot \mathcal{P}(t, s) \quad \text{\{Lem. 5.9\}} \\ &= \sum_{\rho \in \text{GenTorr}(\sigma)} \mathbf{Pr}_{\mathcal{M}}(\langle \rho \rangle) \cdot \mathbf{Pr}_{\mathcal{M}}(\langle ts \rangle) \\ &= \sum_{\rho \in \text{GenTorr}(\sigma)} \mathbf{Pr}_{\mathcal{M}}(\langle \rho \text{tail}(ts) \rangle) \\ & \quad \text{\{Distributivity and } \text{last}(\rho) = t \text{ for all } \rho \in \text{GenTorr}(\sigma)\text{\}} \\ &= \sum_{\rho \in \text{GenTorr}(\sigma), \pi \in \text{Paths}(\text{SCC}_t^+, t, \{s\})} \mathbf{Pr}_{\mathcal{M}}(\langle \rho \text{tail}(\pi) \rangle) \\ &= \sum_{\rho \in \Delta_{\sigma s}} \mu_{\mathcal{M}}(\langle \rho \rangle) \quad \text{\{Dfn. } \Delta\text{\}} \\ &= \sum_{\rho \in \text{GenTorr}(\sigma s)} \mu_{\mathcal{M}}(\langle \rho \rangle) \quad \text{\{Lem. 8.3\}} \\ &= \mathbf{Pr}_{\mathcal{M}}(\bigsqcup_{\rho \in \text{GenTorr}(\sigma s)} \langle \rho \rangle) \\ &= \mathbf{Pr}_{\mathcal{M}}(\text{Torr}(\sigma s)) \quad \text{\{Lem. 5.9\}} \end{aligned}$$

Now suppose that $t \in S_{\text{Inp}}$. We denote by $\text{Ac}(\mathcal{P})$ to the probability matrix of $\text{Ac}(\mathcal{M})$, then

$$\begin{aligned} & \mathbf{Pr}_{\text{Ac}(\mathcal{M})}(\langle \sigma s \rangle) \\ &= \mathbf{Pr}_{\text{Ac}(\mathcal{M})}(\langle \sigma \rangle) \cdot \text{Ac}(\mathcal{P})(t, s) \\ &= \mathbf{Pr}_{\mathcal{M}}(\text{Torr}(\sigma)) \cdot \text{Ac}(\mathcal{P})(t, s) \quad \text{\{HI\}} \\ &= \mathbf{Pr}_{\mathcal{M}}(\bigsqcup_{\rho \in \text{GenTorr}(\sigma)} \langle \rho \rangle) \cdot \text{Ac}(\mathcal{P})(t, s) \quad \text{\{Lem. 5.9\}} \\ &= \left(\sum_{\rho \in \text{GenTorr}(\sigma)} \mathbf{Pr}_{\mathcal{M}}(\langle \rho \rangle) \right) \cdot \text{Ac}(\mathcal{P})(t, s) \\ &= \sum_{\rho \in \text{GenTorr}(\sigma)} \mathbf{Pr}_{\mathcal{M}}(\langle \rho \rangle) \cdot \mathbf{Pr}_{\mathcal{M}, t}(\text{Paths}(\text{SCC}_t^+, t, \{s\})) \\ & \quad \text{\{By definition of } \text{Ac}(\mathcal{P}) \text{ and distributivity\}} \\ &= \sum_{\rho \in \text{GenTorr}(\sigma)} \mathbf{Pr}_{\mathcal{M}}(\langle \rho \rangle) \cdot \sum_{\pi \in \text{Paths}^*(\text{SCC}_t^+, t, \{s\})} \mathbf{Pr}_{\mathcal{M}, t}(\langle \pi \rangle) \\ &= \sum_{\rho \in \text{GenTorr}(\sigma), \pi \in \text{Paths}^*(\text{SCC}_t^+, t, \{s\})} \mathbf{Pr}_{\mathcal{M}}(\langle \rho \text{tail}(\pi) \rangle) \quad \text{\{Dfn. } \mu\text{\}} \\ &= \sum_{\rho \in \Delta_{\sigma s}} \mathbf{Pr}_{\mathcal{M}}(\langle \rho \rangle) \quad \text{\{Dfn. } \Delta\text{\}} \\ &= \sum_{\rho \in \text{GenTorr}(\sigma s)} \mathbf{Pr}_{\mathcal{M}}(\langle \rho \rangle) \quad \text{\{Lem. 8.3\}} \\ &= \mathbf{Pr}_{\mathcal{M}}(\bigsqcup_{\rho \in \text{GenTorr}(\sigma s)} \langle \rho \rangle) \\ &= \mathbf{Pr}_{\mathcal{M}}(\text{Torr}(\sigma s)) \quad \text{\{Lem. 5.9\}} \end{aligned}$$

□