

A Counterexample Guided Abstraction-Refinement Framework for Markov Decision Processes

ROHIT CHADHA and MAHESH VISWANATHAN

Dept. of Computer Science, University of Illinois at Urbana-Champaign

The main challenge in using abstractions effectively, is to construct a suitable abstraction for the system being verified. One approach that tries to address this problem is that of *counterexample guided abstraction-refinement (CEGAR)*, wherein one starts with a coarse abstraction of the system, and progressively refines it, based on invalid counterexamples seen in prior model checking runs, until either an abstraction proves the correctness of the system or a valid counterexample is generated. While CEGAR has been successfully used in verifying non-probabilistic systems automatically, CEGAR has only recently been investigated in the context of probabilistic systems. The main issues that need to be tackled in order to extend the approach to probabilistic systems is a suitable notion of “counterexample”, algorithms to generate counterexamples, check their validity, and then automatically refine an abstraction based on an invalid counterexample. In this paper, we address these issues, and present a CEGAR framework for Markov Decision Processes.

Categories and Subject Descriptors: D.2.4 [Software Engineering]: Program Verification

General Terms: Verification

Additional Key Words and Phrases: Counterexamples, abstraction, refinement, model checking, Markov decision processes

1. INTRODUCTION

Abstraction is an important technique to combat *state space explosion*, wherein a smaller, abstract model that conservatively approximates the behaviors of the original (concrete) system is verified/model checked. The main challenge in applying this technique in practice, is in constructing such an abstract model. *Counterexample guided abstraction-refinement (CEGAR)* [Clarke et al. 2000] addresses this problem by constructing abstractions *automatically* by starting with a coarse abstraction of the system, and progressively refining it, based on invalid counterexamples seen in prior model checking runs, until either an abstraction proves the correctness of the system or a valid counterexample is generated.

While CEGAR has been successfully used in verifying non-probabilistic systems automatically, until recently, CEGAR has not been applied in the context of systems that exhibit probabilistic behavior. In order to extend this approach to probabilistic systems, one needs to identify a family of abstract models, develop a suitable notion of counterexamples, and design algorithms to produce counterexamples from

This work was partially supported by NSF grants CCF 0429639, CCF 0448178, and CNS 0509321. Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 20YY ACM 1529-3785/20YY/0700-0001 \$5.00

erroneous abstractions, check their validity in the original system, and (if needed) automatically refine an abstraction based on an invalid counterexample. In this paper we address these issues, and develop a CEGAR framework for systems described as *Markov Decision Processes (MDP)*.

Abstractions have been extensively studied in the context of probabilistic systems with definitions for good abstractions and specific families of abstractions being identified (see Section 6). In this paper, like Jonsson and Larsen [1991], D’Argenio et al. [2001] and D’Argenio et al. [2002], we use Markov decision processes to abstract other Markov decision processes. The abstraction will be defined by an equivalence relation (of finite index) on the states of the concrete system. The states of the abstract model will be the equivalence classes of this relation, and each abstract state will have transitions corresponding to the transitions of each of the concrete states in the equivalence class.

Crucial to extending the CEGAR approach to probabilistic systems is to come up with an appropriate notion of counterexamples. Clarke et al. [2002] have identified a clear set of metrics by which to evaluate any proposal for counterexamples. Counterexamples must satisfy three criteria: (a) counterexamples should serve as an “explanation” of why the (abstract) model violates the property, (b) must be rich enough to explain the violation of a large class of properties, and (c) must be simple and specific enough to identify bugs, and be amenable to efficient generation and analysis.

With regards to probabilistic systems there are three compelling proposals for counterexamples to consider. The first, originally proposed in [Han and Katoen 2007a] for DTMCs, is to consider counterexamples to be a multi-set of executions. This has been extended to CTMCs [Han and Katoen 2007b], and MDPs [Aljazzar and Leue 2007]. The second is to take counterexamples to be MDPs with a *tree-like* graph structure, a notion proposed by Clarke et al. [2002] for non-probabilistic systems and branching-time logics. The third and final notion, suggested in [Chatterjee et al. 2005; Hermanns et al. 2008], is to view general DTMCs (i.e., models without nondeterminism) as counterexamples. We show that all these proposals are expressively inadequate for our purposes. More precisely, we show that there are systems and properties that do not admit any counterexamples of the above special forms.

Having demonstrated the absence of counterexamples with special structure, we take the notion of counterexamples to simply be “small” MDPs that violate the property and are simulated by the abstract model. Formally, a counterexample for a system \mathcal{M} and property ψ_S will be a pair $(\mathcal{E}, \mathcal{R})$, where \mathcal{E} is a MDP violating the property ψ_S that is *simulated* by \mathcal{M} via the relation \mathcal{R} . The simulation relation has rarely been thought of as being formally part of the counterexample; requiring this addition does not change the asymptotic complexity of counterexample generation, since the simulation relation can be computed efficiently [Baier et al. 2000], and for the specific context of CEGAR, they are merely simple “injection functions”. However, as we shall point out, defining counterexamples formally in this manner makes the technical development of counterexample guided refinement cleaner (and is, in fact, implicitly assumed to be part of the counterexample, in the case of non-probabilistic systems).

One crucial property that counterexamples must exhibit is that they be amenable to efficient generation and analysis [Clarke et al. 2002]. We show that generating the *smallest* counterexample is NP-complete. Moreover it is unlikely to be efficiently approximable. However, in spite of these negative results, we show that there is a very simple polynomial time algorithm that generates a *minimal* counterexample; a minimal counterexample is a pair $(\mathcal{E}, \mathcal{R})$ such that if any state/transition of \mathcal{E} is removed, the resulting MDP no longer violates the property.

Intuitively, a counterexample is valid if the original system can exhibit the “behavior” captured by the counterexample. For non-probabilistic systems [Clarke et al. 2000; Clarke et al. 2002], a valid counterexample is not simply one that is simulated by the original system, even though simulation is the formal concept that expresses the notion of a system exhibiting a behavior. One requires that the original system simulate the counterexample, “in the same manner as the abstract system”. More precisely, if \mathcal{R} is the simulation relation that witnesses \mathcal{E} being simulated by the abstract system, then $(\mathcal{E}, \mathcal{R})$ is valid if the original system simulates \mathcal{E} through a simulation relation that is “contained within” \mathcal{R} . This is one technical reason why we consider the simulation relation to be part of the concept of a counterexample. Thus the algorithm for checking validity is the same as the algorithm for checking simulations between MDPs [Baier et al. 2000; Zhang et al. 2007] except that we have to ensure that the witnessing simulation be “contained within \mathcal{R} ”. However, because of the special nature of counterexamples, better bounds on the running time of the algorithm can be obtained.

Finally, we outline how the abstraction can be automatically refined. Once again the algorithm is a natural generalization of the refinement algorithm in the non-probabilistic case, though it is different from the refinement algorithms proposed in [Chatterjee et al. 2005; Hermanns et al. 2008]; detailed comparison can be found in Section 6. We also state and prove precisely what the refinement algorithm achieves.

1.1 Our Contributions

We now detail our main technical contributions, roughly in the order in which they appear in the paper.

(1) For MDPs, we identify safety and liveness fragments of PCTL. Our fragment is syntactically different than that presented in [Desharnais 1999b; Baier et al. 2005] for DTMCs. Though the two presentations are semantically the same for DTMCs, they behave differently for MDPs.

(2) We demonstrate the expressive inadequacy of all relevant proposals for counterexamples for probabilistic systems, thus demonstrating that counterexamples with special graph structures are unlikely to be rich enough for the safety fragment of PCTL.

(3) We present formal definitions of counterexamples, their validity and consistency, and the notion of good counterexample-guided refinements. We distill a precise statement of what the CEGAR-approach achieves in a single abstraction-refinement step. Thus, we generalize concepts that have been hitherto only defined for “path-like” structures [Clarke et al. 2000; Clarke et al. 2002; Han and Katoen 2007a; 2007b; Aljazzar and Leue 2007; Hermanns et al. 2008] to general graph-like

structures¹, and for the first time formally articulate, what is accomplished in a single abstraction-refinement step.

(4) We present algorithmic solutions to all the computational problems that arise in the CEGAR loop: we give lower bounds as well as upper bounds for counterexample generation, and algorithms to check validity and to refine an abstraction.

(5) A sub-logic of our safe-PCTL, which we call weak safety, does indeed admit counterexamples that have a tree-like structure. For this case, we present an on-the-fly algorithm to unroll the minimal counterexample that we generate and check validity. This algorithm may perform better than the algorithm based on checking simulation for some examples in practice.

Though our primary contributions are to clarify the definitions and concepts needed to carry out CEGAR in the context of probabilistic systems, our effort also sheds light on implicit assumptions made by the CEGAR approach for non-probabilistic systems.

1.2 Outline of the Paper

The rest of the paper is organized as follows. We recall some definitions and notations in Section 2. We also present safety and liveness fragments of PCTL for MDPs in Section 2. We discuss various proposals of counterexamples for MDPs in Section 3, and also present our definition of counterexamples along with algorithmic aspects of counterexample generation. We recall the definition of abstractions based on equivalences in Section 4. We present the definitions of validity and consistency of abstract counterexamples and good counterexample-guided refinement, as well as the algorithms to check validity and refine abstractions in Section 5. Finally, related work is discussed in Section 6.

2. PRELIMINARIES

The paper assumes familiarity with basic probability theory, discrete time Markov chains, Markov decision processes, and the model checking of these models against specifications written in PCTL; the background material can be found in [Rutten et al. 2004]. This section is primarily intended to introduce notation, and to introduce and remind the reader of results that the paper will rely on.

2.1 Relations and Functions.

We assume that the reader is familiar with the basic definitions of relation and functions. We will primarily be interested in binary relations. We will use \mathcal{R}, S, T, \dots to range over relations and f, g, h, \dots to range over functions. We introduce here some notations that will be useful.

Given a set A , we will denote its power-set by 2^A . For a finite set A , the number of elements in A will be denoted by $|A|$.

The identity function on a set A will be often denoted by id_A . Given a function $f : A \rightarrow B$ and set $A' \subseteq A$, the restriction of f to A' will be denoted by $f|_{A'}$.

¹Even when a counterexample is not formally a path, as in [Clarke et al. 2002] and [Hermanns et al. 2008], it is viewed as a collection of paths and simple cycles, and all concepts are defined for the case when the cycles have been unrolled a finite number of times.

For a binary relation $\mathcal{R} \subseteq A \times B$ we will often write $a \mathcal{R} b$ to mean $(a, b) \in \mathcal{R}$. Also, given $a \in A$ we will denote the set $\{b \in B \mid a \mathcal{R} b\}$ by $\mathcal{R}(a)$. Please note \mathcal{R} is uniquely determined by the collection $\{\mathcal{R}(a) \mid a \in A\}$. A binary relation $\mathcal{R}_1 \subseteq A \times B$ is said to be *finer* than $\mathcal{R}_2 \subseteq A \times B$ if $\mathcal{R}_1 \subseteq \mathcal{R}_2$. The composition of two binary relations $\mathcal{R}_1 \subseteq A \times B$ and $\mathcal{R}_2 \subseteq B \times C$, denoted by $\mathcal{R}_2 \circ \mathcal{R}_1$, is the relation $\{(a, c) \mid \exists b \in B. a \mathcal{R}_1 b \text{ and } b \mathcal{R}_2 c\} \subseteq A \times C$.

We say that a binary relation $\mathcal{R} \subseteq A \times B$ is *total* if for all $a \in A$ there is a $b \in B$ such that $a \mathcal{R} b$. We say that a binary relation $\mathcal{R} \subseteq A \times B$ is *functional* if for all $a \in A$ there is at most one $b \in B$ such that $a \mathcal{R} b$. There is a close correspondence between functions and total, functional relations: for any function $f : A \rightarrow B$, the relation $\{(a, f(a)) \mid a \in A\}$ is a total and functional binary relation. Vice-versa, one can construct a unique function from a given total and functional binary relation. We will denote the total and functional relation given by a function f by rel_f .

A *preorder* on a set A is a binary relation that is reflexive and transitive. An equivalence relation on a set A is a preorder which is also symmetric. The *equivalence class* of an element $a \in A$ with respect to an equivalence relation \equiv , will be denoted by $[a]_{\equiv}$; when the equivalence relation \equiv is clear from the context we will drop the subscript \equiv .

2.2 DTMC and MDP

Sequences. For a set X , we will denote the set of non-empty finite sequences of elements of X by X^+ . The set X^ω will be the set of all countably infinite sequences of elements of X . For a finite sequence $\eta \in X^+$ and an infinite sequence $\alpha \in X^\omega$ we denote by $\eta\alpha \in X^\omega$ the sequence obtained by concatenating the sequences η and α in order.

Basic Probability Theory. Let X be a set and E be a set of subsets of X . We say that E is a σ -algebra on X if E contains the empty set, is closed under complementation and also under countable unions. Let F be a set of subsets of X . A σ -algebra generated by F is the smallest σ -algebra that contains F . A (sub)-probability space is a triple (X, E, μ) where E is a σ -algebra on X and μ is a (sub)-probability function [Papoulis and Pillai 2002] on E .

For (finite or countable) set X with σ -field 2^X , the collection all sub-probability measures (i.e., where measure of $X \leq 1$) will be denoted by $\text{Prob}_{\leq 1}(X)$. For $\mu \in \text{Prob}_{\leq 1}(X)$ and $A \subseteq X$, $\mu(A)$ denotes the measure of set A .

Kripke structures. A *Kripke structure* over a set of propositions AP , is formally a tuple $\mathcal{K} = (\mathbb{Q}, q_{\mathcal{I}}, \rightarrow, \text{L})$ where \mathbb{Q} is a set of states, $q_{\mathcal{I}} \in \mathbb{Q}$ is the initial state, $\rightarrow \subseteq \mathbb{Q} \times \mathbb{Q}$ is the transition function, and $\text{L} : \mathbb{Q} \rightarrow 2^{\text{AP}}$ is a labeling function that labels each state with the set of propositions true in it. DTMC and MDP are generalizations of Kripke structures where transitions are replaced by probabilistic transitions.

Discrete Time Markov Chains. A *discrete time Markov chain* (DTMC) over a set of propositions AP , is formally a tuple $\mathcal{M} = (\mathbb{Q}, q_{\mathcal{I}}, \delta, \text{L})$ where \mathbb{Q} is a (finite or countable) set of states, $q_{\mathcal{I}} \in \mathbb{Q}$ is the initial state, $\delta : \mathbb{Q} \rightarrow \text{Prob}_{\leq 1}(\mathbb{Q})$ is the transition function, and $\text{L} : \mathbb{Q} \rightarrow 2^{\text{AP}}$ is a labeling function that labels each state

with the set of propositions true in it. A DTMC is said to be *finite* if the set \mathbf{Q} is finite. Unless otherwise explicitly stated, DTMCs in this paper will be assumed to be finite.

DTMCs generate a probability space (X, E, μ) as follows [Kemeny and Snell 1976; Rutten et al. 2004]. First of all, we pick a new state $\perp \notin \mathbf{Q}$. The set X is $(\mathbf{Q} \cup \{\perp\})^\omega$ and will henceforth be called the set of *paths*. Let E be the σ -algebra generated by the set $\{C_\eta \mid \eta \in (\mathbf{Q} \cup \{\perp\})^+\}$ where $C_\eta = \{\eta\alpha \mid \alpha \in (\mathbf{Q} \cup \{\perp\})^+\}$. The measure μ is the unique measure such that for each $\eta = q_0q_1 \cdots q_n \in Q^+$, $\mu(C_\eta)$ is

- 0 if $q_0 \neq q_{\mathcal{I}}$
- 1 if $n = 0$ and $q_0 = q_{\mathcal{I}}$, and
- $\prod_{0 \leq i < n} pr(q_i, q_{i+1})$ otherwise, where $pr(q_i, q_{i+1})$ is
 - $\delta(q_i)(q_{i+1})$ if $q_i, q_{i+1} \in \mathbf{Q}$,
 - $1 - \delta(q_i)(\mathbf{Q})$ if $q_i \in \mathbf{Q}$ and $q_{i+1} = \perp$,
 - 1 if $q_i = q_{i+1} = \perp$, and
 - 0 if $q_i = \perp$ and $q_{i+1} \in \mathbf{Q}$.

We omit the details of the construction and refer the reader to [Rutten et al. 2004] for the complete construction.

Remark 2.1. We point out here that sometimes instead of the states, the labeling function labels the transitions [Segala and Lynch 1995; Desharnais et al. 2000].

Markov Decision Processes. A *finite Markov decision process* (MDP) over a set of propositions AP , is formally a tuple $\mathcal{M} = (\mathbf{Q}, q_{\mathcal{I}}, \delta, \mathbf{L})$ where $\mathbf{Q}, q_{\mathcal{I}}, \mathbf{L}$ are as in the case for finite DTMCs, and $\delta : \mathbf{Q} \rightarrow 2^{\text{Prob}_{\leq 1}(\mathbf{Q})}$ maps each state to a *finite* non-empty collection of sub-probability measures. We will sometimes say that there is no transition out of $q \in \mathbf{Q}$ if $\delta(q)$ consists of exactly one sub-probability measure $\mathbf{0}$ which assigns 0 to all states in \mathbf{Q} . For this paper, we will assume that for every q, q' and every $\mu \in \delta(q)$, $\mu(q')$ is a rational number. From now on, we will explicitly drop the qualifier “finite” for MDPs. Note that a finite DTMC can be viewed as a special case of a MDP, namely one in which the $\delta(q)$ consists of exactly one element for each state q .

We assume that the reader is familiar with the concept of *schedulers*. A scheduler resolves nondeterminism. Formally, a scheduler is a function $\mathcal{S} : Q^+ \rightarrow \text{Prob}_{\leq 1}(\mathbf{Q})$ such that for any $q_0q_1 \cdots q_n \in Q^+$, $\mathcal{S}(q_0q_1 \cdots q_n) \in \delta(q_n)$. Given a scheduler and a MDP \mathcal{M} , we get a probability space $(X, E, \mu_{\mathcal{S}})$ as follows. As in the case of DTMCs, the set X is $(\mathbf{Q} \cup \{\perp\})^\omega$ where $\perp \notin \mathbf{Q}$. E is the σ -algebra generated by the set $\{C_\eta \mid \eta \in (\mathbf{Q} \cup \{\perp\})^+\}$ where $C_\eta = \{\eta\alpha \mid \alpha \in (\mathbf{Q} \cup \{\perp\})^+\}$. The measure $\mu_{\mathcal{S}}$ is the unique measure such that for each $\eta = q_0q_1 \cdots q_n \in Q^+$, $\mu_{\mathcal{S}}(C_\eta)$ is

- 0 if $q_0 \neq q_{\mathcal{I}}$
- 1 if $n = 0$ and $q_0 = q_{\mathcal{I}}$, and
- $\prod_{0 \leq i < n} pr_{i,i+1}$ otherwise, where $pr_{i,i+1}$ is
 - $\mathcal{S}(q_0q_1 \cdots q_i)(q_{i+1})$ if $q_i, q_{i+1} \in \mathbf{Q}$,
 - $1 - \mathcal{S}(q_0q_1 \cdots q_i)(\mathbf{Q})$ if $q_i \in \mathbf{Q}$ and $q_{i+1} = \perp$,
 - 1 if $q_i = q_{i+1} = \perp$, and
 - 0 if $q_i = \perp$ and $q_{i+1} \in \mathbf{Q}$.

A scheduler \mathcal{S} is said to be *memoryless* if for each $\eta = q_0q_1 \cdots q_n$ and $\eta' = q'_0q'_1 \cdots q'_l$, $\mathcal{S}(\eta) = \mathcal{S}(\eta')$ whenever $q_n = q'_l$. So, a memoryless scheduler is completely specified by a function $tr_{\mathcal{S}} : \mathbb{Q} \rightarrow \text{Prob}_{\leq 1}(\mathbb{Q})$, where $tr_{\mathcal{S}}(q)$ gives the transition $\mathcal{S}(q_0q_1 \cdots q_n)$ whenever $q_n = q$. We call $tr_{\mathcal{S}}(q)$, the *transition out of q* .

Remark 2.2. In the presence of scheduler \mathcal{S} the MDP can also be thought of as a (countable) DTMC, $\mathcal{M}^{\mathcal{S}} = \{\mathbb{Q}^+, q_{\mathcal{I}}, \delta^{\mathcal{S}}, L^{\mathcal{S}}\}$ where

- $L^{\mathcal{S}}(q_0q_1 \cdots q_n) = L(q_n)$ and
- $\delta^{\mathcal{S}}(q_0q_1 \cdots q_n)$ is the sub-probability measure $\mu^{\mathcal{S}} \in \text{Prob}_{\leq 1}(\mathbb{Q}^+)$ such that for each $q'_0q'_1 \cdots q'_l \in \mathbb{Q}^+$, $\mu^{\mathcal{S}}(q'_0q'_1 \cdots q'_l)$ is
 - $\mathcal{S}(q_0q_1 \cdots q_n)(q'_l)$ if $l = n + 1$ and $q_i = q'_i$ for each $0 \leq i \leq n$, and
 - 0 otherwise.

In the presence of a memoryless scheduler \mathcal{S} , the resulting DTMC $\mathcal{M}^{\mathcal{S}}$ is *bisimilar* to a finite DTMC which has the same set of states as \mathcal{M} , the same initial state and the same labeling function, while the transition out of a state q is the one given by the memoryless scheduler \mathcal{S} . We will henceforth confuse this finite DTMC with $\mathcal{M}^{\mathcal{S}}$.

Suppose there are at most k nondeterministic choices from any state in \mathcal{M} . For some ordering of the nondeterministic choices out of each states, the *labeled underlying graph* of a MDP is the directed graph $G = (\mathbb{Q}, \{E_i\}_{i=1}^k)$, where $(q_1, q_2) \in E_i$ iff $\mu(q_2) > 0$, where μ is the i th choice out of q_1 ; we will denote the labeled underlying graph of \mathcal{M} by $G_{\ell}(\mathcal{M})$. The *unlabeled underlying graph* will be $G' = (\mathbb{Q}, \cup_{i=1}^k E_i)$ and is denoted by $G(\mathcal{M})$. The following notation will be useful.

Notation 2.3. Given a MDP $\mathcal{M} = (\mathbb{Q}, q_{\mathcal{I}}, \delta, L)$, a state $q \in \mathbb{Q}$ and a transition $\mu \in \delta(q)$, we say that $\text{post}(\mu, q) = \{q' \in \mathbb{Q} \mid \mu(q') > 0.\}$

PCTL. We recall here the logic PCTL [Rutten et al. 2004] briefly. The formulas of the logic PCTL are generated over a set of propositions AP as follows:

$$\psi := \text{tt} \mid \text{ff} \mid P \mid (\neg\psi) \mid (\psi \wedge \psi) \mid (\psi \vee \psi) \mid \mathcal{P}_{\triangleleft p}(\mathbf{X}\psi) \mid \mathcal{P}_{\triangleleft p}(\psi \mathbf{U} \psi)$$

where $P \in \text{AP}$, $p \in [0, 1]$ is a rational number and $\triangleleft \in \{<, \leq\}$. The logic is interpreted over DTMCs and MDPs and the details of the interpretation are given in an Electronic Appendix, for the sake of the flow of the paper. Given a MDP \mathcal{M} and a state q of \mathcal{M} , we say $q \Vdash_{\mathcal{M}} \psi$ if q satisfies the formula ψ . We will drop \mathcal{M} when clear from the context. We will say that $\mathcal{M} \Vdash \psi$ if the initial state of \mathcal{M} satisfies the formula. The model checking problem for MDPs and PCTL is known to be in polynomial time [Bianco and de Alfaro 1995]. We also point out PCTL is also defined for MDPs in which the labels are on the transitions instead of the state, as in [Segala and Lynch 1994]. In [Segala and Lynch 1994], transition labels instead of propositions are used as formulas of PCTL, and a state q satisfies a transition label a if all the actions out of q are labeled a .

Unrolling of a MDP. Given a MDP $\mathcal{M} = (\mathbb{Q}, q_{\mathcal{I}}, \delta, L)$, natural number $k \geq 0$, and $q \in \mathbb{Q}$ we will define a MDP $\mathcal{M}_k^q = (\mathbb{Q}_k^q, (q, k), \delta_k^q, L_k^q)$ obtained by unrolling the underlying labeled graph of \mathcal{M} upto depth k . Formally, $\mathcal{M}_k^q = (\mathbb{Q}_k^q, (q, k), \delta_k^q, L_k^q)$, the k -th *unrolling of \mathcal{M} rooted at q* is defined as follows.

- $\mathbb{Q}_k^q = \{(q, k)\} \cup (\mathbb{Q} \times \{j \in \mathbb{N} \mid 0 \leq j < k\})$.
- For all $(q', j) \in \mathbb{Q}_k^q$, $L((q', j)) = L(q')$.
- For all $(q', j) \in \mathbb{Q}_k^q$, $\delta_k^q((q', j)) = \{\mu^j \mid \mu \in \delta(q')\}$ where μ^j is defined as–
 - (1) for $j = 0$, $\mu^j(q'') = 0$ for all $q'' \in \mathbb{Q}_k^q$, and
 - (2) for $0 < j < k$, $\mu^j(q'') = \mu(q_1)$ if $q'' = (q_1, j - 1)$ for some $q_1 \in \mathbb{Q}$ and $\mu^j(q'') = 0$ otherwise.

Please note that the underlying unlabeled graph of \mathcal{M}_k^q is (directed) acyclic.

Direct Sum of MDPs. Given MDPs $\mathcal{M} = (\mathbb{Q}, q_{\mathcal{I}}, \delta, L)$ and $\mathcal{M}' = (\mathbb{Q}', q'_{\mathcal{I}}, \delta', L')$ over the set of propositions AP, let $\mathbb{Q} + \mathbb{Q}' = \mathbb{Q} \times \{0\} \cup \mathbb{Q}' \times \{1\}$ be the disjoint sum of \mathbb{Q} and \mathbb{Q}' . Now, define $\delta + \delta' : \mathbb{Q} + \mathbb{Q}' \rightarrow \text{Prob}_{\leq 1}(\mathbb{Q} + \mathbb{Q}')$ and $L + L' : \mathbb{Q} + \mathbb{Q}' \rightarrow \text{AP}$ as follows. For all $q \in \mathbb{Q}$ and $q' \in \mathbb{Q}'$,

- $(\delta + \delta')((q, 0)) = \{\mu \times \{0\} \mid \mu \in \delta(q)\}$ and $(\delta + \delta')((q', 1)) = \{\mu' \times \{1\} \mid \mu' \in \delta'(q')\}$ where $\mu \times \{0\}$ and $\mu' \times \{1\}$ are defined as follows.
 - $\mu \times \{0\}(q_1, 0) = \mu(q_1)$ and $\mu \times \{0\}(q'_1, 1) = 0$ for all $q_1 \in \mathbb{Q}$ and $q'_1 \in \mathbb{Q}'$.
 - $\mu' \times \{1\}(q_1, 0) = 0$ and $\mu' \times \{1\}(q'_1, 1) = \mu'(q'_1)$ for all $q_1 \in \mathbb{Q}$ and $q'_1 \in \mathbb{Q}'$.
- $(L + L')(q, 0) = L(q)$ and $(L + L')(q', 1) = L'(q')$.

Now given $q \in \mathbb{Q} + \mathbb{Q}'$, the MDP $(\mathcal{M} + \mathcal{M}')_q = (\mathbb{Q} + \mathbb{Q}', q, \delta + \delta', L + L')$ is said to be the *direct sum* of \mathcal{M} and \mathcal{M}' with q as the initial state.

Remark 2.4. MDPs $\mathcal{M} = (\mathbb{Q}, q_{\mathcal{I}}, \delta, L)$ and $\mathcal{M}' = (\mathbb{Q}', q'_{\mathcal{I}}, \delta', L')$ are said to be *disjoint* if $\mathbb{Q} \cap \mathbb{Q}' = \emptyset$. If MDPs $\mathcal{M} = (\mathbb{Q}, q_{\mathcal{I}}, \delta, L)$ and $\mathcal{M}' = (\mathbb{Q}', q'_{\mathcal{I}}, \delta', L')$ are disjoint, then $\mathbb{Q} + \mathbb{Q}'$ can be taken to be the union $\mathbb{Q} \cup \mathbb{Q}'$. In such cases, we will confuse $(q, 0)$ with q , $(q', 1)$ with q' , $\mu \times \{0\}$ with μ and $\mu' \times \{1\}$ with μ' (with the understanding that $\mu \in \delta(q)$ takes value 0 on any $q' \in \mathbb{Q}'$ and $\mu' \in \delta'(q')$ takes value 0 on any $q \in \mathbb{Q}$).

2.3 Simulation

Given a binary relation \mathcal{R} on the set of states \mathbb{Q} , a set $A \subseteq \mathbb{Q}$, is said to be \mathcal{R} -closed if the set $\mathcal{R}(A) = \{t \mid \exists q \in A, q \mathcal{R} t\}$ is the same as A . For two sub-probability measures $\mu, \mu' \in \text{Prob}_{\leq 1}(\mathbb{Q})$, we say μ' *simulates* μ with respect to a preorder \mathcal{R} (denoted as $\mu \preceq_{\mathcal{R}} \mu'$) iff for every \mathcal{R} -closed set A , $\mu(A) \leq \mu'(A)$. For a MDP $\mathcal{M} = (\mathbb{Q}, q_{\mathcal{I}}, \delta, L)$, a preorder \mathcal{R} on \mathbb{Q} is said to be a *simulation relation*² if for every $q \mathcal{R} q'$, we have that $L(q) = L(q')$ and for every $\mu \in \delta(q)$ there is a $\mu' \in \delta(q')$ such that $\mu \preceq_{\mathcal{R}} \mu'$. We say that $q \preceq q'$ if there is a simulation relation \mathcal{R} such that $q \mathcal{R} q'$.

Given an equivalence relation \equiv on the set of states \mathbb{Q} , and two sub-probability measures $\mu, \mu' \in \text{Prob}_{\leq 1}(\mathbb{Q})$ we say that μ *is equivalent to* μ' with respect to \equiv (denoted as $\mu \approx_{\equiv} \mu'$) iff for every \equiv -closed set A , $\mu(A) = \mu'(A)$. For a MDP $\mathcal{M} = (\mathbb{Q}, q_{\mathcal{I}}, \delta, L)$, an equivalence \equiv on \mathbb{Q} is said to be a *bisimulation* if for every $q \equiv q'$, we have that $L(q) = L(q')$ and for every $\mu \in \delta(q)$ there is a $\mu' \in \delta(q')$ such

²It is possible to require only that $L(q) \subseteq L(q')$ instead of $L(q) = L(q')$ in the definition of simulation. The results and proofs of the paper could be easily adapted for this definition. One has to modify the definition of safety and liveness fragments of PCTL appropriately.

that $\mu \approx_{\equiv} \mu'$. We say that $q \approx q'$ if there is a bisimulation relation \equiv such that $q \equiv q'$.

Remark 2.5. The ordering on probability measures used in the definition of simulation presented in [Jonsson and Larsen 1991; Segala and Lynch 1994; Baier et al. 2005] is based on *weight functions*. However, the definition presented here, was originally proposed in [Desharnais 1999a] and shown to be equivalent in [Desharnais 1999a; Segala 2006].

We say that MDP $\mathcal{M} = (\mathbb{Q}, q_{\mathcal{I}}, \delta, \mathbb{L})$ is simulated by $\mathcal{M}' = (\mathbb{Q}', q'_{\mathcal{I}}, \delta', \mathbb{L}')$ (denoted by $\mathcal{M} \preceq \mathcal{M}'$) if there is a simulation relation \mathcal{R} on the direct sum of \mathcal{M} and \mathcal{M}' (with any initial state) such that $(q_{\mathcal{I}}, 0) \mathcal{R} (q'_{\mathcal{I}}, 1)$. The MDP \mathcal{M} is said to be bisimilar to \mathcal{M}' (denoted by $\mathcal{M} \approx \mathcal{M}'$) if there is a bisimulation relation \equiv on the direct sum of \mathcal{M} and \mathcal{M}' (with any initial state) such that $(q_{\mathcal{I}}, 0) \equiv (q'_{\mathcal{I}}, 1)$.

As an example of simulations, we have that every MDP $\mathcal{M} = (\mathbb{Q}, q_{\mathcal{I}}, \delta, \mathbb{L})$ simulates its k -th unrolling. Furthermore, we also have that if $k \leq k'$ then the k' -th unrolling simulates the k -th unrolling.

PROPOSITION 2.6. *Given a MDP \mathcal{M} with initial state $q_{\mathcal{I}}$ and natural numbers $k, k' \geq 0$ such that $k \leq k'$. Let $\mathcal{M}_k^{q_{\mathcal{I}}}$ and $\mathcal{M}_{k'}^{q_{\mathcal{I}}}$ be the k -th and k' -th unrolling of \mathcal{M} rooted at $q_{\mathcal{I}}$ respectively. Then $\mathcal{M}_k^{q_{\mathcal{I}}} \preceq \mathcal{M}$ and $\mathcal{M}_k^{q_{\mathcal{I}}} \preceq \mathcal{M}_{k'}^{q_{\mathcal{I}}}$.*

Simulation between disjoint MDPs. We will be especially interested in simulation between disjoint MDPs (in which case we can just take the union of state spaces of the MDPs as the state space of the direct sum). The simulations will also take a certain form which we will call the *canonical form* for our purposes. In order to define this precisely, recall that for any set A , id_A is the identity function on A and that rel_{id_A} is the relation $\{(a, a) \mid a \in A\}$.

Definition 2.7. Given disjoint MDPs $\mathcal{M} = (\mathbb{Q}, q_{\mathcal{I}}, \delta, \mathbb{L})$ and $\mathcal{M}' = (\mathbb{Q}', q'_{\mathcal{I}}, \delta', \mathbb{L}')$, we say that a simulation relation $\mathcal{R} \subseteq (\mathbb{Q} + \mathbb{Q}') \times (\mathbb{Q} + \mathbb{Q}')$ on the direct sum of MDPs \mathbb{Q} and \mathbb{Q}' (with any initial state) is in *canonical form* if there exists a relation $\mathcal{R}_1 \subseteq \mathbb{Q} \times \mathbb{Q}'$ such that $\mathcal{R} = \text{rel}_{id_{\mathbb{Q}}} \cup \mathcal{R}_1 \cup \text{rel}_{id_{\mathbb{Q}'}}$.

The following proposition states that any simulation contains a largest canonical simulation and hence canonical simulations are sufficient for reasoning about simulation between disjoint MDPs. The proof of the proposition has been moved to an Electronic Appendix for the sake of the flow of the paper.

PROPOSITION 2.8. *Given disjoint MDPs $\mathcal{M} = (\mathbb{Q}, q_{\mathcal{I}}, \delta, \mathbb{L})$ and $\mathcal{M}' = (\mathbb{Q}', q'_{\mathcal{I}}, \delta', \mathbb{L}')$, let $\mathcal{R} \subseteq (\mathbb{Q} + \mathbb{Q}') \times (\mathbb{Q} + \mathbb{Q}')$ be a simulation relation on the direct sum of \mathbb{Q} and \mathbb{Q}' . Let $\mathcal{R}_1 = \mathcal{R} \cap (\mathbb{Q} \times \mathbb{Q}')$. Then the relation $\mathcal{R}_0 = \text{rel}_{id_{\mathbb{Q}}} \cup \mathcal{R}_1 \cup \text{rel}_{id_{\mathbb{Q}'}}$ is a simulation relation.*

Notation 2.9. In order to avoid clutter, we will often denote a simulation $\text{rel}_{id_{\mathbb{Q}}} \cup \mathcal{R}_1 \cup \text{rel}_{id_{\mathbb{Q}'}}$ in the canonical form by just \mathcal{R}_1 as in the following proposition. Further, if $\mathcal{R} \subseteq \mathbb{Q} \times \mathbb{Q}'$ is a canonical simulation, then we say that any set $A \subseteq \mathbb{Q} \cup \mathbb{Q}'$ is \mathcal{R} -closed iff it is $\text{rel}_{id_{\mathbb{Q}}} \cup \mathcal{R} \cup \text{rel}_{id_{\mathbb{Q}'}}$ -closed.

PROPOSITION 2.10. *Given disjoint MDPs $\mathcal{M}_0 = (\mathbb{Q}_0, q_0, \delta_0, \mathbb{L}_0)$, $\mathcal{M}_1 = (\mathbb{Q}_1, q_1, \delta_1, \mathbb{L}_1)$ and $\mathcal{M}_2 = (\mathbb{Q}_2, q_2, \delta_2, \mathbb{L}_2)$, if $\mathcal{R}_{01} \subseteq \mathbb{Q}_0 \times \mathbb{Q}_1$ and $\mathcal{R}_{12} \subseteq \mathbb{Q}_1 \times \mathbb{Q}_2$ are canonical simulations then the relation $\mathcal{R}_{02} = \mathcal{R}_{12} \circ \mathcal{R}_{01} \subseteq \mathbb{Q}_0 \times \mathbb{Q}_2$ is a canonical simulation.*

2.4 PCTL-safety and PCTL-liveness.

We define a fragment of PCTL which we call the *safety fragment*. The *safety fragment* of PCTL (over a set of propositions AP) is defined in conjunction with the *liveness fragment* as follows.

$$\begin{aligned} \psi_S &:= \text{tt} \parallel \text{ff} \parallel P \parallel (\neg P) \parallel (\psi_S \wedge \psi_S) \parallel (\psi_S \vee \psi_S) \parallel \mathcal{P}_{\triangleleft p}(\mathbf{X} \psi_L) \parallel \mathcal{P}_{\triangleleft p}(\psi_L \mathbf{U} \psi_L) \\ \psi_L &:= \text{tt} \parallel \text{ff} \parallel P \parallel (\neg P) \parallel (\psi_L \wedge \psi_L) \parallel (\psi_L \vee \psi_L) \parallel (\neg \mathcal{P}_{\triangleleft p}(\mathbf{X} \psi_L)) \parallel (\neg \mathcal{P}_{\triangleleft p}(\psi_L \mathbf{U} \psi_L)) \end{aligned}$$

where $P \in \text{AP}$, $p \in [0, 1]$ is a rational number and $\triangleleft \in \{<, \leq\}$.

Note that for any safety formula ψ_S there exists a liveness formula ψ_L such that for state q of a MDP \mathcal{M} , $q \Vdash_{\mathcal{M}} \psi_S$ iff $q \not\vdash_{\mathcal{M}} \psi_L$. Restricting \triangleleft to be \leq in the above grammar, yields the *strict liveness* and *weak safety* fragments of PCTL. Finally recall that $\diamond\psi$ is an abbreviation for $\text{tt} \mathbf{U} \psi$.

There is a close correspondence between simulation and the liveness and safety fragments of PCTL—simulation preserves liveness and reflects safety. The proof of this correspondence has also been moved to an Electronic Appendix.

LEMMA 2.11. *Let $\mathcal{M} = (\mathbb{Q}, q_{\mathcal{I}}, \delta, \mathbb{L})$ be a MDP. For any states $q, q' \in \mathbb{Q}$, $q \preceq q'$ implies that for every liveness formula ψ_L , if $q \Vdash_{\mathcal{M}} \psi_L$ then $q' \Vdash_{\mathcal{M}} \psi_L$ and that for every safety formula ψ_S , if $q' \Vdash_{\mathcal{M}} \psi_S$ then $q \Vdash_{\mathcal{M}} \psi_S$.*

Remark 2.12.

(1) The fragment presented here is syntactically different than the safety and liveness fragments presented in [Desharnais 1999b; Baier et al. 2005] for DTMCs, where the safety fragment and the liveness fragment takes the following form.

$$\begin{aligned} \psi_S &:= \text{tt} \parallel \text{ff} \parallel P \parallel (\neg P) \parallel (\psi_S \wedge \psi_S) \parallel (\psi_S \vee \psi_S) \parallel \mathcal{P}_{\triangleleft p}(\mathbf{X} \psi_L) \parallel \mathcal{P}_{\triangleleft p}(\psi_L \mathbf{U} \psi_L) \\ \psi_L &:= \text{tt} \parallel \text{ff} \parallel P \parallel (\neg P) \parallel (\psi_L \wedge \psi_L) \parallel (\psi_L \vee \psi_L) \parallel \mathcal{P}_{\triangleright p}(\mathbf{X} \psi_L) \parallel \mathcal{P}_{\triangleright p}(\psi_L \mathbf{U} \psi_L) \end{aligned}$$


where $P \in \text{AP}$, $p \in [0, 1]$ is a rational number and $\triangleleft \in \{<, \leq\}$ and $\triangleright \in \{>, \geq\}$.

The two presentations have the same semantics for DTMCs, but behave differently for MDPs (see Example 2.13 below). As far as we know, the safety fragment of PCTL for state-labeled MDPs has not been discussed previously in the literature.

(2) For MDPs in which the labels are on transitions and not on states, Segala and Lynch [1994] also have a preservation result for the corresponding PCTL. However, their result crucially relies on the fact that all the probabilistic transitions have the full measure 1 and their preservation result does not hold otherwise.

(3) Please note that, unlike the case of DTMCs [Desharnais 1999b; Baier et al. 2005], logical simulation does not characterize simulation for MDPs. One can recover the correspondence between logical simulation and simulation, if each nondeterministic choice is labeled uniquely and the logic allows one to refer to the label of transitions [Desharnais et al. 2000].

Example 2.13. This example highlights the motivation for the syntactic restriction of PCTL safety that we consider here. Consider the two MDPs \mathcal{M}_1 and \mathcal{M}_2

Fig. 1. MDPs \mathcal{M}_1 and \mathcal{M}_2 from Example 2.13.

shown in Figure 1. The initial state of \mathcal{M}_1 is q_0 and the initial state of \mathcal{M}_2 is q_3 . The states q_2 and q_5 are labeled by proposition P . No other states are labeled by any proposition. Note that \mathcal{M}_1 is simulated by \mathcal{M}_2 . Consider the PCTL formula $\psi = \mathcal{P}_{\leq \frac{1}{2}}(\mathbf{X}(\mathcal{P}_{> \frac{1}{2}}(\mathbf{X}P)))$. Now $\mathcal{M}_2 \models \psi$, but $\mathcal{M}_1 \not\models \psi$. This is because \mathcal{M}_2 has a transition from the state q_4 in which the probability of transitioning to q_5 is exactly $\frac{1}{2}$. Therefore, ψ is not a safety formula (although if we were only considering DTMCs, ψ would be a safety formula; see remark above for the PCTL safety fragment for DTMCs presented in [Desharnais 1999b; Baier et al. 2005]). On the other hand, the formula $\psi_S = \mathcal{P}_{\leq \frac{1}{2}}(\mathbf{X}(\neg \mathcal{P}_{\leq \frac{1}{2}}(\mathbf{X}P)))$ is a safety formula and not satisfied by both \mathcal{M}_1 and \mathcal{M}_2 . Note that ψ_S and ψ are logically equivalent on DTMCs.

Example 2.14. We now give examples of safety and liveness properties in the context of PCTL. Consider a leader election protocol that elects a leader from among a collection of n (peer) processes $\text{proc}_1, \text{proc}_2, \dots, \text{proc}_n$. Such protocols are central to achieving a number of goals in a distributed system, including Firewire Root Contention. The model of the protocol has probabilistic transitions because of randomization employed by the leader election algorithm, and nondeterminism due to the asynchrony present in a concurrent, distributed system. Let us assume that the proposition `ElectedLeader` labels exactly those global states where all the participants agree upon a common leader and the proposition `ProtocolFinished` labels exactly those global states where the protocol has finished for all the n -processes. Now one of the properties that one may want to verify about such a protocol is that a leader is *almost surely* elected. This can be written as $\mathcal{P}_{\leq 0}(\text{tt } \mathbf{U} ((\neg \text{ElectedLeader}) \wedge \text{ProtocolFinished}))$, and as can be seen it is a PCTL-safety property as per our characterization. Consider now a stronger property, where we want to ensure that a leader is elected with high probability within a certain time bound, where the deadline is either a real-time deadline or based on the number of steps taken by the protocol. Let us assume that the proposition `DeadlinePassed` labels exactly those global states where the time is greater than the deadline. Then our stronger correctness requirement can be written as $\mathcal{P}_{< 0.05}(\text{tt } \mathbf{U} (\text{DeadlinePassed} \wedge \neg \text{ElectedLeader}))$. This is a safety property by our classification.

Now, assume we want to ensure *weak fairness*— it is possible for each process proc_i to be elected leader. Let us assume that the proposition `Leaderi` labels those




Fig. 2. DTMC with a large set of counterexample executions

states in which the process proc_i is elected leader. Then weak fairness can be written as the PCTL-liveness formula $\psi = \bigwedge_{1 \leq i \leq n} \neg \mathcal{P}_{\leq 0}(\diamond \text{Leader}_i)$.

3. COUNTEREXAMPLES

What is a counterexample? Clarke et al. [2002] say that counterexamples must (a) serve as an “explanation” of why the (abstract) model violates the property, (b) must be rich enough to explain the violation of a large class of properties, and (c) must be simple and specific enough to identify bugs, and be amenable to efficient generation and analysis.

In this section, we discuss three relevant proposals for counterexamples. The first one is due to [Han and Katoen 2007a], who present a notion of counterexamples for DTMCs. This has been recently extended to MDPs by Aljazzar and Leue [2007]. The second proposal for counterexamples was suggested in the context of non-probabilistic systems and branching time properties by Clarke et al. [2002]. Finally the third one has been recently suggested by Chatterjee et al. [2005; Hermanns et al. [2008] for MDPs. We examine all these proposals in order and identify why each one of them is inadequate for our purposes. We then present the definition of counterexamples that we consider in this paper.


3.1 Set of Traces as Counterexamples

The problem of defining a notion of counterexamples for probabilistic systems was first considered in [Han and Katoen 2007a]. Han and Katoen [2007a] present a notion of counterexamples for DTMCs and define a counterexample to be a *finite* set of executions such that the measure of the set is greater than some threshold (they consider weak safety formulas only). The problem to compute the smallest set of executions is intractable, and Han and Katoen present algorithms to generate such a set of executions. This definition has recently been extended for MDPs in [Aljazzar and Leue 2007].

The “set of executions as a counterexample” proposal has a few drawbacks. Consider the DTMC \mathcal{M} shown in Figure 2, where proposition P is true only in state q_3 and the state q_1 is the initial state. This DTMC violates the property $\psi_S = \mathcal{P}_{<1}(\diamond P)$. However, there is no finite set of executions that witnesses this violation. Such properties are not considered in [Han and Katoen 2007a; Aljazzar and Leue 2007].

Even if the property being considered is a weak safety property, the number of paths serving as the counterexample can be large. For example, let $\psi_S = \mathcal{P}_{\leq 1-\delta}(\diamond P)$. The Markov chain \mathcal{M} violates property ψ_S for all values of $\delta > 0$. However, one can show that the smallest set of counterexamples is large due to the following observations.

- Any execution, starting from q_1 , reaching q_3 is of the form $(q_1 q_2)^k q_3$ with measure $(1-\epsilon)^{k-1} \epsilon$. Thus the measure of the set $Exec = \{(q_1 q_2)^k q_3 \mid k \leq n\}$ is $1 - (1-\epsilon)^n$,

Fig. 3. DTMC \mathcal{M}_{notree} : No tree-like counterexamples

and the set $Exec$ has size $O(n^2)$.

—Thus the smallest set of examples that witnesses the violation of ψ_S has at least $r = \frac{\log \delta}{\log(1-\epsilon)}$ elements and the number of nodes in this set is $O(r^2)$.

r can be very large (for example, take $\epsilon = \frac{1}{2}$ and $\delta = \frac{1}{2^{2^r}}$). In such circumstances, it is unclear whether such a set of executions can serve as a comprehensible explanation for why the system violates the property ψ_S .

3.2 Tree-like Counterexamples

In the context of non-probabilistic systems and branching-time properties, Clarke et al. [2002] suggest that counterexamples should be “tree-like”. The reason to consider this proposal carefully is because probabilistic logics like PCTL are closely related to branching-time logics like CTL. *Tree-like counterexamples* for a Kripke structure \mathcal{K} and property φ are defined to be a Kripke structure \mathcal{E} such that (a) \mathcal{E} violates the property φ , (b) \mathcal{E} is simulated by \mathcal{K} , and (c) the underlying graph of \mathcal{E} is *tree-like*, i.e., (i) every non-trivial maximal strongly connected component is a cycle, and (ii) the graph of maximal strongly connected components forms a tree. Clarke et al. [2002] argue that this is the appropriate notion of counterexamples because tree-like counterexamples are easy to comprehend. Moreover, they show that for any Kripke structure \mathcal{K} that violates an ACTL* formula φ , there is a tree-like counterexample \mathcal{E} .

The notion of tree-like counterexamples can be naturally extended to the case of MDPs. Formally, a *tree-like counterexample* for a MDP \mathcal{M} and property ψ_S will be a (disjoint) MDP \mathcal{E} such that the unlabeled underlying graph $G(\mathcal{E})$ is tree-like, \mathcal{E} violates property ψ_S and is simulated by \mathcal{M} . However, surprisingly, unlike the case for Kripke structures and ACTL*, the family of tree-like counterexamples is not rich enough.

Example 3.1. Consider the DTMC \mathcal{M}_{notree} shown in Figure 3 with initial state q_1 , proposition P being true only in state q_3 , and propositions P_1, P_2 and P_4 being true only in states q_1, q_2 and q_4 respectively. Consider the formula $\psi_S = \mathcal{P}_{<1}((P_1 \vee P_2 \vee P_4) \mathcal{U} P)$. Clearly, the DTMC \mathcal{M}_{notree} violates ψ_S .

We will show that there is no tree-like counterexample for \mathcal{M}_{notree} and formula ψ_S , defined in Example 3.1. This is done in two steps. First we prove that there is no tree-like DTMC counterexample for \mathcal{M}_{notree} and ψ_S ; this is the content of Proposition 3.2. Next, by using Proposition 3.2, we show in Lemma 3.3, that there is no tree-like MDP counterexample as well, which gives us our desired result.

The absence of tree-like DTMC counterexamples (Proposition 3.2) is proved as follows.

(1) First, assume that there is a tree-like DTMC \mathcal{T} such that \mathcal{T} is simulated by \mathcal{M}_{notree} and $\mathcal{T} \not\models \psi_S$. Without loss of generality, the initial state of \mathcal{T} can be taken to be in the strongly connected component that is at the root of the component tree of \mathcal{T} .

(2) From the fact that \mathcal{M}_{notree} simulates \mathcal{T} and each reachable state of \mathcal{M} is labeled by a unique proposition, we deduce that each state of \mathcal{T} is simulated by exactly one state of \mathcal{M} . Thus the set of states of \mathcal{T} can be partitioned into 4 sets— $\mathbb{Q}_i = \{q \mid q \text{ is a state of } \mathcal{T} \text{ and } q \text{ is simulated by } q_i\}$ for $i = 1, 2, 3, 4$.

(3) Next we show that for each state $q \in \mathbb{Q}_i$, it is the case that the probability of transitioning from q to \mathbb{Q}_j is the probability of transitioning from q_i to q_j in \mathcal{M}_{notree} which allows to conclude that each $q \in \mathbb{Q}_i$ is in fact bisimilar to q_i and not just simulated by q_i .

(4) Next, we observe that since there is no transition out of q_3 , each state of \mathbb{Q}_3 must be a leaf node of underlying unlabeled graph of \mathcal{T} . Consider the DTMC \mathcal{T}_1 obtained by deleting the set of states \mathbb{Q}_3 from \mathcal{T} and the DTMC \mathcal{M}_1 obtained by deleting the state q_3 from \mathcal{M}_{notree} . \mathcal{T}_1 is also tree-like and bisimilar to \mathcal{M}_1 .

(5) Consider a strongly connected component \mathcal{C} of \mathcal{T}_1 which is a leaf node of the component graph of \mathcal{T}_1 . Now, this strongly connected component must be bisimilar to \mathcal{M}_1 . However, the strongly connected component $\{q_1, q_2, q_4\}$ of \mathcal{M}_1 has two cycles $q_1 \rightarrow q_4 \rightarrow q_2 \rightarrow q_1$ and $q_1 \rightarrow q_2 \rightarrow q_1$. This implies that \mathcal{C} must also contain two different cycles. Hence \mathcal{C} cannot be a cycle which violates the tree-like structure of \mathcal{T}_1 .

Using this observation, the absence of tree-like MDP counterexample (Lemma 3.3), is proved based on the following idea. If there is a MDP \mathcal{M} such that \mathcal{M} violates ψ_S and $\mathcal{M} \preceq \mathcal{M}_{notree}$, then there must be a memoryless scheduler \mathcal{S} such that the DTMC $\mathcal{M}^{\mathcal{S}}$ violates ψ_S and is also simulated \mathcal{M}_{notree} . The result now follows from the observation that $\mathcal{M}^{\mathcal{S}}$ is tree-like if \mathcal{M} is tree-like. We are now ready to present the formal details of these statements and their proofs.

PROPOSITION 3.2. *Consider the DTMC \mathcal{M}_{notree} and safety formula ψ_S defined in Example 3.1. If $\mathcal{T} = (\mathbb{Q}, q_{\mathcal{I}}, \delta, \mathbb{L})$ is a DTMC (disjoint from \mathcal{M}_{notree}) such that $\mathcal{T} \preceq \mathcal{M}_{notree}$ and $\mathcal{T} \not\models \psi_S$ then \mathcal{T} is not tree-like.*

PROOF. Assume, by way of contradiction, that \mathcal{T} is tree-like. Let $\mathbb{Q}_0 = \{q_1, q_2, q_3, q_4\}$ and for each $1 \leq i \leq 4$, let μ_{q_i} denote the transition out of q_i in \mathcal{M}_{notree} .

For each $q \in \mathbb{Q}$, let μ_q denote the transition out of q in \mathcal{T} . Let $\mathcal{R} \subseteq \mathbb{Q} \times \mathbb{Q}_0$ be a canonical simulation that witnesses the fact that $\mathcal{T} \preceq \mathcal{M}_{notree}$. We have by definition, $q_{\mathcal{I}} \mathcal{R} q_1$. Please note that since \mathcal{T} violates ψ_S the measure of all paths of \mathcal{T} starting with $q_{\mathcal{I}}$ and satisfying $(P_1 \vee P_2 \vee P_4) \mathcal{U} P$ is 1.

As \mathcal{T} is tree-like, any non-trivial strongly connected components of $G(\mathcal{T})$ is a cycle and \mathbb{G} , the graph of the strongly connected components (trivial or non-trivial) of $G(\mathcal{T})$ form a tree. Without loss of generality, we can assume that the strongly connected component that forms the root of \mathbb{G} contains $q_{\mathcal{I}}$ (otherwise we can just consider the DTMC restricted to the states reachable from $q_{\mathcal{I}}$).

Hence, we have that every state in \mathbb{Q} is reachable from $q_{\mathcal{I}}$ with non-zero probability. From this and the fact \mathcal{R} is a canonical simulation, we can show that for any state $q \in \mathbb{Q}$ there is a $q' \in \mathbb{Q}_0$ such that $q \mathcal{R} q'$. Also since each state in \mathbb{Q}_0 is

labeled by a unique proposition, it follows that for each $q \in \mathbb{Q}$ there is a unique $q' \in \mathbb{Q}_0$ such that $q \mathcal{R} q'$ (in other words, \mathcal{R} is total and functional).

Now, for each $1 \leq i \leq 4$, let $\mathbb{Q}_i \subseteq \mathbb{Q}$ be the set $\{q \in \mathbb{Q} \mid q \mathcal{R} q_i.\}$ By the above observations, we have that \mathbb{Q}_i 's are pairwise disjoint; $\mathbb{Q} = \mathbb{Q}_1 \cup \mathbb{Q}_2 \cup \mathbb{Q}_3 \cup \mathbb{Q}_4$; and for each $1 \leq i \leq 4$, $\mathbb{Q}_i \cup \{q_i\}$ is a \mathcal{R} -closed set. Since \mathcal{R} is a canonical simulation, whenever $q \mathcal{R} q'$, we have $\mu_q(\mathbb{Q}_i) \leq \mu_{q'}(q_i)$, for each $1 \leq i \leq 4$. Moreover, we can, in fact, prove the following stronger claim.

Claim: $\mu_q(\mathbb{Q}_i) = \mu_{q'}(q_i)$ for each $q \mathcal{R} q'$ and $1 \leq i \leq 4$.

Proof of the claim: Consider some q, q' such that $q \mathcal{R} q'$. We proceed by contradiction. Assume that there is some i such $\mu_q(\mathbb{Q}_i) < \mu_{q'}(q_i)$. Please note that in this case $q' \neq q_3$ (as $\mu_{q_3}(q_i) = 0, \forall 1 \leq i \leq 4$).

There are several possible cases (depending q' and i). We just discuss the case when q' is q_4 and i is 2. The other cases are similar. For this case we have that $\mu_q(\mathbb{Q}_2) < 1$. Also note that for $j \neq 2$, $\mu_q(\mathbb{Q}_j) \leq \mu_{q_4}(q_j) = 0$. Hence, $\mu_q(\mathbb{Q}) < 1$. Now, pick two new states q_{new_2} and q_{new_3} not occurring in $\mathbb{Q} \cup \mathbb{Q}_0$. Construct a new tree-like DTMC \mathcal{T}' extending \mathcal{T} as follows. The states of \mathcal{T}' are $\mathbb{Q} \cup \{q_{new_2}, q_{new_3}\}$. Only proposition P_2 is true in q_{new_2} and only proposition P is true in q_{new_3} . The labeling function for other states remains the same. We extend the probabilistic transition μ_q by letting $\mu_q(q_{new_2}) = 1 - \mu_q(\mathbb{Q})$ and $\mu_q(q_{new_3}) = 0$ (transition probabilities to other states do not get affected). The state q_{new_2} has a probabilistic transition $\mu_{q_{new_2}}$ such that $\mu_{q_{new_2}}(q_{new_3}) = \frac{1}{4}$ and $\mu_{q_{new_2}}(\bar{q}) = 0$ for any $\bar{q} \neq q_{new_3}$. The transition probability from q_{new_3} to any state is 0. For all other states the transitions remain the same.

Now, please note that there is a path π from $q_{\mathcal{T}}$ to q (in \mathcal{T} and hence in \mathcal{T}' also) with non-zero “measure” such that $P_1 \vee P_2 \vee P_4$ is true at each point in this path. Furthermore, at each point in this path, P is false. Consider the path π' in \mathcal{T}' obtained by extending π by q_{new_2} followed by q_{new_3} . Now, by construction π' satisfies $(P_1 \vee P_2 \vee P_4) \mathcal{U} P$ and the “measure” of this path > 0 (as $1 - \mu_q(\mathbb{Q}) > 0$). Now the path π' is not present in \mathcal{T} and hence the measure of all paths in \mathcal{T}' that satisfy $(P_1 \vee P_2 \vee P_4) \mathcal{U} P$ is strictly greater than the measure of all paths in \mathcal{T} that satisfy $(P_1 \vee P_2 \vee P_4) \mathcal{U} P$. The latter number is 1 and thus the measure of all paths in \mathcal{T}' that satisfy $(P_1 \vee P_2 \vee P_4) \mathcal{U} P > 1$. Impossible! \square (End proof of the claim)

We proceed with the proof of the main proposition. Let $\mathcal{R}_1 \subseteq (\mathbb{Q} \cup \mathbb{Q}_0) \times (\mathbb{Q} \cup \mathbb{Q}_0)$ be the reflexive, symmetric and transitive closure of \mathcal{R} (in other words, the smallest equivalence that contains \mathcal{R}). It is easy to see that the equivalence classes of \mathcal{R}_1 are exactly $\mathbb{Q}_i \cup \{q_i\}, 1 \leq i \leq 4$. From this fact and our claim above, we can show that \mathcal{R}_1 is a bisimulation.

Observe now that each element of $\mathbb{Q}_3 \subseteq \mathbb{Q}$ must be a leaf node of \mathbb{G} , the graph of the strongly connected components of $G(\mathcal{T})$. Using this, one can easily show that if \mathcal{T}_1 is the DTMC obtained from \mathcal{T} by restricting the state space to $\mathbb{Q}_1 \cup \mathbb{Q}_2 \cup \mathbb{Q}_4$, then \mathcal{T}_1 is tree-like. Let \mathcal{M}_1 be the DTMC obtained from \mathcal{M}_{notree} by restricting the state space to $\mathbb{Q}_0 \setminus \{q_3\}$ and let $\tilde{\mathbb{Q}} = (\mathbb{Q}_1 \cup \mathbb{Q}_2 \cup \mathbb{Q}_4) \cup (\mathbb{Q}_0 \setminus \{q_3\})$. It is easy to see that the equivalence relation $\mathcal{R}_2 = \mathcal{R}_1 \cap (\tilde{\mathbb{Q}} \times \tilde{\mathbb{Q}})$ is also a bisimulation.

Now, let \mathbb{G}_1 be the graph of strongly connected components of $G(\mathcal{T}_1)$. Now, fix a strongly connected component of $G(\mathcal{T}_1)$, say \mathbb{C} , that is a leaf node of \mathbb{G}_1 . Fix a state q which is a node of \mathbb{C} . Since \mathbb{G}_1 is tree-like and \mathbb{C} is a leaf node, it is

easy to see that $\text{post}(\mu_q, q)$ can contain at most 1 element. Also, we have that $q \in Q_1 \cup Q_2 \cup Q_4$. Now if $q \in Q_1$, we have that q (as a state of \mathcal{T}_1) is bisimilar to q_1 (as a state of \mathcal{M}_1). However, this implies that $\text{post}(\mu_q, q)$ must be at least 2 as q_1 has a non-zero probability of transitioning to 2 states labeled by different propositions. Hence $q \notin Q_1$. If $q \in Q_2$, then please note that $\text{post}(\mu_q, q)$ must contain an element in Q_1 which should also be in C . By the above observation this is not possible. Hence $q \notin Q_2$. Similarly, we can show that $q \notin Q_4$. Hence $q \notin Q_1 \cup Q_2 \cup Q_4$. A contradiction. \square

We are ready to show that $\mathcal{M}_{\text{notree}}$ has no tree-like counterexamples.

LEMMA 3.3. *Consider the DTMC $\mathcal{M}_{\text{notree}}$ and formula ψ_S defined in Example 3.1. There is no tree-like counterexample witnessing the fact that $\mathcal{M}_{\text{notree}}$ violates ψ_S .*

PROOF. First, since $((P_1 \vee P_2 \vee P_3) \mathcal{U} P)$ is a simple reachability formula, if there is a MDP $\mathcal{E} \preceq \mathcal{M}_{\text{notree}}$ which violates ψ_S , then there is a memoryless scheduler S such that \mathcal{E}^S violates the property [Bianco and de Alfaro 1995]. Now note that if we were to just consider the states of \mathcal{E}^S reachable from the initial state then \mathcal{E}^S is also tree-like. In other words, there is a tree-like DTMC that is simulated by $\mathcal{M}_{\text{notree}}$ and which violates the property ψ_S . The result now follows from Proposition 3.2. \square

Tree-like graph structures are not rich enough for PCTL-safety. However, it can be shown that if we restrict our attention to weak safety formulas, then we have tree counterexamples. However, such trees can be very big as they depend on the actual transition probabilities.

THEOREM 3.4. *If ψ_{WS} is a weak safety formula and $\mathcal{M} \not\models \psi_{WS}$, then there is a \mathcal{M}' such that $G(\mathcal{M}')$ is a tree, $\mathcal{M}' \preceq \mathcal{M}$, and $\mathcal{M}' \not\models \psi_{WS}$.*

PROOF. The result follows from the following two observations.

- If the underlying graph $G(\mathcal{M}_1)$ of a MDP \mathcal{M}_1 is acyclic then there is a MDP \mathcal{M}_2 such that $G(\mathcal{M}_2)$ is a tree and $\mathcal{M}_1 \approx \mathcal{M}_2$.
- For any strict liveness formula ψ_{SL} , and a state $q \in \mathcal{M}$ if $q \Vdash_{\mathcal{M}} \psi_{SL}$ then there is a k such that $\mathcal{M}_k^q \Vdash \psi_{SL}$ where \mathcal{M}_k^q is the k -th unrolling of \mathcal{M} rooted at q .

The latter observation can be proved by a straightforward induction on the structure of strict liveness formulas. We consider the case when ψ_{SL} is $(\neg \mathcal{P}_{\leq p}(\psi_{SL_1} \mathcal{U} \psi_{SL_2}))$.

Now if $q \Vdash \neg \mathcal{P}_{\leq p}(\psi_{SL_1} \mathcal{U} \psi_{SL_2})$ then note that there is a memoryless scheduler \mathcal{S} [Bianco and de Alfaro 1995; Han and Katoen 2007a] such that there is a finite set of finite paths of the DTMC $\mathcal{M}^{\mathcal{S}}$ starting from q such that 1) each path satisfies $\psi_{SL_1} \mathcal{U} \psi_{SL_2}$ and 2) the measure of these paths $> p$. Now, these paths can be arranged in a tree T nodes of which are labeled by the corresponding state of \mathcal{M} . If a state q' labels a leaf node of T then we have $q' \Vdash_{\mathcal{M}} \psi_{SL_2}$; otherwise $q' \Vdash_{\mathcal{M}} \psi_{SL_1}$. Note for a non-leaf node, we can assume that $q' \not\Vdash_{\mathcal{M}} \psi_{SL_2}$.

For any state $q' \in Q$ labeling a node in T define $\text{maxdepth}(q') = \max\{\text{depth}(t) \mid t \text{ is a node of } T \text{ and } t \text{ is labeled by } q'\}$. If q' labels a non-leaf node of T fix $k_{q'}$ such that $\mathcal{M}_{k_{q'}}^{q'} \Vdash \psi_{SL_1}$ ($k_{q'}$ exists by induction hypothesis). If q' labels a leaf node of




Fig. 4. MDP \mathcal{M} used in the proof of Proposition 3.5. Transition μ_1 is shown as dashed edges, and transition μ_2 is shown as solid edges.

Then fix $k_{q'}$ such that $\mathcal{M}_{k_{q'}}^{q'} \Vdash \psi_{SL_2}$. Now, let $k = \max\{\text{maxdepth}(q') + k_{q'} \mid q' \text{ labels a node of } \mathbb{T}\}$. It can now be shown easily that $\mathcal{M}_k^q \Vdash \psi_{SL}$. \square

3.3 DTMCs as counterexamples

We now consider the third and final proposal for a notion of counterexamples that is relevant for MDPs. Chatterjee et al. [2005] use the idea of abstraction-refinement to synthesize winning strategies for stochastic 2-player games. They abstract the game graph, construct winning strategies for the abstracted game, and check the validity of those strategies for the original game. They observe that for discounted reward objectives and average reward objectives, the winning strategies are *memoryless*, and so “counterexamples” can be thought of as finite-state models without nondeterminism (which is resolved by the strategies constructed).

This idea is also used in [Hermanns et al. 2008]. They observe that for weak safety formulas of the form $\mathcal{P}_{\leq p}(\psi_1 \mathcal{U} \psi_2)$ where ψ_1 and ψ_2 are propositions (or boolean combinations of propositions), if a MDP \mathcal{M} violates the property then there is a memoryless scheduler \mathcal{S} such that the DTMC $\mathcal{M}^{\mathcal{S}}$ violates $\mathcal{P}_{\leq p}(\psi_1 \mathcal{U} \psi_2)$ (see [Bianco and de Alfaro 1995]). Therefore, they take the pair $(\mathcal{S}, \mathcal{M}^{\mathcal{S}})$ to be the counterexample.

Motivated by these proposals and our evidence of the inadequacies of sets of executions and tree-like systems as counterexamples, we ask whether DTMCs (or rather purely probabilistic models) could serve as an appropriate notion for counterexamples of MDPs. We answer this question in the negative. The reason why DTMCs fail to suffice is the following: since each probabilistic operator can refer to different resolutions of the nondeterminism in the MDP, if a formula contains more than one probabilistic operator then in a counterexample more than one resolution of the nondeterminism may be required.

PROPOSITION 3.5. *There is a MDP \mathcal{M} and a safety formula ψ_S such that $\mathcal{M} \not\Vdash \psi_S$ but there is no DTMC \mathcal{M}' that violates ψ_S and is simulated by \mathcal{M} .*

PROOF. The MDP \mathcal{M} will have three states q_0, q_1, q_2 with q_0 as the initial state, and is shown in Figure 4. The transition probability from q_1 and q_2 to any other state is 0. There will be two transitions out of q_0 , μ_1 and μ_2 , where $\mu_1(q_0) = 0, \mu_1(q_1) = \frac{3}{4}, \mu_1(q_2) = \frac{1}{4}$ and $\mu_2(q_0) = 0, \mu_2(q_1) = \frac{1}{4}, \mu_2(q_2) = \frac{3}{4}$. For the labeling function, we pick two distinct propositions P_1 and P_2 and let $L(q_0) = \emptyset, L(q_1) = \{P_1\}$ and $L(q_2) = \{P_2\}$. Consider the safety formula $\psi_S = \mathcal{P}_{< \frac{3}{4}}(X(P_1 \wedge \neg P_2)) \vee \mathcal{P}_{< \frac{3}{4}}(X(\neg P_1 \wedge P_2))$. Now \mathcal{M} violates ψ_S .

Suppose that $\mathcal{M}' = (Q', q_I, \delta', L')$ is a counterexample for \mathcal{M} and ψ_S . Then we

must have $q_{\mathcal{I}} \Vdash \neg \mathcal{P}_{<\frac{3}{4}}(\mathbf{X}(P_1 \wedge \neg P_2)) \wedge \neg \mathcal{P}_{<\frac{3}{4}}(\mathbf{X}(\neg P_1 \wedge P_2))$. Now, if \mathcal{M}' is a DTMC, $\delta'(q_{\mathcal{I}})$ must contain exactly one element $\mu_{q_{\mathcal{I}}}$. Also since $q_{\mathcal{I}} \Vdash \neg \mathcal{P}_{<\frac{3}{4}}(\mathbf{X}(P_1 \wedge \neg P_2))$ there must be a set of states \mathbf{Q}'_1 such that $\mu_{q_{\mathcal{I}}}(\mathbf{Q}'_1) \geq \frac{3}{4}$ and for each $q'_1 \in \mathbf{Q}'_1$, $P_1 \in \mathbf{L}'(q'_1)$ but $P_2 \notin \mathbf{L}'(q'_1)$. Similarly, there must also be a set of states \mathbf{Q}'_2 such that $\mu_{q_{\mathcal{I}}}(\mathbf{Q}'_2) \geq \frac{3}{4}$ and for each $q'_2 \in \mathbf{Q}'_2$, $P_2 \in \mathbf{L}'(q'_2)$ but $P_1 \notin \mathbf{L}'(q'_2)$. Now, clearly $\mathbf{Q}'_1 \cap \mathbf{Q}'_2 = \emptyset$. Thus, we have that $\mu_{q_{\mathcal{I}}}(\mathbf{Q}') \geq \frac{3}{4} + \frac{3}{4} > 1$. A contradiction. \square

Remark 3.6. Please note that in the proof of Proposition 3.5, we have used a safety formula which is a disjunction of safety formulas. We could also construct examples in which the safety formula consists of nested reachability formulas, *i.e.*, a formula of the form $\mathcal{P}_{\leq p}(\phi \mathcal{U} \phi')$ where ϕ and/or ϕ' are not boolean combinations of propositional symbols. For example, consider the MDP \mathcal{M} in the proof of Proposition 3.5 and let ψ be safety formula $\mathcal{P}_{<\frac{3}{4}}((\neg \mathcal{P}_{<\frac{3}{4}}(\mathbf{X}P_2)) \mathcal{U} P_1)$. Then $\mathcal{M} \not\models \psi$, and it can be shown along the lines of the proof of Proposition 3.5 that there is no DTMC counterexample for \mathcal{M} and ψ .

3.4 Our Proposal: MDPs as Counterexamples

Counterexamples for MDPs with respect to safe PCTL formulas cannot have any special structure. We showed that there are examples of MDPs and properties that do not admit any tree-like counterexample (Section 3.2). We also showed that there are examples that do not admit collections of executions, or general DTMCs (*i.e.*, models without nondeterminism) as counterexamples (Sections 3.1, 3.3). Therefore in our definition, counterexamples will simply be general MDPs. We will further require that counterexamples carry a “proof” that they are counterexamples in terms of a canonical simulation which witnesses the fact the given MDP simulates the counterexample. Although we do not really need to have this simulation in the definition for discussing counterexamples (one can always compute a simulation), this slight extension will prove handy while discussing counterexample guided refinement. Formally,

Definition 3.7. For a MDP $\mathcal{M} = (\mathbf{Q}, q_{\mathcal{I}}, \delta, \mathbf{L})$ and safety property ψ_S such that $\mathcal{M} \not\models \psi_S$, a counterexample is pair $(\mathcal{E}, \mathcal{R})$ such that $\mathcal{E} = \{\mathbf{Q}_{\mathcal{E}}, q_{\mathcal{E}}, \delta_{\mathcal{E}}, \mathbf{L}_{\mathcal{E}}\}$ is a MDP disjoint from \mathcal{M} , $\mathcal{E} \not\models \psi_S$ and $\mathcal{R} \subseteq \mathbf{Q}_{\mathcal{E}} \times \mathbf{Q}$ is a canonical simulation.

For the counterexample to be useful we will require that it be “small”. Our definition of what it means for a counterexample to be “small” will be driven by another requirement outlined by in [Clarke et al. 2002], namely, that it should efficiently generatable. These issues will be considered next.

3.5 Computing Counterexamples

Since we want the counterexample to be small, one possibility would be to consider the smallest counterexample. The size of a counterexample $(\mathcal{E}, \mathcal{R})$ can be taken to be the sum of sizes of the underlying labeled graph of \mathcal{E} , the size of the numbers used as probabilities in \mathcal{E} and the cardinality of the set \mathcal{R} ; the smallest counterexample is then the one that has the smallest size. However, it turns out that computing the smallest counterexample is a computationally hard problem. This is the formal content of our next result. For this section, we assume the standard definition of the size of a PCTL formula.

Notation 3.8. Given a safety formula, ψ_S , we denote the size of ψ_S by $|\psi_S|$.

We now formally define the size of the counterexample.

Definition 3.9. Let $\mathcal{M} = (\mathbf{Q}, q_I, \delta, L)$ be a MDP. The size of \mathcal{M} , denoted as $|\mathcal{M}|$, is the sum of the size (vertices+edges) of the labeled underlying graph $G_\ell(\mathcal{M})$ and the sum $\sum_{q \in \mathbf{Q}, q' \in \mathbf{Q}, \mu \in \delta(q), \mu(q') > 0} \text{size}(\mu(q'))$. The size of a counterexample $(\mathcal{E}, \mathcal{R})$, denoted as $|(\mathcal{E}, \mathcal{R})|$, is the sum of the size of \mathcal{E} and the cardinality (number of elements) of the relation \mathcal{R} .

Please note that any MDP \mathcal{M} of size n has a counterexample of size $\leq 2n$ (just take an isomorphic copy of \mathcal{M} as the counterexample MDP and take the obvious “injection” as the canonical simulation relation). The next theorem shows that the problem of finding the smallest counterexample is NP-complete. The proof of the theorem has been moved to an Electronic Appendix so as to not disturb the flow of the paper.

THEOREM 3.10. *Given a MDP \mathcal{M} , a safety formula ψ_S such that $\mathcal{M} \not\models \psi_S$, and a number $k \leq 2|\mathcal{M}|$, deciding whether there is a counterexample $(\mathcal{E}, \mathcal{R})$ of size $\leq k$ is NP-complete.*

Not only is the problem of finding the smallest counterexample NP-hard, it also unlikely to be efficiently approximable. This is the content of the next theorem whose proof has also been moved to an Electronic Appendix.

THEOREM 3.11. *Given a MDP \mathcal{M} , a safety formula ψ_S and $n = |\mathcal{M}| + |\psi_S|$ such that $\mathcal{M} \not\models \psi_S$. The smallest counterexample for \mathcal{M} and ψ_S cannot be approximated in polynomial time to within $O(2^{\log^{1-\epsilon} n})$ unless $\text{NP} \subseteq \text{DTIME}(n^{\text{poly} \log(n)})$.*

Remark 3.12. A few points about our hardness and inapproximability results are in order.

- (1) Please note that we did not take the size of the labeling function into account. One can easily modify the proofs to take this into account.
- (2) The same reduction also shows a lower bound for the safety fragment of ACTL* properties as the reduction does not rely on any important features of quantitative properties.

Since finding the smallest counterexamples is computationally hard, we consider the problem of finding *minimal counterexamples*. Intuitively, a minimal counterexample has the property that removing any edge from the labeled underlying graph of the counterexample, results in a MDP that is not a counterexample. In order to be able to define this formally, we need the notion of when one MDP is contained in the other.

Definition 3.13. We say that a MDP $\mathcal{M}' = (\mathbf{Q}', q_I, \delta', L')$ is contained in a MDP $\mathcal{M} = (\mathbf{Q}, q_I, \delta, L)$ if $\mathbf{Q}' \subseteq \mathbf{Q}$, $L'(q') = L(q')$ for all $q' \in \mathbf{Q}'$, and for each $q' \in \mathbf{Q}'$ there is a 1-to-1 function $f_{q'} : \delta'(q') \rightarrow \delta(q')$ with the following property: for each $\mu' \in \delta'(q')$ and $q'' \in \mathbf{Q}'$, either $\mu'(q'') = f_{q'}(\mu')(q'')$ or $\mu'(q'') = 0$. We denote this by $\mathcal{M}' \subseteq \mathcal{M}$.

```

Initially  $M_{curr} = \bar{\mathcal{M}}$ 
For each edge  $(\bar{q}, \bar{q}_1) \in E_i$  in  $G_\ell(M_{curr})$ 
  Let  $M'$  be the MDP obtained from  $M_{curr}$  by setting  $\mu(\bar{q}_1) = 0$ , where  $\mu$  is the
   $i$ th choice out of  $\bar{q}$ , in  $M_{curr}$ 
  If  $M' \not\models \psi_S$  then  $M_{curr} = M'$ 
od  $\leftarrow$  End of For loop
Let  $\mathcal{E}$  be the MDP obtained from  $M_{curr}$  by removing the set of states from  $M_{curr}$ 
which are not reachable in the underlying unlabeled graph of  $M_{curr}$ 
If  $\mathcal{Q}_\mathcal{E}$  is the set of states of  $\mathcal{E}$ , then let  $\text{rel}_{inj} = \{(\bar{q}, q) \mid \bar{q} \in \mathcal{Q}_\mathcal{E}\}$ 
return( $\mathcal{E}, \text{rel}_{inj}$ )

```

Fig. 5. Algorithm for computing the minimal counterexample

Observe that if $\mathcal{M}' \subseteq \mathcal{M}$ then $\mathcal{M}' \preceq \mathcal{M}$. We present the definition of minimal counterexamples obtained by lexicographic ordering on pairs $(\mathcal{E}, \mathcal{R})$.

Definition 3.14. For a MDP \mathcal{M} and a safety property ψ_S , $(\mathcal{E}, \mathcal{R})$ is a minimal counterexample iff

- $(\mathcal{E}, \mathcal{R})$ is a counterexample for \mathcal{M} and ψ_S and
- If $(\mathcal{E}_1, \mathcal{R}_1)$ is also a counterexample for \mathcal{M} and ψ_S , then
 - $\mathcal{E}_1 \subseteq \mathcal{E}$ implies that $\mathcal{E}_1 = \mathcal{E}$; and
 - if $\mathcal{E}_1 = \mathcal{E}$ then $\mathcal{R}_1 \subseteq \mathcal{R}$ implies that $\mathcal{R}_1 = \mathcal{R}$.



Though finding the smallest counterexample is NP-complete and is unlikely to be efficiently approximable, there is a very simple polynomial time algorithm to compute a minimal counterexample. In fact the counterexample computed by our algorithm is going to be contained in the original MDP (upto “renaming” of states). Before we proceed, we fix some notation for the rest of the paper.

Notation 3.15. Given a MDP $\mathcal{M} = (\mathcal{Q}, q_I, \delta, \mathbf{L})$, for each $q \in \mathcal{Q}$ fix a unique element \bar{q} not occurring in \mathcal{Q} . Define an *isomorphic* MDP $\bar{\mathcal{M}} = (\bar{\mathcal{Q}}, \bar{q}_I, \bar{\delta}, \bar{\mathbf{L}})$ as follows.

- $\bar{\mathcal{Q}} = \{\bar{q} \mid q \in \mathcal{Q}\}$.
- $\bar{\delta}(\bar{q}) = \{\bar{\mu} \mid \mu \in \delta(q)\}$ where
 - for each $\bar{q} \in \bar{\mathcal{Q}}$, $\bar{\mu}(\bar{q}) = \mu(q)$.
- $\bar{\mathbf{L}}(\bar{q}) = \mathbf{L}(q)$.

We are ready to give the counterexample generation algorithm. The algorithm shown in Figure 5 clearly computes a minimal counterexample contained in the original MDP upto “renaming” of states (note that the minimality of rel_{inj} is a direct consequence of the fact that every state in \mathcal{E} is reachable from initial state and hence must be simulated by some state in \mathcal{M}). Its running time is polynomial because model checking problem for MDPs is in P [Bianco and de Alfaro 1995].

THEOREM 3.16. *Given a MDP \mathcal{M} and a safety formula ψ_S such that $\mathcal{M} \not\models \psi_S$, the algorithm in Figure 5 computes a minimal counterexample and runs in time polynomial in the size of \mathcal{M} and ψ_S .*

Fig. 6. Kripke structure \mathcal{K}_{ex} Fig. 7. Its abstraction \mathcal{K}_{ab}

Please note that for safety properties of the form $\mathcal{P}_{\leq p}(\psi_S \mathcal{U} \psi'_S)$ and $\mathcal{P}_{< p}(\psi_S \mathcal{U} \psi'_S)$ where ψ_S, ψ'_S are boolean combinations of propositions, if $(\mathcal{E}, \text{rel}_{\text{inj}})$ is the counterexample generated by Figure 5 then \mathcal{E} must be a DTMC. This is because if \mathcal{E} violates such a property then there is a memoryless scheduler \mathcal{S} such that $\mathcal{E}^{\mathcal{S}}$ violates the same property (see [Bianco and de Alfaró 1995]). For such properties; model checking algorithm also computes the memoryless scheduler witnessing the violation. Thus for such properties, one could initialize M_{curr} to be $\bar{\mathcal{M}}^{\mathcal{S}_1}$ where \mathcal{S}_1 is the memoryless scheduler generated when $\bar{\mathcal{M}}$ is model checked for violation of the given safety property.

There are a couple of important research questions to be investigated in connection with making the algorithm in Figure 5 scalable. First, the algorithm makes multiple calls to the model checker which can be expensive. It is therefore imperative that techniques to reuse the results of previous model checking runs be leveraged. Second, the counterexample returned by the algorithm, clearly depends on the order in which edges of $G_\ell(\bar{\mathcal{M}})$ are considered. Discovering heuristics for this ordering, based on the property and $\bar{\mathcal{M}}$, is an important research question.

4. ABSTRACTIONS

Usually, in counterexample guided abstraction-refinement framework, the abstract model is defined with the help of an equivalence relation on the states of the system [Clarke et al. 2000]. Informally, the construction for non-probabilistic systems proceeds as follows. Given a Kripke structure $\mathcal{K} = (\mathbf{Q}, q_{\mathcal{I}}, \rightarrow, \mathbf{L})$ and equivalence relation \equiv on \mathbf{Q} such that $\mathbf{L}(q) = \mathbf{L}(q')$ for $q \equiv q'$; the abstract Kripke structure for \mathcal{K} and \equiv is defined as the Kripke structure $\mathcal{K}_{\mathcal{A}} = (\mathbf{Q}_{\mathcal{A}}, q_{\mathcal{A}}, \rightarrow_{\mathcal{A}}, \mathbf{L}_{\mathcal{A}})$ where

- $\mathbf{Q}_{\mathcal{A}} = \{[q]_{\equiv} \mid q \in \mathbf{Q}\}$ is the set of equivalence classes under \equiv ,
- $q_{\mathcal{A}} = [q_{\mathcal{I}}]_{\equiv}$,
- $[q]_{\equiv} \rightarrow_{\mathcal{A}} [q']_{\equiv}$ if there is some $q_1 \in [q]_{\equiv}$ and $q'_1 \in [q']_{\equiv}$ such that $q_1 \rightarrow q'_1$, and
- $\mathbf{L}_{\mathcal{A}}([q]_{\equiv}) = \mathbf{L}(q)$.

Example 4.1. Consider the Kripke \mathcal{K}_{ex} structure given in Figure 6 where q_0 is the initial state and the state q_{11} is labeled by proposition P (no other state is labeled

by any proposition). Consider the equivalence relation \equiv which partitions the set $\{q_0, q_1, q_2, q_3, q_4, q_5, q_6, q_7, q_8, q_9, q_{10}, q_{11}\}$ into the equivalence classes $\{q_0, q_1, q_3\}$, $\{q_2\}$, $\{q_4\}$, $\{q_5, q_6\}$, $\{q_7, q_8\}$, $\{q_9\}$, $\{q_{10}\}$ and $\{q_{11}\}$. Then the abstract Kripke structure, \mathcal{K}_{ab} for \mathcal{K}_{ex} and \equiv is given by the Kripke structure in Figure 7. Here $\{q_0, q_1, q_3\}$ is the initial state and $\{q_{11}\}$ is labeled by proposition P .

This construction is generalized for MDPs in [Jonsson and Larsen 1991; Huth 2005; D'Argenio et al. 2001]. To describe this generalized construction formally, we first need to lift distributions on a set with an equivalence relation \equiv to a distribution on the equivalence classes of \equiv .³

Definition 4.2. Given $\mu \in \text{Prob}_{\leq 1}(\mathbf{Q})$ and an equivalence \equiv on \mathbf{Q} , the lifting of μ (denoted by $[\mu]_{\equiv}$) to the set of equivalence classes of \mathbf{Q} under \equiv is defined as $[\mu]_{\equiv}([q]_{\equiv}) = \mu(\{q' \in \mathbf{Q} \mid q' \equiv q\})$.

For a MDP $\mathcal{M} = (\mathbf{Q}, q_{\mathcal{I}}, \delta, \mathbf{L})$, we will say a binary relation \equiv is an *equivalence relation compatible with \mathcal{M}* if \equiv is an equivalence relation on \mathbf{Q} such that $\mathbf{L}(q) = \mathbf{L}(q')$ for all $q \equiv q'$. The abstract models used in our framework are then formally defined as follows.

Definition 4.3. Given a set of propositions AP, let $\mathcal{M} = (\mathbf{Q}, q_{\mathcal{I}}, \delta, \mathbf{L})$ be an AP labeled MDP. Let \equiv be an equivalence relation compatible with \mathcal{M} . The *abstract MDP for \mathcal{M} with respect to the equivalence relation \equiv* is a MDP $\mathcal{M}_{\equiv} = (\mathbf{Q}_{\equiv}, q_{\equiv}, \delta_{\equiv}, \mathbf{L}_{\equiv})$ where



- (1) $\mathbf{Q}_{\equiv} = \{[q]_{\equiv} \mid q \in \mathbf{Q}\}$.
- (2) $q_{\equiv} = [q_{\mathcal{I}}]_{\equiv}$.
- (3) $\delta_{\equiv}([q]_{\equiv}) = \{\mu \mid \exists q' \in [q]_{\equiv} \text{ and } \mu_1 \in \delta(q') \text{ such that } \mu = [\mu_1]_{\equiv}\}$.
- (4) $\mathbf{L}_{\equiv}([q]_{\equiv}) = \mathbf{L}(q)$.

The elements of \mathbf{Q} will henceforth be called *concrete states* and the elements \mathbf{Q}_{\equiv} will henceforth be called *abstract states*. The relation $\text{rel}_{\equiv}^{\alpha} \subseteq \mathbf{Q} \times \mathbf{Q}_{\equiv}$ defined as $\text{rel}_{\equiv}^{\alpha} = \{(q, [q]_{\equiv}) \mid q \in \mathbf{Q}\}$ will henceforth be called the *abstraction relation*. The relation $\text{rel}_{\equiv}^{\gamma} \subseteq \mathbf{Q}_{\equiv} \times \mathbf{Q}$ defined as $\text{rel}_{\equiv}^{\gamma} = \{([q]_{\equiv}, q') \mid [q]_{\equiv} \in \mathbf{Q}_{\equiv}, q \equiv q'\}$ will henceforth be called the *concretization relation*.

Remark 4.4. The relation $\text{rel}_{\equiv}^{\alpha}$ is total and functional and hence represents a function α which is often called the *abstraction map* in literature. Please note that one can define the equivalence \equiv via the function α . The relation $\text{rel}_{\equiv}^{\gamma}$ is total (not necessarily functional) and hence represents a map into the power-set $2^{\mathbf{Q}}$. The function $\gamma : \mathbf{Q}_{\equiv} \rightarrow 2^{\mathbf{Q}}$ defined as $\gamma(a) = \text{rel}_{\equiv}^{\gamma}(a)$ is often called the *concretization map* in literature.

We conclude this section by making a couple of observations about the construction of the abstract MDP. First notice that the abstract MDP \mathcal{M}_{\equiv} has been defined

³It is possible to avoid lifting distributions if one assumes that each transition in the system is uniquely labeled, and has the property that the target sub-probability measure has non-zero measure for at most one state. This does not affect the expressive power of the model and is used in [Hermanns et al. 2008]. However the disadvantage is that the abstract model may be larger as fewer transitions will be collapsed.

Fig. 8. DTMC \mathcal{M}_{ex} Fig. 9. MDP \mathcal{M}_{ab} . The state $\{q_1, q_2\}$ (and $\{q_3, q_4\}$) has a nondeterministic choice between two transitions corresponding to q_1 's (q_4 's) transition and q_2 's (q_3 's) transition in \mathcal{M}_{ex} . The former is shown with dashed edge and the latter as a solid edge.

to ensure that it simulates \mathcal{M} via the canonical simulation relation $\text{rel}_{\equiv}^{\alpha}$. Next, we show that we can obtain a “refinement” of the abstract MDP \mathcal{M}_{\equiv} , by considering the abstraction of \mathcal{M} with respect to another equivalence \simeq that is finer than \equiv . This is stated next.

Definition 4.5. Let $\mathcal{M} = (\mathbf{Q}, q_{\mathcal{I}}, \delta, \mathbf{L})$ be a MDP over the set of atomic propositions AP. Further let \equiv and \simeq be two equivalence relations compatible with \mathcal{M} such that $\simeq \subseteq \equiv$. The abstract MDP \mathcal{M}_{\simeq} is said to be a *refinement* of \mathcal{M}_{\equiv} . The relation $\text{rel}_{\simeq, \equiv}^{\alpha} \subseteq \mathbf{Q}_{\simeq} \times \mathbf{Q}_{\equiv}$ defined as $\{([q']_{\simeq}, [q]_{\equiv}) \mid [q']_{\simeq} \subseteq [q]_{\equiv}\}$ is said to be a refinement relation for $(\mathcal{M}_{\simeq}, \mathcal{M}_{\equiv})$.

The following is an immediate consequence of the definition.

PROPOSITION 4.6. *Let \simeq and \equiv be two equivalence relations compatible with the MDP \mathcal{M} such that $\simeq \subseteq \equiv$. Recall that the refinement relation for $(\mathcal{M}_{\simeq}, \mathcal{M}_{\equiv})$ is denoted by $\text{rel}_{\simeq, \equiv}^{\alpha}$. Then $\text{rel}_{\simeq, \equiv}^{\alpha}$ is a canonical simulation and $\text{rel}_{\simeq, \equiv}^{\alpha} \circ \text{rel}_{\simeq}^{\alpha} = \text{rel}_{\equiv}^{\alpha}$.*

Example 4.7. Consider for example, the DTMC \mathcal{M}_{ex} from Figure 8 where q_0 is the initial state and proposition P labels q_5 . Let \equiv be the equivalence relation which partitions the set $\{q_0, q_1, q_2, q_3, q_4, q_5\}$ into the equivalence classes $\{q_0\}$, $\{q_1, q_2\}$, $\{q_3, q_4\}$ and $\{q_5\}$. The resulting MDP \mathcal{M}_{ab} is given in Figure 9. Here the initial state is $\{q_0\}$ and P labels $\{q_5\}$.

5. COUNTEREXAMPLE GUIDED REFINEMENT

As described in Section 4, in our framework, a MDP \mathcal{M} will be abstracted by another MDP \mathcal{M}_{\equiv} defined on the basis of an equivalence relation \equiv on the states of \mathcal{M} . Model checking \mathcal{M}_{\equiv} against a safety property ψ_S will either tell us that ψ_S is satisfied by \mathcal{M}_{\equiv} (in which case, it is also satisfied by \mathcal{M} as shown in Lemma 2.11) or it is not. If $\mathcal{M}_{\equiv} \not\models \psi_S$ then \mathcal{M}_{\equiv} can be analyzed to obtain a minimal counterexample $(\mathcal{E}, \text{rel}_{\text{inj}})$, using the algorithm in Figure 5. The counterexample $(\mathcal{E}, \text{rel}_{\text{inj}})$ must be analyzed to decide whether $(\mathcal{E}, \text{rel}_{\text{inj}})$ proves that \mathcal{M} fails to satisfy ψ_S , or the counterexample is *spurious* and the abstraction (or rather the equivalence relation




Fig. 10. The three counterexamples $\mathcal{K}_{\text{cex}_1}$ (shown with short dashed edges), $\mathcal{K}_{\text{cex}_2}$ (shown with long dashed edges), and $\mathcal{K}_{\text{cex}_3}$ (shown with short and long dashed edges)

\equiv) must be refined to “eliminate” it. In order to carry out these steps, we need to first identify what it means for a counterexample to be *valid and consistent* for \mathcal{M} , describe and analyze an algorithm to check validity, and then demonstrate how the abstraction can be refined if the counterexample is spurious. In this section, we will outline our proposal to carry out these steps. We will frequently recall how these steps are carried out in the non-probabilistic case through a running example to convince the reader that our definitions are a natural generalization to the probabilistic case.

5.1 Checking Counterexamples

Checking if a counterexample proves that the system \mathcal{M} fails to meet its requirements ψ_S , intuitively, requires one to check if the “behavior” (or behaviors) captured by the counterexample are indeed exhibited by the system. The formal concept that expresses when a systems exhibits certain behaviors is *simulation*. Thus, one could potentially consider defining a valid counterexample to be one that is simulated by the MDP \mathcal{M} . However, as we illustrate in this section, the notion of valid counterexamples that is used in the context of non-probabilistic systems [Clarke et al. 2000; Clarke et al. 2002] is stronger. We, therefore, begin by motivating and formally defining when a counterexample is valid and consistent (Section 5.1.1), and then present and analyze the algorithm for checking validity (Section 5.1.2).

5.1.1 Validity and Consistency of Counterexamples. In the context of non-probabilistic systems, a valid counterexample is not simply one that is simulated by the original system. This is illustrated by the following example; we use this to motivate our generalization to probabilistic systems.

Example 5.1. Recall the Kripke structure \mathcal{K}_{ex} given in Example 4.1 along with the abstraction \mathcal{K}_{ab} (these structures are given in Figures 6 and 7, respectively). The LTL safety-property $\Box(\neg P)$ is violated by \mathcal{K}_{ab} . For such safety properties, counterexamples are just paths in \mathcal{K}_{ab} (which of course can be viewed as Kripke structures in their own right). The counterexample generation algorithms in [Clarke et al.

2000; Clarke et al. 2002] could possibly generate any one of three paths in \mathcal{K}_{ab} shown in Fig 10: $\mathcal{K}_{\text{cex}_1} = \{q_0, q_1, q_3\} \rightarrow \{q_2\} \rightarrow \{q_9\} \rightarrow \{q_{11}\}$, $\mathcal{K}_{\text{cex}_2} = \{q_0, q_1, q_3\} \rightarrow \{q_4\} \rightarrow \{q_{10}\} \rightarrow \{q_{11}\}$ and $\mathcal{K}_{\text{cex}_3} = \{q_0, q_1, q_3\} \rightarrow \{q_5, q_6\} \rightarrow \{q_7, q_8\} \rightarrow \{q_{11}\}$. Now each of the counterexamples $\mathcal{K}_{\text{cex}_1}$, $\mathcal{K}_{\text{cex}_2}$ and $\mathcal{K}_{\text{cex}_3}$ is simulated by \mathcal{K}_{ex} because \mathcal{K}_{ex} has a path, starting from the initial state, having 4 states, where only the fourth state satisfies proposition P . However, the algorithm outlined in [Clarke et al. 2000; Clarke et al. 2002] only considers $\mathcal{K}_{\text{cex}_1}$ to be valid. In order to see this, let us recall how the algorithm proceeds. The algorithm starts from the last state of the counterexample and proceeds backwards, checking at each point whether any of the concrete states corresponding to the abstract state in the counterexample can exhibit the counterexample from that point onwards. It also checks that the initial concrete state can exhibit the counterexample. Thus, $\mathcal{K}_{\text{cex}_2}$ is invalid because q_0 does not have a transition to q_4 and $\mathcal{K}_{\text{cex}_3}$ is invalid because none of q_0, q_1 , or q_3 have a transition to q_6 (the only state among q_5 and q_6 that can exhibit $\{q_5, q_6\} \rightarrow \{q_7, q_8\} \rightarrow \{q_{11}\}$).

The example above illustrates that in order to check validity, the algorithm searches for a simulation relation, wherein each (abstract) state of the counterexample is mapped to one of the concrete states that correspond to it, rather than an arbitrary simulation relation. Thus the “proof” for the validity of a counterexample in a concrete system, must be “contained” in the proof that demonstrates the validity of the counterexample in the abstract system. Based on this intuition we formalize the notion of when a counterexample is valid and consistent.

Definition 5.2. Let \mathcal{M} be a MDP with set of states \mathbf{Q} , and \equiv be an equivalence relation that is compatible with \mathcal{M} . Let ψ_S be a PCTL-safety formula such that $\mathcal{M}_{\equiv} \not\models \psi_S$ and let $(\mathcal{E}, \mathcal{R}_0)$ be a counterexample for \mathcal{M}_{\equiv} and ψ_S with set of states $\mathbf{Q}_{\mathcal{E}}$. We say that the counterexample $(\mathcal{E}, \mathcal{R}_0)$ is *valid* and *consistent with* (\mathcal{M}, \equiv) if there is a relation $\mathcal{R} \subseteq \mathbf{Q}_{\mathcal{E}} \times \mathbf{Q}$ such that

- (1) \mathcal{R} is a canonical simulation (validity); and
- (2) $\text{rel}_{\equiv}^{\alpha} \circ \mathcal{R} \subseteq \mathcal{R}_0$ (consistency).

The relation \mathcal{R} is said to be a *validating simulation*. If no such \mathcal{R} exists then $(\mathcal{E}, \mathcal{R}_0)$ is said to be *invalid for* (\mathcal{M}, \equiv) .

The above definition provides one technical reason for why it is convenient to view a counterexample as not just a MDP but rather as a MDP along with a simulation relation; we will see another justification for this when we discuss refinement.

Remark 5.3. When the counterexample $(\mathcal{E}, \text{rel}_{\text{inj}})$ is generated as in Theorem 3.16, $\mathbf{Q}_{\mathcal{E}} \subseteq \{\bar{a} \mid a \in \mathbf{Q}_{\equiv}\}$ and the relation $\text{rel}_{\text{inj}} = \{(\bar{a}, a) \mid \bar{a} \in \mathbf{Q}_{\mathcal{E}}\}$. In this case, please note that consistency is equivalent to requiring that $\mathcal{R} \subseteq \{(\bar{a}, q) \mid \bar{a} \in \mathbf{Q}_{\mathcal{E}} \text{ and } q \in \text{rel}_{\equiv}^{\gamma}(a)\}$. In other words, consistency is equivalent to requiring that $\mathcal{R} \subseteq \text{rel}_{\equiv}^{\gamma} \circ \text{rel}_{\text{inj}}$.

We conclude this section by showing that for minimal counterexamples the containment in the consistency requirement can be taken to be equality.

PROPOSITION 5.4. *Let \mathcal{M} be a MDP, \equiv an equivalence relation compatible with \mathcal{M} and ψ_S be a safety formula such that $\mathcal{M}_{\equiv} \not\models \psi_S$. If $(\mathcal{E}, \mathcal{R}_0)$ is a minimal*

```

Let  $\mathcal{E} = (\mathcal{Q}_{\mathcal{E}}, q_{\mathcal{E}}, \delta_{\mathcal{E}}, L_{\mathcal{E}})$  and  $\mathcal{M} = (\mathcal{Q}, q_{\mathcal{I}}, \delta, L)$ 
Initially  $R = \{(\bar{a}, q) \mid \bar{a} \in \mathcal{Q}_{\mathcal{E}} \text{ and } q \in \text{rel}_{\equiv}^{\gamma}(\bar{a})\}$  and  $R_{old} = \emptyset$ 
while ( $R_{old} \neq R$ ) do
   $R_{old} = R$ 
  For each state  $\bar{a} \in \mathcal{Q}_{\mathcal{E}}$  and each  $\mu \in \delta_{\mathcal{E}}(\bar{a})$  do
     $R = \{(\bar{b}, q) \in R \mid \bar{b} \neq \bar{a}\} \cup \{(\bar{a}, q) \in R \mid \exists \mu' \in \delta(q). \mu \preceq_{R_{old}} \mu'\}$ 
    If  $\mathcal{R}(\bar{a}) = \emptyset$  then return (“invalid”,  $\bar{a}, \mu, R_{old}, R$ )
    If  $\bar{a} = q_{\mathcal{E}}$  and  $q_{\mathcal{I}} \notin R(\bar{a})$  then return (“invalid”,  $\bar{a}, \mu, R_{old}, R$ )
  od  $\leftarrow$  end of For loop
od  $\leftarrow$  end of while loop
Return (“valid”)

```

Fig. 11. Algorithm for checking validity and consistency of counterexamples

counterexample for \mathcal{M}_{\equiv} and ψ_S and $(\mathcal{E}, \mathcal{R}_0)$ is consistent and valid for (\mathcal{M}, \equiv) with validating simulation \mathcal{R} then $\text{rel}_{\equiv}^{\alpha} \circ \mathcal{R} = \mathcal{R}_0$.

PROOF. We have that \mathcal{R} is a simulation and $\text{rel}_{\equiv}^{\alpha} \circ \mathcal{R} \subseteq \mathcal{R}_0$. Since $\text{rel}_{\equiv}^{\alpha}$ and \mathcal{R} are simulations, so is $\text{rel}_{\equiv}^{\alpha} \circ \mathcal{R}$. As $\mathcal{E} \not\models \psi_S$, we get that $(\mathcal{E}, \text{rel}_{\equiv}^{\alpha} \circ \mathcal{R})$ is also a counterexample for \mathcal{M}_{\equiv} and ψ_S . Thus, if $\text{rel}_{\equiv}^{\alpha} \circ \mathcal{R} \subsetneq \mathcal{R}_0$, then $(\mathcal{E}, \text{rel}_{\equiv}^{\alpha} \circ \mathcal{R})$ is not minimal. Hence $\text{rel}_{\equiv}^{\alpha} \circ \mathcal{R} = \mathcal{R}_0$. \square

5.1.2 *Algorithm to check Validity of Counterexamples.* We now present the algorithm to check the validity and consistency of a counterexample. We will assume that the counterexample is a minimal one, generated by the algorithm in Theorem 3.16. Thus the counterexample is of the form $(\mathcal{E}, \text{rel}_{\text{inj}})$, where the set of states $\mathcal{Q}_{\mathcal{E}}$ is a subset of $\{\bar{a} \mid a \in \mathcal{Q}_{\equiv}\}$ and the relation rel_{inj} is $\{(\bar{a}, a) \mid \bar{a} \in \mathcal{Q}_{\mathcal{E}}\}$. The algorithm for counterexample checking is then the standard simulation checking algorithm [Baier et al. 2000] that computes the validating simulation through progressive refinement, except that in our case we start with $R = \{(\bar{a}, q) \mid \bar{a} \in \mathcal{Q}_{\mathcal{E}} \text{ and } q \in \text{rel}_{\equiv}^{\gamma}(a)\}$. The algorithm is shown in Figure 11. Please note that for the rest of the paper (and in the algorithm), for $\mu \in \text{Prob}_{\leq 1}(\mathcal{Q}_{\mathcal{E}})$, $\mu' \in \text{Prob}_{\leq 1}(\mathcal{Q})$ and $\mathcal{R} \subseteq \mathcal{Q}_{\mathcal{E}} \times \mathcal{Q}$, by $\mu \preceq_{\mathcal{R}} \mu'$ we mean $\mu \preceq_{\text{id}_{\mathcal{Q}_{\mathcal{E}}} \cup \mathcal{R} \cup \text{id}_{\mathcal{Q}}} \mu'$.

We will now show that the algorithm in Figure 11 is correct. We start by showing that the algorithm terminates.

PROPOSITION 5.5. *The counterexample checking algorithm shown in Figure 11 terminates.*

PROOF. Let \mathcal{R}_0^n and \mathcal{R}_1^n be respectively the relations denoted by the variable R at the beginning and the end of the n th iteration of the while loop. A simple inspection of the algorithm tells us that either $\mathcal{R}_1^n = \mathcal{R}_0^n$ or $\mathcal{R}_1^n \subsetneq \mathcal{R}_0^n$. If $\mathcal{R}_1^n = \mathcal{R}_0^n$ then the while loop terminates (and returns “valid”). Otherwise the size of relation denoted by variable R decreases by at least one. Hence, if \mathcal{R}_1^n is never equal to \mathcal{R}_0^n then it must be case that $\mathcal{R}_1^n(\bar{a})$ becomes empty for some n and some $\bar{a} \in \mathcal{Q}_{\mathcal{E}}$ and then the algorithm terminates. \square

We now show that if the algorithm returns “valid” then the counterexample $(\mathcal{E}, \text{rel}_{\text{inj}})$ is valid and consistent.

PROPOSITION 5.6. *If the algorithm in Figure 11 returns “valid” then the counterexample $(\mathcal{E}, \text{rel}_{\text{inj}})$ is valid and consistent for (\mathcal{M}, \equiv) .*

PROOF. Please note that the algorithm returns “valid” only when the while loop terminates. At that point the variables R_{old} and R denote the same relation, which we will call \mathcal{R} for the rest of the proof. Please note as $\mathcal{R} \subseteq \{(\bar{a}, q) \mid \bar{a} \in \mathbb{Q}_{\mathcal{E}} \text{ and } q \in \text{rel}_{\equiv}^{\gamma}(\bar{a})\}$, $\text{rel}_{\equiv}^{\alpha} \circ \mathcal{R} \subseteq \text{rel}_{\text{inj}}$. Hence, the result will follow if we can show that \mathcal{R} is a canonical simulation. Consider the last iteration of the while loop. Now, a simple inspection of the algorithm says that the variable R does not change its value in this iteration. The only way this is possible is if for each $\bar{a} \in \mathbb{Q}_{\mathcal{E}}$, $\mu \in \delta_{\mathcal{E}}(\bar{a})$ and each $(\bar{a}, q) \in \mathcal{R}$, there exists $\mu' \in \delta(q)$ such that $\mu \preceq_{\mathcal{R}} \mu'$. Also, $q_{\mathcal{I}} \in \mathcal{R}(q_{\mathcal{E}})$. Thus \mathcal{R} is a canonical simulation. \square

We now show that the if $(\mathbb{Q}_{\mathcal{E}}, \text{rel}_{\text{inj}})$ is valid and consistent, then the algorithm in Figure 11 must return “valid”.

PROPOSITION 5.7. *If the counterexample $(\mathbb{Q}_{\mathcal{E}}, \text{rel}_{\text{inj}})$ is valid and consistent with (\mathcal{M}, \equiv) then the algorithm in Figure 11 returns “valid”.*

PROOF. Assume that the counterexample $(\mathbb{Q}_{\mathcal{E}}, \text{rel}_{\text{inj}})$ is valid and consistent. Thus there is a canonical simulation $\mathcal{R} \subseteq \mathbb{Q}_{\mathcal{E}} \times \mathbb{Q}_{\equiv}$ such that $\text{rel}_{\equiv}^{\alpha} \circ \mathcal{R} = \text{rel}_{\text{inj}}$ (equality is a consequence of minimality; see Proposition 5.4). We make the following observations:

- (1) As $\text{rel}_{\equiv}^{\alpha} \circ \mathcal{R} = \text{rel}_{\text{inj}}$, $\mathcal{R} \subseteq \{(\bar{a}, q) \mid \bar{a} \in \mathbb{Q}_{\mathcal{E}} \text{ and } q \in \text{rel}_{\equiv}^{\gamma}(a)\}$.
- (2) For each $\bar{a} \in \mathbb{Q}_{\mathcal{E}}$, $\mathcal{R}(\bar{a}) \neq \emptyset$ (otherwise (\bar{a}, a) will be present in rel_{inj} but not in $\text{rel}_{\equiv}^{\alpha} \circ \mathcal{R}$).
- (3) $q_{\mathcal{E}} \mathcal{R} q_{\mathcal{I}}$ (\mathcal{R} is a simulation).
- (4) For each $\bar{a} \in \mathbb{Q}_{\mathcal{E}}$, each $\mu \in \delta_{\mathcal{E}}(\bar{a})$ and each $q \in \mathbb{Q}$ such that $\bar{a} \mathcal{R} q$ there exists a $\mu' \in \delta(q)$ such that $\mu \preceq_{\mathcal{R}} \mu'$ (\mathcal{R} is a simulation).
- (5) For any relation $\mathcal{R}_1 \subseteq \{(\bar{a}, q) \mid \bar{a} \in \mathbb{Q}_{\mathcal{E}} \text{ and } q \in \text{rel}_{\equiv}^{\gamma}(a)\}$ such that $\mathcal{R} \subseteq \mathcal{R}_1$, $\mu \in \text{Prob}_{\leq 1}(\mathbb{Q}_{\mathcal{E}})$ and $\mu' \in \text{Prob}_{\leq 1}(\mathbb{Q})$ we have that $\mu \preceq_{\mathcal{R}} \mu'$ implies that $\mu \preceq_{\mathcal{R}_1} \mu'$.

Now, the first observation above implies that in the algorithm in Figure 11 initially \mathcal{R} is contained within the relation denoted by the variable R . \mathcal{R} is also contained in the relation denoted by the variable R_{old} , the first time the variable R_{old} takes a non-empty value. From this point on, we claim \mathcal{R} is always contained in the relations R_{old} and R . This claim is a consequence of the fourth and the fifth observations which ensure that every time R is updated, \mathcal{R} is contained in the relation denoted by R . Finally note that second and third observations ensure that the algorithm can never declare the counterexample to be invalid and hence by termination (Proposition 5.5), the algorithm must return “valid”. \square

A careful analysis of the special structure of the validating simulation yields better bounds than that reported in [Baier et al. 2000] for general simulation.

THEOREM 5.8. *Let \mathcal{M} be a MDP, \equiv an equivalence relation compatible with \mathcal{M} , and ψ_S a safety property such that $\mathcal{M}_{\equiv} \not\models \psi_S$. Let $(\mathcal{E}, \text{rel}_{\text{inj}})$ be a counterexample for \mathcal{M}_{\equiv} and ψ_S generated using Theorem 3.16. Then the algorithm in Figure 11*

returns “valid” iff $(\mathcal{E}, \text{rel}_{\text{inj}})$ is valid and consistent with (\mathcal{M}, \equiv) . Let $n_{\mathcal{M}}$ and $m_{\mathcal{M}}$ be the number of vertices and edges, respectively, in the underlying labeled graph $G_{\ell}(\mathcal{M})$. The algorithm shown in Figure 11 runs in time $O(n_{\mathcal{M}}^2 m_{\mathcal{M}}^2)$.

PROOF. Thanks to Propositions 5.6 and 5.7, the result will follow if we show that the running time of the algorithm is $O(n_{\mathcal{M}}^2 m_{\mathcal{M}}^2)$.

Observe that the outermost while loop runs for as long as R changes. If for each $\bar{a} \in \mathbf{Q}_{\mathcal{E}}$, we define $s_a = |\text{rel}_{\equiv}^{\gamma}(a)| = |\text{rel}_{\equiv}^{\gamma} \circ \text{rel}_{\text{inj}}(\bar{a})|$, a bound on the number of iterations of the outermost loop is $\sum_{\bar{a} \in \mathbf{Q}_{\mathcal{E}}} s_a = n_{\mathcal{M}}$, as each state of \mathbf{Q} belongs to at most one $\text{rel}_{\equiv}^{\gamma}(a)$. Next let us define d_a to be number of outgoing edges from the set $\text{rel}_{\equiv}^{\gamma}(a)$ and d_q to be the out-degree of q in $G_{\ell}(\mathcal{M})$. Clearly for each state $\bar{a} \in \mathbf{Q}_{\mathcal{E}}$, the total number of tests of the form $\mu \preceq_{R_{\text{old}}} \mu'$ is bounded by $\sum_{q \in \text{rel}_{\equiv}^{\gamma}(a)} d_a d_q = d_a^2$. Thus, the total number of tests $\mu \preceq_{R_{\text{old}}} \mu'$ in a single iteration of the outermost loop is bounded by $\sum_{\bar{a} \in \mathbf{Q}_{\mathcal{E}}} d_a^2 \leq m_{\mathcal{M}}^2$. Now, because each state of \mathbf{Q} belongs to at most one $\text{rel}_{\equiv}^{\gamma} \circ \text{rel}_{\text{inj}}(\bar{a})$, the test $\mu \preceq_{R_{\text{old}}} \mu'$ simply requires one to check that for each $\bar{b} \in \mathbf{Q}_{\mathcal{E}}$, $\mu(\bar{b}) \leq \sum_{q: (\bar{b}, q) \in R_{\text{old}}} \mu'(q)$ and thus can be done in $O(n_{\mathcal{M}})$ time. Thus we don't need to compute flows in bipartite networks, as is required in the general case [Baier et al. 2000]. The total running time is, therefore, $O(n_{\mathcal{M}}^2 m_{\mathcal{M}}^2)$. \square


We make some observations about the validity checking algorithm.

(1) For linear time properties and non-probabilistic systems, checking validity of a counterexample simply determines if the first state in the counterexample can be simulated by the initial state of the system by going backwards from the last state of the counterexample. The same idea can be exploited for probabilistic systems as well if the underlying unlabeled graph of the counterexample \mathcal{E} is a tree (or more generally a DAG); the resulting algorithm will depend on the height of the counterexample and cut the running time of the algorithm by a factor of $n_{\mathcal{M}}$.

(2) Theorem 3.4 observes that for weak safety formulas, counterexamples whose underlying graph is a tree can be found by “unrolling” the minimal MDP counterexample. Instead of first explicitly unrolling the counterexample and then checking, one can unroll the counterexample “on the fly” while checking validity. This algorithm is presented in Section 5.3, after our discussion on refinement so as to not interrupt the flow of the paper. The crucial idea is to decide when to stop unrolling which is made by keeping track of the satisfaction of subformulas at various states. The running time of the algorithm will be $O(h \cdot n_{\mathcal{M}} m_{\mathcal{M}}^2)$, where h is the height of the unrolled tree. Thus depending on $n_{\mathcal{M}}$ and h , one could either compute the actual simulation relation, or simply check whether the tree of height h is simulated.

(3) One can construct the graph of maximal strongly connected components of $G_{\ell}(\mathcal{E})$, and compute the simulation relation on each maximal strongly connect component, in the order of their topological sort. While this new algorithm will not yield better asymptotic bounds, it may work better in practice.

To complete the description of the CEGAR approach, all we need to do is describe the refinement step. However, before we proceed, we describe a result and give some notations for the case when the counterexample generated by the algorithm in Figure 5 is declared by the counterexample checking algorithm to be invalid.

Fig. 12. Counterexample \mathcal{E} for MDP \mathcal{M}_{ab} in Figure 9

PROPOSITION 5.9. *If the algorithm in Figure 11 returns (“invalid”, \bar{a} , μ , R_{old} , R), then for all $q \in R_{old}(\bar{a}) \setminus R(\bar{a})$ and all $\mu_1 \in \delta(q)$, $\mu \not\leq_{R_{old}} \mu_1$.*

PROOF. Immediate consequence of the algorithm. \square

Notation 5.10. If the algorithm in Figure 11 returns (“invalid”, \bar{a} , μ , R_{old} , R) then

- \bar{a} is said to be an *invalidating abstract state*;
- $\mu \in \delta_{Q_\varepsilon}(\bar{a})$ is said to be an *invalidating transition*; and
- the pair (R_{old}, R) is said to be the *invalidating witness* for \bar{a} and μ .


Example 5.11. Consider for example, the DTMC \mathcal{M}_{ex} and its abstraction \mathcal{M}_{ab} from Example 4.7, shown in Figures 8 and 9 respectively. Recall that q_0 is the initial state and proposition P labels q_5 . Consider the safety formula $\psi_S = \mathcal{P}_{\leq \frac{2}{3}}(\mathbf{X}(\neg \mathcal{P}_{< \frac{1}{2}}(\mathbf{X}P)))$. The MDP \mathcal{M}_{ab} violates ψ_S and the counterexample generation algorithm will output the counterexample $(\mathcal{E}, \text{rel}_{inj})$ where \mathcal{E} shown in Figure 12 is a DTMC.⁴ The relation rel_{inj} is $\{(\bar{a}, a) \mid a \text{ is a state of } \mathcal{M}_{ab}\}$. Now, $(\mathcal{E}, \text{rel}_{inj})$ is not valid for \mathcal{M}_{ab} and the counterexample checking algorithm will output “invalid” in the second iteration of the while loop in Figure 11. The invalidating abstract state is the initial state $\overline{\{q_0\}}$, the invalidating transition is the transition out of $\overline{\{q_0\}}$ and invalidating witness will be (R_{old}, R) where

- $R_{old} = \{(\overline{\{q_0\}}, \{q_0\}), (\overline{\{q_1, q_2\}}, \{q_2\}), (\overline{\{q_3, q_4\}}, \{q_3\}), (\overline{\{q_5\}}, \{q_5\})\}$, and
- $R = \{(\{q_1, q_2\}, \{q_2\}), (\{q_3, q_4\}, \{q_3\}), (\{q_5\}, \{q_5\})\}$.

5.2 Refining Abstractions

The last step in the abstraction-refinement loop is to refine the abstraction in the case when the counterexample is invalid. The algorithm in Figure 11, concludes the invalidity of the counterexample, when it finds some abstract state a such that \bar{a} is a state of the counterexample and \bar{a} is not simulated by any concrete state in $\text{rel}_\gamma(a)$, or when \bar{a} is the initial state of the counterexample and it is not simulated by the initial state of \mathcal{M} . At this point, we will refine the abstraction by refining the equivalence \equiv that was used to construct the abstract model in the first place. The goal of the refinement step is for it to be “counterexample guided”. The ideal

⁴Please note that in the figures illustrating counterexamples, we will confuse \bar{a} with a .

Fig. 13. The abstraction \mathcal{K}'_{ab}

situation is one where the spurious counterexample is “eliminated” by the refinement step. However, as we remind the reader, this is not achieved in the CEGAR approach for non-probabilistic systems. We, therefore, begin (Section 5.2.1) by motivating and defining the notion of a “good refinement”. We show that good refinements do indeed lead to progress in the CEGAR approach. After this, in Section 5.2.2, we present a refinement algorithm along with a proof that it results in good refinements.

5.2.1 Good Refinements. We begin by recalling the refinement step in the CEGAR approach for non-probabilistic systems through an example, to demonstrate that refinement does not lead to the elimination of the counterexample. The example, however, motivates what the refinement step does indeed achieve, leading us to the notion of good refinements.

Example 5.12. As in Example 5.1, consider the Kripke structure \mathcal{K}_{ex} from Example 4.1 along with the abstraction \mathcal{K}_{ab} (these structures are also given in Figures 6 and 7). The LTL safety-property $\phi = \Box(\neg P)$ is violated by \mathcal{K}_{ab} and the counterexample generation algorithm may generate the spurious counterexample $\mathcal{K}_{cex_3} = \{q_0, q_1, q_3\} \rightarrow \{q_5, q_6\} \rightarrow \{q_7, q_8\} \rightarrow \{q_{11}\}$. Now, the counterexample checking algorithm in [Clarke et al. 2000; Clarke et al. 2002] starts from the last state of the counterexample and proceeds backwards, checking at each point whether any of the concrete states corresponding to the abstract state in the counterexample can exhibit the counterexample from that point onwards. The algorithm finds that there is a path $q_6 \rightarrow q_7 \rightarrow q_{11}$ in \mathcal{K}_{ex} which simulates $\{q_5, q_6\} \rightarrow \{q_7, q_8\} \rightarrow \{q_{11}\}$ but finds that there is no transition from $\{q_0, q_1, q_3\}$ to q_6 . At this point, it declares the counterexample to be invalid. The refinement step then breaks the equivalence class $\text{post}(\{q_0, q_1, q_3\}) = \{q_5, q_6\}$ into $\{q_6\}$ and $\{q_5\}$. The resulting abstraction \mathcal{K}'_{ab} shown in Figure 13 still has the “spurious” counterexample $\mathcal{K}'_{cex_3} = \{q_0, q_1, q_3\} \rightarrow \{q_5\} \rightarrow \{q_7, q_8\} \rightarrow \{q_{11}\}$ “contained” within \mathcal{K}_{cex_3} .

Though, in Example 5.12, the counterexample is not eliminated by the refinement, progress has been, nonetheless, made. Considering the example carefully, one notes that breaking $\{q_5, q_6\}$ could have yielded two possible new paths: $\mathcal{K}'_{cex_3} =$

$\{q_0, q_1, q_3\} \rightarrow \{q_5\} \rightarrow \{q_7, q_8\} \rightarrow \{q_{11}\}$ and $\mathcal{K}''_{\text{ce}x_3} = \{q_0, q_1, q_3\} \rightarrow \{q_6\} \rightarrow \{q_7, q_8\} \rightarrow \{q_{11}\}$. However, only one path $\mathcal{K}'_{\text{ce}x_3}$ is exhibited by \mathcal{K}'_{ab} while the other path $\mathcal{K}''_{\text{ce}x_3}$ is not exhibited by \mathcal{K}'_{ab} and hence has been “eliminated”. Thus, what is eliminated is at least one simulation relation that is “contained” in the original spurious counterexample. We capture this concept for MDPs as follows.

Definition 5.13. Let \mathcal{M} be a MDP with states \mathbf{Q} , \simeq and \equiv be equivalence relations on \mathbf{Q} compatible with \mathcal{M} such that $\simeq \subseteq \equiv$, and ψ_S be a safety property. Let $(\mathcal{E}, \mathcal{R})$ be a counterexample for \mathcal{M}_{\equiv} and ψ_S , where $\mathbf{Q}_{\mathcal{E}}$ is the set of states of \mathcal{E} with initial state $q_{\mathcal{E}}$. Finally let q_{\simeq} be the initial state of \mathcal{M}_{\simeq} . We say that \simeq is a *good \equiv -refinement* for $(\mathcal{E}, \mathcal{R})$ if there is *some* $\mathcal{R}' \subseteq \mathbf{Q}_{\mathcal{E}} \times \mathbf{Q}_{\simeq}$ such that $(q_{\mathcal{E}}, q_{\simeq}) \in \mathcal{R}'$, $\text{rel}_{\simeq, \equiv}^{\alpha} \circ \mathcal{R}' = \mathcal{R}$ but \mathcal{R}' is not a canonical simulation (of \mathcal{E} by \mathcal{M}_{\simeq}). If no such \mathcal{R}' exists, we say that \simeq is a *bad \equiv -refinement* for $(\mathcal{E}, \mathcal{R})$.

Observe that the conditions on \mathcal{R}' ensure that \mathcal{R}' is one of the possible proofs that \mathcal{M} violates ψ_S “contained” within the counterexample $(\mathcal{E}, \mathcal{R})$. Intuitively, a good \equiv -refinement \simeq ensures that $(\mathcal{E}, \mathcal{R}')$ is not a counterexample for \mathcal{M}_{\simeq} and ψ_S . This presents yet another justification for formally treating the simulation relation (or proof) as part of the notion of a counterexample. In the absence of the simulation relation, it is difficult to justify why refinement is “counterexample guided” given that the behavior (i.e., \mathcal{E}) itself may not be eliminated.

Remark 5.14. Before presenting the consequences of good refinements, we would like to examine the formal definition more carefully, in order to highlight the subtle reasons why all the points in the definition are needed.

(1) Observe that \mathcal{R}' satisfying $\text{rel}_{\simeq, \equiv}^{\alpha} \circ \mathcal{R}' = \mathcal{R}$ always exist: take \mathcal{R}' to be any relation such that for each $q_0 \in \mathbf{Q}_{\mathcal{E}}$, $\mathcal{R}'(q_0) = \cup_{[q]_{\equiv} \in \mathcal{R}(q_0)} X_{[q]_{\equiv}}$ where $X_{[q]_{\equiv}}$ is any non-empty subset of $\{[q_1]_{\simeq} \in \mathbf{Q}_{\simeq} \mid [q_1]_{\simeq} \subseteq [q]_{\equiv}\}$. Further, any \mathcal{R}' such that $\text{rel}_{\simeq, \equiv}^{\alpha} \circ \mathcal{R}' = \mathcal{R}$ must be of this form.

(2) We demand $\text{rel}_{\simeq, \equiv}^{\alpha} \circ \mathcal{R}' = \mathcal{R}$ rather than $\text{rel}_{\simeq, \equiv}^{\alpha} \circ \mathcal{R}' \subseteq \mathcal{R}$. One can easily come up with \mathcal{R}' which satisfies $\text{rel}_{\simeq, \equiv}^{\alpha} \circ \mathcal{R}' \subseteq \mathcal{R}$ but is not a simulation (take, e.g. $\mathcal{R}' = \emptyset$). Note also that if $(\mathcal{E}, \mathcal{R})$ is a minimal counterexample then the set $\mathcal{R}(q_0)$ is non-empty for each $q_0 \in \mathbf{Q}_{\mathcal{E}}$. Thus, by taking \simeq to be \equiv and taking $X_{[q]_{\equiv}}$ to be empty for some $[q]_{\equiv}$ above, we will ensure that the resulting \mathcal{R}' is not a simulation and $\text{rel}_{\simeq, \equiv}^{\alpha} \circ \mathcal{R}' \subseteq \mathcal{R}$. Hence we would have declared \equiv to be a good \equiv -refinement if we had not required the equality! Thus we may not be able to guarantee any progress in the CEGAR loop (see Proposition 5.16).

(3) We require that $(q_{\mathcal{E}}, q_{\simeq}) \in \mathcal{R}'$. If we had not required this condition then one could have achieved a “good” refinement by just breaking the initial state and taking \mathcal{R}' to be any relation such that $\text{rel}_{\simeq, \equiv}^{\alpha} \circ \mathcal{R}' = \mathcal{R}$ but fails to be a simulation just because \mathcal{R}' relates $q_{\mathcal{E}}$ to equivalence classes that does not contain the initial state of \mathcal{M}_{\simeq} .

Suppose $(\mathcal{E}, \mathcal{R})$ is a counterexample for \mathcal{M}_{\equiv} and ψ_S ; hence $\mathcal{E} \not\models \psi_S$. Therefore, if $\mathcal{R}' \subseteq \mathbf{Q}_{\mathcal{E}} \times \mathbf{Q}_{\simeq}$ is a canonical simulation then $(\mathcal{E}, \mathcal{R}')$ is a counterexample for \mathcal{M}_{\simeq} and ψ_S . The following proposition says that if $\text{rel}_{\simeq, \equiv}^{\alpha} \circ \mathcal{R}' = \mathcal{R}$ and $(\mathcal{E}, \mathcal{R})$ is invalid for (\mathcal{M}, \equiv) , then $(\mathcal{E}, \mathcal{R}')$ is invalid for (\mathcal{M}, \simeq) . Therefore, a good refinement ensures that at least one of counterexamples “contained” within $(\mathcal{E}, \mathcal{R})$ is not a

counterexample, and thereby eliminates at least one spurious counterexample that would not be eliminated by a bad refinement.

PROPOSITION 5.15. *Let \mathcal{M} be a MDP with Q as the set of states, \simeq and \equiv be equivalence relations compatible with \mathcal{M} such that $\simeq \subseteq \equiv$. Let $(\mathcal{E}, \mathcal{R})$ be a counterexample for \mathcal{M}_{\equiv} and ψ_S with $Q_{\mathcal{E}}$ as the set of states. Let $\mathcal{R}' \subseteq Q_{\mathcal{E}} \times Q_{\simeq}$ be a canonical simulation such that $\text{rel}_{\simeq, \equiv}^{\alpha} \circ \mathcal{R}' = \mathcal{R}$. Then $(\mathcal{E}, \mathcal{R}')$ is a counterexample for \mathcal{M}_{\simeq} and ψ_S . Further, if $(\mathcal{E}, \mathcal{R})$ is invalid for (\mathcal{M}, \equiv) then $(\mathcal{E}, \mathcal{R}')$ is invalid for (\mathcal{M}, \simeq) .*

PROOF. From the fact that \mathcal{R}' is a simulation and that $\mathcal{E} \not\models \psi_S$, it follows that $(\mathcal{E}, \mathcal{R}')$ is a counterexample for \mathcal{M}_{\simeq} and ψ_S . Assume, by way of contradiction, that $(\mathcal{E}, \mathcal{R}')$ is valid and consistent with (\mathcal{M}, \simeq) . Then there exists a canonical simulation $\mathcal{R}_0 \subseteq Q_{\mathcal{E}} \times Q$ such that $\text{rel}_{\simeq}^{\alpha} \circ \mathcal{R}_0 \subseteq \mathcal{R}'$. This implies that $(\text{rel}_{\simeq, \equiv}^{\alpha} \circ \text{rel}_{\simeq}^{\alpha}) \circ \mathcal{R}_0 \subseteq \text{rel}_{\simeq, \equiv}^{\alpha} \circ \mathcal{R}'$. But the left hand side is $\text{rel}_{\equiv}^{\alpha} \circ \mathcal{R}_0$ while the right hand side is \mathcal{R} . This implies that $(\mathcal{E}, \mathcal{R})$ is a valid and consistent with (\mathcal{M}, \equiv) (with \mathcal{R}_0 as the validating simulation). A contradiction! \square

We conclude this section by showing that good refinements ensure progress in the CEGAR loop.

PROPOSITION 5.16. *Let \mathcal{M} be a MDP, and \simeq and \equiv be equivalence relations compatible with \mathcal{M} such that $\simeq \subseteq \equiv$. Let $(\mathcal{E}, \mathcal{R})$ be a counterexample for \mathcal{M}_{\equiv} and ψ_S . If \simeq is a good \equiv -refinement for $(\mathcal{E}, \mathcal{R})$, then $\simeq \subseteq \equiv$.*

PROOF. Fix a relation $\mathcal{R}' \subseteq Q_{\simeq} \times Q$ such that $\text{rel}_{\simeq, \equiv}^{\alpha} \circ \mathcal{R}' = \mathcal{R}$ but \mathcal{R}' is not a canonical simulation. We now proceed by contradiction. Assume that $\simeq = \equiv$. Then \mathcal{M}_{\equiv} and \mathcal{M}_{\simeq} are the same abstract MDP and $\text{rel}_{\simeq, \equiv}^{\alpha}$ is the identity relation. Thus \mathcal{R}' and \mathcal{R} are the same relation. But \mathcal{R}' is not a simulation which contradicts that fact that $(\mathcal{E}, \mathcal{R})$ is a counterexample for \mathcal{M}_{\equiv} and ψ_S . \square

5.2.2 Algorithm for Refinement. In this section, we will show how an abstract model can be refined based on a spurious counterexample obtained as in Theorem 3.16. Before presenting our algorithm, we recall how the refinement step proceeds for non-probabilistic systems, through an example. This will help us highlight a couple of key points about the refinement step in the non-probabilistic case.

Example 5.17. Recall in the (non-probabilistic) Example 5.12, the counterexample checking step proceeded from the last state of the counterexample trace $\mathcal{K}_{\text{cex}_3} = \{q_0, q_1, q_3\} \rightarrow \{q_5, q_6\} \rightarrow \{q_7, q_8\} \rightarrow \{q_{11}\}$ and confirmed that the concrete state q_6 can simulate the path $\{q_5, q_6\} \rightarrow \{q_7, q_8\} \rightarrow \{q_{11}\}$. Since there is no “concrete” transition from $\{q_0, q_1, q_3\}$ to q_6 , the counterexample checking step concludes that the counterexample is invalid. The state $\{q_0, q_1, q_3\}$ is the counterpart of “invalidating abstract state”, the transition $t = \{q_0, q_1, q_3\} \rightarrow \{q_5, q_6\}$ is the counterpart of “invalidating transition”, $\{q_6\}$ is counterpart of $R_{\text{old}}(\{q_5, q_6\})$ and \emptyset is counterpart of $R(\{q_0, q_1, q_3\})$. The refinement step for non-probabilistic case is obtained by splitting the equivalence class $\text{post}(t, \{q_0, q_1, q_3\}) = \{q_5, q_6\}$ into $\{q_6\} = R_{\text{old}}(\{q_5, q_6\})$ and $\{q_5\} = \{q_5, q_6\} \setminus R_{\text{old}}(\{q_5, q_6\})$.

On the other hand, suppose for the Kripke structure \mathcal{K}_{ex} and its abstraction \mathcal{K}_{ab} , the counterexample $\mathcal{K}_{\text{cex}_2} = \{q_0, q_1, q_3\} \rightarrow \{q_4\} \rightarrow \{q_{10}\} \rightarrow \{q_{11}\}$ is chosen instead

of $\mathcal{K}_{\text{cex}_3}$. In this case the counterexample generation algorithm declares the counterexample to be invalid when q_0 , the initial state of \mathcal{K}_{ex} , fails to be in $R(\{q_0, q_1, q_3\})$ during the counterexample checking algorithm. In this case, $\{q_0, q_1, q_3\}$ is the “invalidating” abstract state; $\{q_0, q_2, q_3\}$ is the counterpart of $R_{\text{old}}(\{q_0, q_1, q_3\})$ and $\{q_3\}$ is the counterpart of $R(\{q_0, q_1, q_3\})$. For this case, the invalidating abstract state $\{q_0, q_1, q_3\}$ is itself broken into $R_{\text{old}}(\{q_0, q_1, q_3\}) \setminus R(\{q_0, q_1, q_3\}) = \{q_0, q_1\}$ and $\{q_0, q_1, q_3\} \setminus (R_{\text{old}}(\{q_0, q_1, q_3\}) \setminus R(\{q_0, q_1, q_3\})) = \{q_3\}$. Thus, in this case the invalidating abstract state itself is broken into equivalence classes and not its successor!

Let us examine the refinement step outlined in Example 5.17 more carefully. There are two cases to consider: when the invalidating abstract state is not the initial state, where we only split abstract state that is the target of the invalidating transition; and when the invalidating abstract state is the initial state of the counterexample, where we also split the invalidating abstract state.

To generalize to probabilistic systems, we make the following observations. Suppose \mathcal{M}_{\equiv} is the abstraction of \mathcal{M} with respect to \equiv , and let $(\mathcal{E}, \text{rel}_{\text{inj}})$ be a counterexample for \mathcal{M}_{\equiv} and ψ_S obtained as in Theorem 3.16 and which is invalid for (\mathcal{M}, \equiv) . If \bar{a} is the invalidating abstract state, and $\mu \in \delta_{\mathcal{E}}(\bar{a})$ is the invalidating transition, then $\text{post}(\mu, \bar{a})$ may contain several states (including \bar{a} itself). Hence, our refinement step will be forced to split several equivalence classes instead of one as in the case of non-probabilistic systems. Next, we observe that in the case when the counterexample is a “path” (or more generally a DAG), the algorithm to check validity only needs to “process” each state of the counterexample once. Hence, if (R_{old}, R) is the invalidating witness, then $R_{\text{old}}(\bar{a}) = \text{rel}_{\equiv}^{\gamma}(a)$. Therefore, $R_{\text{old}}(\bar{a}) \setminus R(\bar{a})$ is always $\text{rel}_{\equiv}^{\gamma}(a)$ except (possibly) when the invalidating state \bar{a} is the initial state of the counterexample. This is the primary reason why for non-probabilistic systems the invalidating abstract state is never split, except when it is the initial state. However, when analyzing counterexamples that could be general MDPs, the counterexample checking algorithm will need to “process” each state multiple times, and then $R_{\text{old}}(\bar{a})$ need not be $\text{rel}_{\equiv}^{\gamma}(a)$, at the time the counterexample is deemed to be invalid. Thus, in our refinement algorithm, we will be forced to also split the invalidating abstract state.

The above intuitions are formalized in the refinement step shown in Figure 14; recall that for each $\bar{d} \in \mathcal{Q}_{\mathcal{E}}$, we have $R(\bar{d}) \subseteq R_{\text{old}}(\bar{d}) \subseteq \text{rel}_{\gamma}(d)$. We conclude by showing that the resulting refinement is a good refinement (and hence progress is ensured in the CEGAR loop).

THEOREM 5.18. *Let $(\mathcal{E}, \text{rel}_{\text{inj}})$ be a counterexample, generated using Theorem 3.16, for \mathcal{M}_{\equiv} and safety property ψ_S , where \mathcal{M}_{\equiv} is the abstraction of \mathcal{M} with respect to the compatible equivalence relation \equiv . If the counterexample checking algorithm in Figure 11 returns (“invalid”, \bar{a} , μ , R_{old} , R), then the refinement $\simeq_{\subseteq \equiv}$ obtained as in Figure 14 is a good \equiv -refinement for $(\mathcal{E}, \text{rel}_{\text{inj}})$.*

PROOF. Let $\mathcal{E} = (\mathcal{Q}_{\mathcal{E}}, q_{\mathcal{E}}, \delta_{\mathcal{E}}, L_{\mathcal{E}})$. Recall that $\mathcal{M}_{\simeq} = (\mathcal{Q}_{\simeq}, q_{\simeq}, \delta_{\simeq}, L_{\simeq})$ is the abstract MDP for \mathcal{M} and \simeq , where the set of states of \mathcal{M}_{\simeq} are equivalence classes under \simeq . Consider the *functional* relation $\mathcal{R} \subseteq \mathcal{Q}_{\mathcal{E}} \times \mathcal{Q}_{\simeq}$ defined as follows.

— $(\bar{a}, a') \in \mathcal{R}$ iff a' is the \simeq -equivalence class $R_{\text{old}}(\bar{a}) \setminus R(\bar{a})$.

The refinement \simeq is obtained from the equivalence \equiv as follows.
 If \bar{a} is the invalidating abstract state, $\mu \in \delta_{\mathcal{E}}(\bar{a})$ the invalidating transition
 and (R_{old}, R) the invalidating witness then:
 The \equiv -equivalence class $\text{rel}_{\gamma}(a)$ is broken into
 new \simeq -equivalence classes $R_{old}(\bar{a}) \setminus R(\bar{a})$ and $\text{rel}_{\gamma}(a) \setminus (R_{old}(\bar{a}) \setminus R(\bar{a}))$
 For each $\bar{b} \in \text{post}(\mu, \bar{a}) \setminus \bar{a}$, the \equiv -equivalence class $\text{rel}_{\gamma}(b)$ is broken into
 new \simeq -equivalence classes $R_{old}(\bar{b})$ and $\text{rel}_{\gamma}(b) \setminus R_{old}(\bar{b})$
 No other \equiv -equivalence class is refined.

Fig. 14. Refinement algorithm based on invalid counterexamples

- For each $\bar{b} \in \text{post}(\mu, \bar{a}) \setminus \bar{a}$, $(\bar{b}, b') \in \mathcal{R}$ iff b' is the \simeq -equivalence class $R_{old}(\bar{b})$.
- For each $\bar{c} \in \mathcal{Q}_{\mathcal{E}} \setminus (\text{post}(\mu, \bar{a}) \cup \bar{a})$, $(\bar{c}, c') \in \mathcal{R}$ iff $c' = c$.

Please note that it is easy to see that by construction $(q_{\mathcal{E}}, q_{\simeq}) \in \mathcal{R}$; if \bar{a} is not the initial state then the observation follows immediately, and otherwise, observe that $q_{\mathcal{I}} \in R_{old}(\bar{a}) \setminus R(\bar{a})$. Further, we clearly have $\text{rel}_{\simeq, \equiv}^{\alpha} \circ \mathcal{R} = \text{rel}_{\text{inj}}$. We will show that \mathcal{R} is not a canonical simulation and hence we can conclude that \simeq is a good refinement.

Now, let $a_0 \in \mathcal{Q}_{\simeq}$ be the \simeq -equivalence class $\mathcal{R}_{old}(\bar{a}) \setminus R(\bar{a})$. Observe that $(\bar{a}, a_0) \in \mathcal{R}$. Next, recall that the violating transition $\mu \in \delta_{\mathcal{E}}(\bar{a})$. Hence the desired result will follow if we can show that for each $\mu_0 \in \delta_{\simeq}(a_0)$ we have that $\mu \not\preceq_{\mathcal{R}} \mu_0$.

We proceed by contradiction. Let μ' be such that $\mu' \in \delta_{\simeq}(a_0)$ and $\mu \preceq_{\mathcal{R}} \mu'$. By definition of abstractions, there is a $q \in \text{rel}_{\simeq}^{\gamma}(a_0) = R_{old}(\bar{a}) \setminus R(\bar{a})$ and $\mu_1 \in \delta(q)$ such that $\mu' = [\mu_1]_{\simeq}$. From $\mu \preceq_{\mathcal{R}} \mu'$, we can conclude the following.

- $\mu(\bar{a}) \leq \mu_1(R_{old}(\bar{a}) \setminus R(\bar{a})) \leq \mu_1(R_{old}(a))$.
- For each $\bar{b} \in \text{post}(\mu, \bar{a}) \setminus \bar{a}$, $\mu(\bar{b}) \leq \mu_1(R_{old}(\bar{b}))$.
- For $\bar{c} \in \mathcal{Q}_{\mathcal{E}} \setminus (\text{post}(\mu, \bar{a}) \cup \bar{a})$, $\mu(\bar{c}) = 0 \leq \mu_1(R_{old}(\bar{c}))$.

For each $\bar{d} \in \mathcal{Q}_{\mathcal{E}}$, it follows from construction that $R_{old}(\bar{d}) \subseteq \text{rel}_{\equiv}^{\gamma}(d)$. Therefore for all $\bar{d}_0, \bar{d}_1 \in \mathcal{Q}_{\mathcal{E}}$ such that $\bar{d}_0 \neq \bar{d}_1$, $R_{old}(\bar{d}_0) \cap R_{old}(\bar{d}_1) = \emptyset$. It now follows easily from the above observations that $\mu \preceq_{\mathcal{R}_{old}} \mu_1$ which contradicts Proposition 5.9. \square

Example 5.19. Consider the DTMC \mathcal{M}_{ex} and its abstraction \mathcal{M}_{ab} from Example 4.7 and the property $\psi_S = \mathcal{P}_{\leq \frac{2}{3}}(\mathbf{X}(\neg \mathcal{P}_{< \frac{1}{2}}(\mathbf{X}P)))$ from Example 5.11. \mathcal{M}_{ab} violates ψ_S and the counterexample generation algorithm will output the counterexample $(\mathcal{E}, \text{rel}_{\text{inj}})$ where \mathcal{E} is shown in Figure 12. The counterexample is invalid, and the invalidating abstract state is the initial state $\{\bar{q}_0\}$ and the invalidating transition μ is the transition out of $\{\bar{q}_0\}$. Note that $\text{post}(\{\bar{q}_0\}, \mu) = \{\{q_1, q_2\}, \{q_3, q_4\}\}$. Now $R_{old}(\{q_1, q_2\}) = q_2$ and $R_{old}(\{q_3, q_4\}) = q_3$. Thus the refinement procedure breaks two equivalence classes—the class $\{q_1, q_2\}$ is broken into $\{q_1\}$ and $\{q_2\}$ and the class $\{q_3, q_4\}$ is broken into $\{q_3\}$ and $\{q_4\}$ (strictly speaking the class $\{q_0\}$ is also broken $\{q_0\}$ and \emptyset). In contrast, the refinement procedure only breaks one class for non-probabilistic system.

We now exhibit an example in which all the steps of CEGAR are illustrated.

Example 5.20. Consider the DTMC $\mathcal{M} = (\mathbf{Q}, q_{\mathcal{I}}, \delta, \mathbf{L})$ where,

- \mathbf{Q} consists of $n + 3$ states (here n is a some even number ≥ 2). We let $\mathbf{Q} = \{q_f, q_{\text{even}}, q_{\text{odd}}\} \cup \{q_0, q_1 \dots, q_{n-1}\}$.
- $q_{\mathcal{I}} = q_0$.
- δ is defined as follows. For each $1 \leq i < n$, $\delta(q_i) = \{\mu_i\}$ where μ_i is the unique probability distribution on \mathbf{Q}_i that satisfies the following:
 - $\mu_i(q_f) = \frac{1}{4}$,
 - $\mu_i(q_{\text{odd}}) = \frac{1}{4}$ if i is odd,
 - $\mu_i(q_{\text{even}}) = \frac{1}{4}$ if i is 0 or even,
 - $\mu_i(q_{i+1}) = \frac{1}{2}$ if $i < n - 1$; and
 - $\mu_{n-1}(q_0) = \frac{1}{2}$.
- The probabilistic transition out of q_f assigns probability 0 to every state, the transition out of q_{odd} assigns probability $\frac{1}{2}$ to q_f and $\frac{1}{2}$ to q_{odd} ; and the transition out of q_{even} assigns probability 1 to q_{even} .
- The labeling function \mathbf{L} assigns proposition P to q_f . No other state is labeled by proposition P .

Consider the safety formula $\psi = \mathcal{P}_{\leq \frac{1}{2}}(\diamond P)$ and suppose we want to check whether $\mathcal{M} \models \psi$. We will start with a very coarse abstraction by dividing the set \mathbf{Q} into two equivalence classes $\{q_f\}$ and $\mathbf{Q} \setminus \{q_f\}$. The resulting abstraction \mathcal{M}_{ab} is shown in Figure 15 (the initial state $\mathbf{Q} \setminus \{q_f\}$ has three nondeterministic choices shown by long-dashed lines, solid lines and double lines). Now \mathcal{M}_{ab} does not satisfy ψ and there are two possible minimal counterexamples which witness this fact. These counterexamples $\mathcal{E}_{\text{ex}_0}$ and $\mathcal{E}_{\text{ex}_1}$ are shown in Figure 16 (the initial state is $\mathbf{Q} \setminus \{q_f\}$). For this example, we consider the case when the counterexample generated is $\mathcal{E}_{\text{ex}_0}$.

Suppose that the counterexample generating algorithm generates the counterexample $\mathcal{E}_{\text{ex}_0}$. The counterexample turns out to be invalid, and the refinement algorithm will split the equivalence class $\mathbf{Q} \setminus \{q_f\}$ into two classes — $\{q_{\text{even}}, q_{\text{odd}}\}$ and $\mathbf{Q}_n = \{q_0, q_1 \dots, q_{n-1}\}$. The splitting of $\mathbf{Q} \setminus \{q_f\}$ results in the new abstract MDP $\mathcal{M}_{\text{ab}_1}$ given in Figure 17 (the initial state is \mathbf{Q}_n and the state $\{q_{\text{even}}, q_{\text{odd}}\}$ has two outgoing transitions).

Now, $\mathcal{M}_{\text{ab}_1}$ also does not satisfy ψ and there is one minimal counterexample $\mathcal{E}_{\text{ex}_2}$ witnessing this. This counterexample is shown in Figure 18. This counterexample turns out to be invalid and the refinement process now splits the two classes — $\{q_{\text{even}}, q_{\text{odd}}\}$ and $\mathbf{Q}_n = \{q_0, q_1 \dots, q_{n-1}\}$. The class $\{q_{\text{even}}, q_{\text{odd}}\}$ is split into $\{q_{\text{even}}\}$ and $\{q_{\text{odd}}\}$. The class \mathbf{Q}_n is split into $\mathbf{Q}_{\text{even}} = \{q_0, q_2, \dots, q_{n-2}\}$ and $\mathbf{Q}_{\text{odd}} = \{q_1, q_3, \dots, q_{n-1}\}$. The resulting abstract MDP $\mathcal{M}_{\text{ab}_2}$ is shown in Figure 19 (the initial state is \mathbf{Q}_{even}).

Now, $\mathcal{M}_{\text{ab}_2}$ also does not satisfy ψ , and there is exactly one minimal counterexample $\mathcal{E}_{\text{ex}_3}$ which witnesses this fact. This counterexample is shown in Figure 20 (the initial state is \mathbf{Q}_{even}). This turns out to be a valid counterexample and at this point our CEGAR loop terminates with the result — $\mathcal{M} \not\models \psi$.

Remark 5.21. In the Electronic Appendix, we give another example demonstrating our CEGAR approach and show that model checking for safety properties using the CEGAR loop can be faster than model checking using direct methods.




Fig. 15. MDP \mathcal{M}_{ab} .




Fig. 16. \mathcal{E}_{ex_0} (left) and \mathcal{E}_{ex_1} (right).




Fig. 17. \mathcal{M}_{ab_1} .




Fig. 18. \mathcal{E}_{ex_2} .




Fig. 19. \mathcal{M}_{ab_2} .




Fig. 20. \mathcal{M}_{ex_3} .

5.3 Counterexample checking for weak safety

In this section we outline an algorithm that given a counterexample \mathcal{E} for a MDP \mathcal{M} and weak safety property ψ_{WS} either determines that \mathcal{E} is not a valid counterexample, or finds a finite unrolling of \mathcal{E} that is simulated by \mathcal{M} , and witnesses the fact that \mathcal{M} does not satisfy ψ_{WS} . The algorithm unrolls \mathcal{E} on the fly, and does not construct the unrolled MDP explicitly. The running time of the algorithm depends on the height of the unrolling, which if small, can result in faster checking than the algorithm shown in Figure 11. Before presenting the algorithm, we introduce some notation that we will find useful in describing the algorithm. Recall that given a

MDP \mathcal{M} , a state q of \mathcal{M} and $k \in \mathbb{N}$ the k -th unrolling of \mathcal{M} rooted at q is denoted by \mathcal{M}_k^q .

Notation 5.22. Given $k \in \mathbb{N}$ and states $q, q' \in \mathcal{M}$, we say that $q \preceq_k q'$ if $\mathcal{M}_k^q \preceq \mathcal{M}_k^{q'}$. Given a PCTL formula ψ we say that $q \Vdash_k \psi$ if $\mathcal{M}_k^q \Vdash \psi$.

We observe the following two facts. If $q \preceq q'$ then $q \preceq_k q'$ for all k . The proof of Theorem 3.4 implies that for a weak safety formula ψ_{WS} if $q \not\Vdash \psi_{WS}$, then there is a k_0 s.t. $q \not\Vdash_{k_0} \psi_{WS}$. These two facts can be combined to obtain a counterexample checking algorithm for the weak safety fragment of PCTL as we will describe shortly.

For the rest of the Section, we fix the following notation. $\mathcal{M} = (\mathbf{Q}, q_{\mathcal{I}}, \delta, \mathbf{L})$ is the (original) MDP that we are checking against weak safety property ψ_{WS} and \equiv is an equivalence relation that is compatible with \mathcal{M} . Assuming \mathcal{M}_{\equiv} violates the safety property ψ_{WS} , we will denote the minimal counterexample obtained as in Theorem 3.16 by $(\mathcal{E}, \text{rel}_{\text{inj}})$. Let $\mathcal{E} = (\mathbf{Q}_{\mathcal{E}}, q_{\mathcal{E}}, \delta_{\mathcal{E}}, \mathbf{L}_{\mathcal{E}})$. For a state $a = [q]_{\equiv}$ in abstraction \mathcal{M}_{\equiv} , $\text{rel}_{\equiv}^{\downarrow}(a) = \{q' \in \mathbf{Q} \mid q \equiv q'\}$ is the concretization map. The relation $\{(\bar{a}, q) \mid \bar{a} \in \mathbf{Q}_{\mathcal{E}} \text{ and } q \in \text{rel}_{\equiv}^{\downarrow}(a)\}$ will be denoted by $R_{\mathcal{I}}$. Finally, we will use ψ_{SL} to denote the strict liveness formula obtained by negating ψ_{WS} . $\text{SLSubForm}(\psi_{SL})$ will denote the set of PCTL strict liveness subformulas of ψ_{SL} and $\text{PathForm}(\psi_{SL})$ will denote the set of path subformulas of ψ_{SL} .⁵

The proposed algorithm iteratively constructs the relations $R_k = R_{\mathcal{I}} \cap \preceq_k$ and $\text{Sat}_k = \{(\bar{a}, \psi) \mid \bar{a} \Vdash_k \psi, \bar{a} \in \mathbf{Q}_{\mathcal{E}}, \psi \in \text{SLSubForm}(\psi_{SL})\}$. We make the following observations.

- (1) If the set $R_k(\bar{a}) = \{q \mid (\bar{a}, q) \in R_k\}$ becomes \emptyset for some k and $\bar{a} \in \mathbf{Q}_{\mathcal{E}}$, then $(\mathcal{E}, \text{rel}_{\text{inj}})$ is invalid for (\mathcal{M}, \equiv) . We can also call the counterexample invalid if the initial concrete state $q_{\mathcal{I}}$ is not contained in $R_k(q_{\mathcal{E}})$.
- (2) If $\text{Sat}_k(q_{\mathcal{E}}, \psi_{SL})$ and $q_{\mathcal{I}} \in R_k(q_{\mathcal{E}})$ then $q_{\mathcal{I}} \Vdash_k \psi_{SL}$ also. Thus, the concrete MDP violates the given safety property and we can report this.

This iteration must end as a consequence of Theorem 3.4. The computation also needs to compute the function MaxProb_k defined as follows. Given $\bar{a} \in \mathbf{Q}_{\mathcal{E}}$ and a path formula $\phi \in \text{PathForm}(\psi_{SL})$, MaxProb_k gives the maximum probability (over all schedulers) of ϕ being true in $\mathcal{E}_{\bar{a}}$. The relations R_{k+1} , Sat_{k+1} and MaxProb_{k+1} can be computed using R_k , Sat_k , and MaxProb_k and do not need other previous values.

Figure 21 gives the details of this algorithm. At the beginning of $(k+1)$ -th unrolling of the while loop, the relations R_{curr} and Sat_{curr} are the relations R_k and Sat_k respectively. The (doubly-indexed) array $\text{MaxProb}_{\text{curr}}[\bar{a}][\psi]$ is the function $\text{MaxProb}_k[\bar{a}][\psi]$. Within the while loop, $R_{\text{tmp}, \bar{a}}$ is the set $R_{k+1}(\bar{a})$ while $\text{Sat}_{\text{tmp}, \bar{a}}$ is the set $\{\psi \mid \bar{a} \Vdash_{k+1} \psi \in \text{SLSubForm}(\psi_{SL})\}$, and $\text{MaxProb}_{\text{tmp}, \bar{a}}[\psi]$ is $\text{MaxProb}_{k+1}[\bar{a}][\psi]$. R_{curr} , Sat_{curr} , and $\text{MaxProb}_{\text{curr}}$ are updated *after* $R_{\text{tmp}, \bar{a}}$, $\text{Sat}_{\text{tmp}, \bar{a}}$ and $\text{MaxProb}_{\text{tmp}, \bar{a}}$ are computed for *all* states $\bar{a} \in \mathbf{Q}_{\mathcal{E}}$. The following proposition follows easily from the observations made in the section.

PROPOSITION 5.23. *The algorithm in Figure 21 terminates. If the algorithm returns “Invalid Counterexample” then the counterexample obtained using Theorem 3.16 is not valid. If the algorithm returns “Safety Violated” then $\mathcal{M} \not\Vdash \psi_{WS}$.*

⁵A path subformula of a PCTL-liveness formula ψ_L are formulas of the kind $X\psi_L$ and $\psi_L U \psi_L$.

```

Initially
   $R_{curr} = \{(\bar{a}, q) \mid q \in \text{rel}_\gamma(\bar{a}), \bar{a} \in \mathcal{Q}_\mathcal{E}\}$ 
   $\text{Sat}_{curr} = \{(\bar{a}, \psi) \mid \bar{a} \Vdash_0 \psi, \psi \in \text{SLSubForm}(\psi_{SL}), \bar{a} \in \mathcal{Q}_\mathcal{E}\}$ 
   $\text{MaxProb}_{curr}[\bar{a}][\phi] = \text{MaxProb}_0(\bar{a})(\phi)$  for  $\bar{a} \in \mathcal{Q}_\mathcal{E}, \phi \in \text{PathForm}(\psi_{SL})$ 
While (true)
  do
    If  $q_I \notin R_{curr}(q_\mathcal{E})$  return “Invalid Counterexample”
    If  $\text{Sat}_{curr}(q_\mathcal{E}, \psi_{SL})$  return “Safety Violated”
    for each  $\bar{a} \in \mathcal{Q}_\mathcal{E}$ 
      do
         $R_{tmp, \bar{a}} = \{q \mid (\bar{a}, q) \in R_{curr} \text{ and } \forall \mu \in \delta_\mathcal{E}(\bar{a}) \exists \mu' \in \delta(q). \mu \preceq_{R_{curr}} \mu'\}$ 
        If  $R_{tmp, \bar{a}} = \emptyset$  return “Invalid Counterexample”
        COMPUTE( $\bar{a}, \text{Sat}_{curr}, \text{MaxProb}_{curr}, \text{Sat}_{tmp, \bar{a}}, \text{MaxProb}_{tmp, \bar{a}}$ )
      od
     $\text{Sat}_{curr} = \{(\bar{a}, \psi) \mid \bar{a} \in \mathcal{Q}_\mathcal{E}, \psi \in \text{Sat}_{tmp, \bar{a}}\}$ 
     $R_{curr} = \{(\bar{a}, q) \mid \bar{a} \in \mathcal{Q}_\mathcal{E}, q \in R_{tmp, \bar{a}}\}$ 
     $\text{MaxProb}_{curr}[\bar{a}][\phi] = \text{MaxProb}_{tmp, \bar{a}}[\phi]$  for each  $a \in \mathcal{Q}_\mathcal{E}, \phi \in \text{PathForm}(\psi)$ 
  od

```

The procedure COMPUTE returns $\text{Sat}_{tmp, \bar{a}}$, the set of sub-formulas of ψ_{SL} satisfied by \bar{a} in the “next” unrolling of the tree with root \bar{a} . It also returns $\text{MaxProb}_{tmp, \bar{a}}$, that given a path formula ϕ gives the maximum probability of ϕ being true in the next unrolling. COMPUTE is defined as follows.

```

COMPUTE( $\bar{a}, \text{Sat}_{curr}, \text{MaxProb}_{curr}, \text{Sat}_{tmp, \bar{a}}, \text{MaxProb}_{tmp, \bar{a}}$ )
  Fix an enumeration  $\varphi_1, \dots, \varphi_n$  of the set  $\{\varphi \mid \varphi \in \text{PathForm}(\psi_{SL}) \cup \text{SLSubForm}(\psi_{SL})\}$  such that  $\text{size}(\varphi_i) \leq \text{size}(\varphi_j)$  for  $i \leq j$ .
  Initially
     $\text{Sat}_{tmp, \bar{a}} = \emptyset$ 
     $\text{MaxProb}_{tmp, \bar{a}}[\phi] = 0$  for all  $\phi \in \text{PathForm}(\psi_{SL})$ .
  For  $i = 1$  to  $n$ 
    If  $\varphi_i$  is  $p$  and  $p \in \mathcal{L}_\mathcal{E}(\bar{a})$  then  $\text{Sat}_{tmp, \bar{a}} = \text{Sat}_{tmp, \bar{a}} \cup \{p\}$ 
    If  $\varphi_i$  is  $\neg p$  and  $p \notin \mathcal{L}_\mathcal{E}(\bar{a})$  then  $\text{Sat}_{tmp, \bar{a}} = \text{Sat}_{tmp, \bar{a}} \cup \{\neg p\}$ 
    If  $\varphi_i$  is  $\psi_1 \vee \psi_2$  and ( $\psi_1 \in \text{Sat}_{tmp, \bar{a}}$  or  $\psi_2 \in \text{Sat}_{tmp, \bar{a}}$ ) then
       $\text{Sat}_{tmp, \bar{a}} = \text{Sat}_{tmp, \bar{a}} \cup \{\varphi_i\}$ 
    If  $\varphi_i$  is  $\psi_1 \wedge \psi_2$ ,  $\psi_1 \in \text{Sat}_{tmp, \bar{a}}$  and  $\psi_2 \in \text{Sat}_{tmp, \bar{a}}$  then
       $\text{Sat}_{tmp, \bar{a}} = \text{Sat}_{tmp, \bar{a}} \cup \{\varphi_i\}$ 
    If  $\varphi_i$  is  $\neg \mathcal{P}_{\leq p}(\phi)$  and  $\text{MaxProb}_{tmp, \bar{a}}[\phi] > p$ 
      then  $\text{Sat}_{tmp, \bar{a}} = \text{Sat}_{tmp, \bar{a}} \cup \{\varphi_i\}$ 

    If  $\varphi_i$  is  $X\psi$  then
       $\text{MaxProb}_{tmp, \bar{a}}[\varphi_i] = \max_{\mu \in \delta_\mathcal{E}(\bar{a})} \mu(\{\bar{b} \mid (\bar{b}, \psi) \in \text{Sat}_{curr}\})$ 
    If  $\varphi_i$  is  $\psi_1 \mathcal{U} \psi_2$  then
      If  $\psi_2 \in \text{Sat}_{tmp, \bar{a}}$  then  $\text{MaxProb}_{tmp, \bar{a}}[\varphi_i] = 1$ 
      else
        If  $\psi_1 \notin \text{Sat}_{tmp, \bar{a}}$  then  $\text{MaxProb}_{tmp, \bar{a}}[\varphi_i] = 0$ 
        else  $\text{MaxProb}_{tmp, \bar{a}}[\varphi_i] = \max_{\mu \in \delta_\mathcal{E}(\bar{a})} \sum_{\bar{b} \in \mathcal{Q}_\mathcal{E}} (\mu(\bar{b})) (\text{MaxProb}_{curr}[\psi_1 \mathcal{U} \psi_2][\bar{b}])$ 

```

Fig. 21. On the fly algorithm for checking Strict Liveness

Finally, we observe that the algorithm in Figure 21 may be made more efficient in practice as follows. First, since we are dealing with strict liveness fragment, the sequence Sat_k is an increasing sequence and the function $\text{MaxProb}_k[\bar{a}][\psi] \leq \text{MaxProb}_{k+1}[\bar{a}][\psi]$. Hence, only needs to compute $\text{Sat}_{k+1} \setminus \text{Sat}_k$ and $\text{MaxProb}_{k+1} - \text{MaxProb}_k$. This optimization will be explored in future work.

6. RELATED WORK

Abstraction Schemes: Abstractions have been extensively studied in the context of probabilistic systems. General issues in defining good abstractions as well as specific proposals for families of abstract models are presented in [Jonsson and Larsen 1991; Huth 2004; Norman 2004; Huth 2005; D’Argenio et al. 2001; 2002; Fecher et al. 2006; Katoen et al. 2007; Monniaux 2005; Kwiatkowska et al. 2006; McIver and Morgan 2004]. MDPs as abstract models of probabilistic systems have been considered in a number of papers. They are natural models to abstract to because when one abstracts (whether in the non-probabilistic case or the probabilistic case) the main feature that abstract models have is nondeterminism. In the context of probabilistic systems this means considering transition systems that have both nondeterminism and probability, which is exactly the model of MDPs. Instead of explicit nondeterministic choices, Markov Chains with interval bounds on transition probabilities have been proposed as abstract models in [Fecher et al. 2006]. Such “interval” Markov Chains are abstractions of the models we consider in this paper (and hence less precise), and are useful when the number of nondeterministic choices out of an abstract state are more than an exponential in the number of abstract states; in such cases model checking the interval Markov Chain can be faster, even if it yields less precise results. Another important scheme for abstract models proposed is where the system is abstracted as a game between two players. Like our construction of abstract MDPs, the states of the game correspond to equivalence classes of some relation on the concrete states. In the ensuing game, one player picks a concrete state in the equivalence class associated with the abstract state, and the second player picks one of the transitions out of the concrete state chosen by player 1. Such a game abstraction has an advantage over the MDP abstraction in that both upper and lower bounds on the probability of reaching states can be obtained by analyzing the game; the MDP abstraction will only allow one to obtain upper bounds. However, the game abstraction approach has so far only been applied to simple quantitative properties, and it is not clear how to extend it to a general logic fragment like PCTL-safety.

Recently, theorem-prover based algorithms for constructing abstractions of probabilistic systems based on predicates have been presented [Wachter et al. 2007]. Another notion that has been recently proposed is the notion of a “magnifying-lens abstraction” [de Alfaro and Roy 2007], which can be used to assist in the model checking process, by approximating the measure of the satisfaction of path formulas for sets of concrete states; the method is not an abstraction in the traditional sense in that neither is an abstract model explicitly constructed, nor is the model used for reasoning, one that simulates the concrete model.

Counterexamples: The notion of counterexamples is critical for the approach of counterexample guided abstraction-refinement. Criteria for defining counterexam-

ples are identified in [Clarke et al. 2002], along with a notion of counterexamples for branching-time properties and non-probabilistic systems. The problem of defining counterexamples for probabilistic systems has received considerable attention recently. Starting from the seminal papers [Aljazzar et al. 2005; Han and Katoen 2007a], the notion of sets of paths with high measure as counterexamples has been used for DTMCs, CTMCs, and MDPs [Han and Katoen 2007a; 2007b; Aljazzar and Leue 2007]. Another definition that has been proposed is that of DTMCs (or purely probabilistic models) in [Chatterjee et al. 2005; Hermanns et al. 2008]. Our notion of counterexample is different from these proposals and we demonstrate that these other proposals are not rich enough for the class of properties we consider.

Automatic Abstraction-Refinement: The foundations of automatic abstraction-refinement in the context of non-probabilistic systems were laid down in [Clarke et al. 2000; Clarke et al. 2002], and the technique has been applied to analyze software [Ball and Rajamani 2002; Majumdar et al. 2002; Kroening et al. 2004; Cook et al. 2005] and hardware designs [Clarke et al. 2002; Glussman et al. 2003; Bjesse and Kukula 2004] with success. CEGAR has also been extended to consider 3-valued abstractions [Gurfinkel and Chechik 2006; Shoham and Grumberg 2007]. A number of subtle heuristics have been identified that assist in the construction of abstractions, counterexample analysis, and refinement. Primarily these are identifying special equivalences on concrete states based on predicates [Das 2003; Lakhnech et al. 2001] and automata [Bouajjani et al. 2004] to construct abstractions efficiently; using SAT solvers to analyze counterexamples [Das 2003; Clarke et al. 2002; Glussman et al. 2003; Bjesse and Kukula 2004]; using SAT solvers, ILP, and learning [Clarke et al. 2002] to identify heuristics to split abstract states during refinement. In order for the CEGAR approach to scale to large systems, these heuristics, identified for non-probabilistic systems, will need to be borrowed and experimented with in the future. This paper focuses primarily on laying down the definitions and algorithms needed to serve as the foundation for the CEGAR approach for probabilistic systems.

In the context of probabilistic systems, automatic abstraction-refinement was first considered in [D’Argenio et al. 2001; 2002]. There are two main differences with our work. First, they consider only reachability properties. Second, the refinement process outlined in [D’Argenio et al. 2001; 2002] is not counterexample based, but rather based on partition refinement. What this means is that their refinement is biased towards separating states that are not bisimilar, rather than states that are “distinguished” by the property, and so it is likely that their method refines more than needed.

Counterexample guided refinement has been used as the basis of synthesizing winning strategies for 2-player stochastic games in [Chatterjee et al. 2005]. Though the problem of 2-player games is more general than verification, the specific model considered by Chatterjee et al. [2005] has some peculiarities and so does not subsume the problem or its solution presented here. First, in their model, states are partitioned into “nondeterministic” states that have purely nondeterministic transitions and “probabilistic” states that have purely probabilistic transitions. The abstraction does not abstract any of the “probabilistic states”; only the “nondeterministic” states are collapsed. This results in larger abstract models, and obviates

certain issues in counterexample checking that we deal with. Second, they take counterexamples to be finite models without nondeterminism. They can do this because they consider a simpler class of properties than we do, and as we show in Section 3.3, DTMCs are not rich enough for all of safe-PCTL. Next, their counterexample checking algorithm is different than even the one used in the context of non-probabilistic systems. They consider a counterexample to be valid only if *all* the concrete states corresponding to the abstract states in the counterexample can simulate the behavior captured by the counterexample. Thus, they deem certain counterexamples to be spurious, even if they will be recognized as providing enough evidence for the violation of the property by other CEGAR schemes (including ours). Finally, they do not have a precise statement characterizing the qualities of their refinement algorithm.

In a recent paper, Hermanns et al. [2008] consider CEGAR for probabilistic systems. They consider very special types of reachability properties namely, those that can be expressed by formulas of the form $\mathcal{P}_{\leq p}(\psi_1 \mathcal{U} \psi_2)$ where ψ_1 and ψ_2 are propositions (or boolean combinations of propositions). For this class of properties, they use DTMCs as the notion of counterexamples: the counterexample obtained is a pair $(\mathcal{S}, \mathcal{M}^{\mathcal{S}})$ where \mathcal{S} is a memoryless scheduler. As we show, DTMCs cannot serve as counterexample for the richer class of properties considered here. For the counterexample checking algorithm, they generate a finite set, \mathbf{ap} , of *abstract* paths of $\mathcal{M}_{\equiv}^{\mathcal{S}}$ in decreasing order of measures such that the total measure of these paths is $> p$. Then, they build (on-the-fly) a “concrete” scheduler which maximizes the measure of the paths in \mathbf{ap} that are simulated by the original MDP. Let $\mathbf{p}_{\text{total}}$ be the maximum probability (under all schedulers) of $\psi_1 \mathcal{U} \psi_2$ being satisfied by the abstract MDP, $\mathbf{p}_{\mathbf{ap}}$ be the total measure of \mathbf{ap} and \mathbf{p}_{max} be the maximum measure of “abstract” paths in \mathbf{ap} simulated by the concrete MDP. If $\mathbf{p}_{\text{max}} > p$ then the counterexample is declared to be valid and if $\mathbf{p}_{\text{total}} - \mathbf{p}_{\mathbf{ap}} + \mathbf{p}_{\text{max}} \leq p$ then the counterexample is declared to be invalid and the abstraction refined. If neither is the case, then Hermanns et al. [2008] heuristically decide either to generate more abstract paths or to refine the abstraction. The refinement is based upon refining some “spurious” abstract path (namely, a path that is not simulated by the concrete system). There is, however, no formal statement characterizing progress based on the refinement algorithm outlined in [Hermanns et al. 2008].

Another automatic abstraction-refinement approach was recently proposed in [Kattenbelt et al. 2009], in which the abstract MDP is represented as a two-person stochastic game [Kwiatkowska et al. 2006]—one player picks the nondeterminism of the concrete state and the other player picks the nondeterminism introduced by the abstraction. They do not verify the game against a formula in a logic, instead they compute both upper and lower bounds on quantitative properties of the original MDP upto a predefined precision. The quantitative properties considered are the maximum and minimum probabilities of reaching a state (they can also deal with expected rewards). Model checking the two-person stochastic game not only yields upper and lower bounds on quantitative properties of the original MDP, but also the strategies for both the players that achieve the bounds. If the difference between the bounds is less than the predefined precision, the abstraction-refinement loop terminates otherwise they use the strategies to further refine the abstraction. The

refinement is achieved by identifying an abstract state for which the strategy for achieving the upper and lower bounds is witnessed by distinct concrete states and refinement step separates these concrete states.

7. CONCLUSIONS AND FUTURE WORK

We presented a CEGAR framework for MDPs, where a MDP \mathcal{M} is abstracted by another MDP \mathcal{A} defined using an equivalence on the states of \mathcal{M} . Our main contributions when presenting this framework were a definition for the notion of a counterexample, along with algorithms to compute counterexamples, check their validity and perform automatic refinement based on an invalid counterexample. Our discussion in this paper considered only the analysis of finite state MDPs. A natural question to ask is how these ideas would translate when analyzing infinite state MDPs. As in the case of non-probabilistic systems, the notions of counterexamples, the validity of counterexamples, principles of refinement etc. will remain the same when moving to infinite state systems. However, just like for non-probabilistic systems, for even one iteration of the CEGAR loop to terminate for infinite state systems, the basic problems of constructing abstractions, and analyzing counterexamples must be decidable. These problems are in general undecidable for infinite state systems, and therefore CEGAR can only be applied to certain special tractable classes of infinite state MDPs.

There are a number of interesting questions left open for future investigation. First these ideas need to be implemented and experimented with. In order for this approach to be scalable, symbolic algorithms for the steps outlined here, will be required. In this paper, we considered abstractions defined using a general equivalence relations on concrete states. For non-probabilistic systems, it has been observed that the most useful type of equivalences are those defined using predicates [Das 2003] on concrete states. While the construction of abstractions for probabilistic systems using predicates has been explored [Hermanns et al. 2008], it will be important to see how counterexamples can be analyzed and predicates refined using SAT solvers and theorem provers. Further a number of key heuristics have been identified analyzing non-probabilistic systems using SAT solvers, and it will be interesting to see how those heuristics perform in the context of probabilistic systems.

Next, ideas and heuristics to improve the running time of our counterexample construction algorithm need to be explored. First, our algorithm makes multiple calls to a model checker which can be expensive, and it will be important to explore to leverage the results of previous model checking runs in subsequent model checking runs. Second, the order in which transitions are considered for elimination, crucially affects the final size of the counterexample. Good heuristics for ordering transitions to obtain small counterexamples, must be identified.

ELECTRONIC APPENDIX

The electronic appendix for this article can be accessed in the ACM Digital Library by visiting the following URL: <http://www.acm.org/pubs/citations/journals/tocl/20YY-V-N/p1-Chadha>.

ACKNOWLEDGMENTS

The authors thank the following people: Chandra Chekuri for many fruitful discussions especially relating to the algorithmic aspects of finding counterexamples and checking validity of counterexamples; Radha Jagadeesan for discussions on simulations and the safety fragment of PCTL; anonymous referees for sending pointers to [Aljazzar and Leue 2007; Hermanns et al. 2008; Kattenbelt et al. 2009], for encouraging us to formally articulate the guarantees of our refinement algorithm, and improving the general presentation of the results.

REFERENCES

- ALJAZZAR, H., HERMANN, H., AND LEUE, S. 2005. Counterexamples for timed probabilistic reachability. In *Proceedings of International Conference on Formal Modelling and Analysis of Timed Systems*. 177–195.
- ALJAZZAR, H. AND LEUE, S. 2007. Counterexamples for model checking of Markov Decision Processes. Tech. Rep. soft-08-01, University of Konstanz.
- BAIER, C., ENGELEN, B., AND MAJSTER-CEDERBAUM, M. E. 2000. Deciding bisimilarity and similarity for probabilistic processes. *Journal of Computer and System Sciences* 60, 1, 187–231.
- BAIER, C., KATOEN, J.-P., HERMANN, H., AND WOLF, V. 2005. Comparative branching-time semantics for Markov chains. *Information and Computation* 200, 149–214.
- BALL, T. AND RAJAMANI, S. 2002. Generating abstract explanations of spurious counterexamples in C programs. Tech. Rep. 2002-09, Microsoft Research.
- BIANCO, A. AND DE ALFARO, L. 1995. Model checking of probabilistic and nondeterministic systems. In *Proceedings of the International Conference on the Foundations of Software Technology and Theoretical Computer Science*. 499–513.
- BJESSE, P. AND KUKULA, J. 2004. Using counter example guided abstraction refinement to find complex bugs. In *Proceedings of Design, Automation, and Test in Europe Conference and Exposition*. 156–161.
- BOUAJJANI, A., HABERMEHL, P., AND VOJNAR, T. 2004. Abstract Regular Model Checking. In *Proceedings of the International Conference on Computer Aided Verification*. 372–286.
- CHATTERJEE, K., HENZINGER, T., JHALA, R., AND MAJUMDAR, R. 2005. Counter-example Guided Planning. In *Proceedings of Uncertainty in Artificial Intelligence*. 104–111.
- CLARKE, E., GRUMBERG, O., JHA, S., LU, Y., AND VEITH, H. 2000. Counterexample-guided abstraction refinement. In *Proceedings of the International Conference on Computer Aided Verification*. 154–169.
- CLARKE, E., GUPTA, A., KUKULA, J., AND STRICHMAN, O. 2002. SAT based abstraction-refinement using ILP and machine learning techniques. In *Proceedings of the International Conference on Computer Aided Verification*. 265–279.
- CLARKE, E., JHA, S., LU, Y., AND VEITH, H. 2002. Tree-like counterexamples in model checking. In *Proceedings of the IEEE Symposium on Logic in Computer Science*. 19–29.
- COOK, B., PODELSKI, A., AND RYBALCHENKO, A. 2005. Abstraction Refinement for Termination. In *Proceedings of the International Static Analysis Symposium*. 87–101.
- D’ARGENIO, P. R., JEANNET, B., JENSEN, H. E., AND LARSEN, K. G. 2001. Reachability analysis of probabilistic systems by successive refinements. In *Proceedings of the International Workshop on Performance Modeling and Verification*. 39–56.
- D’ARGENIO, P. R., JEANNET, B., JENSEN, H. E., AND LARSEN, K. G. 2002. Reduction and refinement strategies for probabilistic analysis. In *Proceedings of the International Workshop on Performance Modeling and Verification*. 57–76.
- DAS, S. 2003. Predicate abstraction. Ph.D. thesis, Stanford University.
- DE ALFARO, L. AND ROY, P. 2007. Magnifying-lens abstraction for Markov Decision Processes. In *Proceedings of the International Conference on Computer Aided Verification*. 325–338.
- DESHARNAIS, J. 1999a. Labelled Markov processes. Ph.D. thesis.

- DESHARNAIS, J. 1999b. Logical characterization of simulation for Markov chains. Tech. Rep. CSR-99-8, University of Birmingham. Proceedings of the International Workshop on Performance Modeling and Verification.
- DESHARNAIS, J., GUPTA, V., JAGADEESAN, R., AND PANANGADEN, P. 2000. Approximating labeled Markov processes. In *Proceedings of the IEEE Symposium on Logic in Computer Science*. 95–106.
- DODIS, Y. AND KHANNA, S. 1999. Designing networks with bounded pairwise distance. In *Proceedings of the ACM Symposium on the Theory of Computing*. 750–759.
- FECHER, H., LEUCKER, M., AND WOLF, V. 2006. Don't know in probabilistic systems. In *Proceedings of the International SPIN Workshop on Model Checking Software*. 71–88.
- GAREY, M. AND JOHNSON, D. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman.
- GLUSSMAN, M., KAMHI, G., MADOR-HAIM, S., FRAER, R., AND VARDI, M. 2003. Multiple-Counterexample Guided Iterative Abstraction Refinement: An Industrial Evaluation. In *Proceedings of the International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. 176–191.
- GURFINKEL, A. AND CHECHIK, M. 2006. Why waste a perfectly good abstraction? In *Proceedings of the International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. 212–226.
- HAN, T. AND KATOEN, J.-P. 2007a. Counterexamples in probabilistic model checking. In *Proceedings of the International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. 72–86.
- HAN, T. AND KATOEN, J.-P. 2007b. Providing evidence of likely being on time — Counterexample generation for CTMC. In *Proceedings of the International Symposium on Automated Technology for Verification and Analysis*.
- HERMANN, H., WACHTER, B., AND ZHANG, L. 2008. Probabilistic CEGAR. In *Proceedings of the International Conference on Computer Aided Verification*. Also available as technical report number 33, AVACS, University of Saarland, 2007.
- HUTH, M. 2004. An abstraction framework for mixed non-deterministic and probabilistic systems. In *Validation of Stochastic Systems: A Guide to Current Research*. 419–444.
- HUTH, M. 2005. On finite-state approximants for probabilistic computation tree logic. *TCS* 346, 113–134.
- JONSSON, B. AND LARSEN, K. G. 1991. Specification and refinement of probabilistic processes. In *Proceedings of the IEEE Symposium on Logic in Computer Science*. 266–277.
- KATOEN, J.-P., KLINK, D., LEUCKER, M., AND WOLF, V. 2007. Three-valued abstraction for continuous-time Markov chains. In *Proceedings of the International Conference on Computer Aided Verification*. 311–324.
- KATTENBELT, M., KWIATKOWSKA, M., NORMAN, G., AND PARKER, D. 2009. Abstraction refinement for probabilistic software. In *Proc. 10th International Conference on Verification, Model Checking, and Abstract Interpretation (VMCAI'09)*, N. Jones and M. Müller-Olm, Eds. LNCS, vol. 5403. Springer, 182–197.
- KEMENY, J. AND SNELL, J. 1976. *Denumerable Markov Chains*. Springer-Verlag.
- KROENING, D., GROCE, A., AND CLARKE, E. 2004. Counterexample Guided Abstraction Refinement via Program Execution. In *Proceedings of the International Conference on Formal Engineering Methods*. 224–238.
- KWIATKOWSKA, M., NORMAN, G., AND PARKER, D. 2006. Game-based abstraction for Markov Decision Processes. In *Proceedings of the International Conference on Quantitative Evaluation of Systems*. 157–166.
- LAKHNECH, Y., BENSALAM, S., BEREZIN, S., AND OWRE, S. 2001. Incremental verification by abstraction. In *Proceedings of the International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. 98–112.
- MAJUMDAR, R., HENZINGER, T., JHALA, R., AND SUTRE, G. 2002. Lazy abstraction. In *Proceedings of the ACM Symposium on Principles of Programming Languages*. 58–70.
- ACM Transactions on Computational Logic, Vol. V, No. N, Month 20YY.

- MCIVER, A. AND MORGAN, C. 2004. *Abstraction, Refinement and Proof for Probabilistic Systems*. Springer.
- MONNIAUX, D. 2005. Abstract interpretation of programs as Markov Decision Processes. *Science of Computer Programming* 58, 179–205.
- NORMAN, G. 2004. Analyzing randomized distributed algorithms. In *Validation of Stochastic Systems: A Guide to Current Research*. 384–418.
- PAPOULIS, A. AND PILLAI, S. U. 2002. *Probability, Random Variables and Stochastic Processes*. McGraw Hill, New York.
- RUTTEN, J. M., KWIATKOWSKA, M., NORMAN, G., AND PARKER, D. 2004. *Mathematical Techniques for Analyzing Concurrent and Probabilistic Systems*. AMS.
- SEGALA, R. 2006. Probability and nondeterminism in operational models of concurrency. In *Proceedings of the International Conference on Concurrency Theory*. 64–78.
- SEGALA, R. AND LYNCH, N. A. 1994. Probabilistic simulations for probabilistic processes. In *Proceedings of the International Conference on Concurrency Theory*. 481–496.
- SEGALA, R. AND LYNCH, N. A. 1995. Probabilistic simulations for probabilistic processes. *Nord. J. Comput.* 2, 2, 250–273.
- SHOHAM, S. AND GRUMBERG, O. 2007. A Game-Based framework for CTL Counterexamples and 3-Valued Abstraction-Refinement. *ACM Transactions on Computational Logic* 9, 1.
- WACHTER, B., ZHANG, L., AND HERMANN, H. 2007. Probabilistic model checking modulo theories. In *Proceedings of the International Conference on Quantitative Evaluation of Systems*.
- ZHANG, L., HERMANN, H., EISENBRAND, F., AND JANSEN, D. N. 2007. Flow faster: Efficient decision algorithms for probabilistic simulation. In *Proceedings of the International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. 155–170.

Received July 2008; revised March 2009; accepted July 2009