

Probabilistic Hyperproperties with Nondeterminism^{*}

Erika Ábrahám¹, Ezio Bartocci², Borzoo Bonakdarpour³, and Oyendrila Dobe³

¹ RWTH-Aachen, Germany, abraham@informatik.rwth-aachen.de

² Technische Universität Wien, Austria, ezio.bartocci@tuwien.ac.at

³ Michigan State University, USA, {borzoo,dobeoyen}@msu.edu

Abstract. We study the problem of formalizing and checking probabilistic hyperproperties for models that allow nondeterminism in actions. We extend the temporal logic HyperPCTL, which has been previously introduced for discrete-time Markov chains, to enable the specification of hyperproperties also for Markov decision processes. We generalize HyperPCTL by allowing explicit and simultaneous quantification over schedulers and probabilistic computation trees and show that it can express important quantitative requirements in security and privacy. We show that HyperPCTL model checking over MDPs is in general undecidable for quantification over probabilistic schedulers with memory, but restricting the domain to memoryless non-probabilistic schedulers turns the model checking problem decidable. Subsequently, we propose an SMT-based encoding for model checking this language and evaluate its performance.

1 Introduction

Hyperproperties [1] extend the conventional notion of *trace properties* [2] from a set of traces to a set of sets of traces. In other words, a hyperproperty stipulates a *system* property and not the property of just individual traces. It has been shown that many interesting requirements in computing systems are hyperproperties and cannot be expressed by trace properties. Examples include (1) a wide range of information-flow security policies such as *noninterference* [3] and *observational determinism* [4], (2) sensitivity and robustness requirements in cyber-physical systems [5], and consistency conditions such as *linearizability* in concurrent data structures [6].

Hyperproperties can describe the requirements of probabilistic systems as well. They generally express probabilistic relations between multiple executions of a system. For example, in information-flow security, adding probabilities is motivated by establishing a connection between information theory and information flow across multiple traces. A prominent example is probabilistic schedulers that

^{*} This research has been partially supported by the United States NSF SaTC Award 181338, by the Vienna Science and Technology Fund ProbInG Grant ICT19-018 and by the DFG Research and Training Group UnRAVeL. The order of authors is alphabetical and all authors made equal contribution.

open up an opportunity for an attacker to set up a probabilistic covert channel. Or, *probabilistic causation* compares the probability of occurrence of an effect between scenarios where the cause is or is not present.

The state of the art on probabilistic hyperproperties has exclusively been studied in the context of discrete-time Markov chains (DTMCs). In [7], we proposed the temporal logic HyperPCTL, which extends PCTL by allowing explicit and simultaneous quantification over computation trees. For example, the DTMC in Fig. 1 satisfies the following HyperPCTL formula:

$$\psi = \forall \hat{s}. \forall \hat{s}'. (init_{\hat{s}} \wedge init_{\hat{s}'} \Rightarrow (\mathbb{P}(\diamond a_{\hat{s}}) = \mathbb{P}(\diamond a_{\hat{s}'}))) \quad (1)$$

which means that the probability of reaching proposition a from any pair of states \hat{s} and \hat{s}' labeled by $init$ should be equal. Other works on probabilistic hyperproperties for DTMCs include parameter synthesis [8] and statistical model checking [5, 9].

An important gap in the spectrum is verification of probabilistic hyperproperties with regard to models that allow *nondeterminism*, in particular, *Markov decision processes* (MDP). Nondeterminism plays a crucial role in many probabilistic systems. For instance, nondeterministic queries can be exploited in order to make targeted attacks to databases with private information [10]. To motivate the idea, consider the MDP in Fig. 2, where h is a high secret and l is a low publicly observable variable. To protect the secret, there should be no probabilistic dependencies between observations on the low variable l and the value of h . However, an attacker that chooses a scheduler that always takes action α from states s_0 and s_1 can learn whether or not $h \leq 0$ by observing the probability of obtaining $l = 1$ (or $l = 2$). On the other hand, a scheduler that always chooses action β , does not leak any information about the value of h . Thus, a natural question to ask is whether a certain property holds for all or some schedulers.

With the above motivation, in this paper, we focus on probabilistic hyperproperties in the context of MDPs. Such hyperproperties inherently need to consider different nondeterministic choices in different executions, and naturally call for quantification over schedulers. There are several challenges to achieve this. In general, there are schedulers whose reachability probabilities cannot be achieved by any memoryless non-probabilistic scheduler, and, hence finding a scheduler is not reducible to checking non-probabilistic memoryless schedulers, as it is done in PCTL mode checking for MDPs. Consider for example the MDP in Fig. 3, for

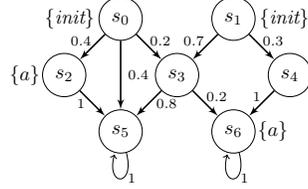


Fig. 1: Example DTMC.

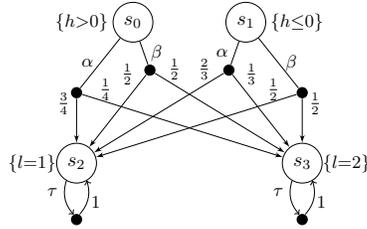


Fig. 2: Example MDP.

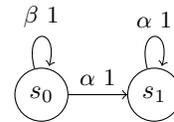


Fig. 3: Example MDP.

which we want to know whether there is a scheduler such that the probability to reach s_1 from s_0 equals 0.5. There are two non-probabilistic memoryless schedulers, one choosing action α and the other, action β in s_0 . The first one is the maximal scheduler for which s_1 is reached with probability 1, and the second one is the minimal scheduler leading to probability 0. However, the probability 0.5 cannot be achieved by any non-probabilistic scheduler. *Memoryless* probabilistic schedulers can neither achieve probability 0.5: if a memoryless scheduler would take action α with any positive probability, then the probability to reach s_1 is always 1. The only way to achieve the reachability probability 0.5 (or any value strictly between 0 and 1) is by a probabilistic scheduler with memory, e.g., taking α and β in s_0 with probabilities 0.5 each when this is the first step on a path, and β with probability 1 otherwise.

Our contributions in this paper are as follows. We first extend the temporal logic HyperPCTL [7] to the context of MDPs. To this end, we augment the syntax and semantics of HyperPCTL to quantify over schedulers and relate probabilistic computation trees for different schedulers. For example, the following formula generalizes (1) by requiring that the respective property should hold for all computation trees starting in any states \hat{s} and \hat{s}' of the DMTC induced by any scheduler $\hat{\sigma}$:

$$\forall \hat{\sigma}. \forall \hat{s}(\hat{\sigma}). \forall \hat{s}'(\hat{\sigma}). \left(\text{init}_{\hat{s}} \wedge \text{init}_{\hat{s}'} \right) \Rightarrow \left(\mathbb{P}(\diamond a_{\hat{s}}) = \mathbb{P}(\diamond a_{\hat{s}'}) \right)$$

On the negative side, we show that the problem to check HyperPCTL properties for MDPs is in general undecidable. On the positive side, we show that the problem becomes decidable when we restrict the scheduler quantification domain to memoryless non-probabilistic schedulers. We also show that this restricted problem is already NP-complete (respectively, coNP-complete) in the size of the given MDP for HyperPCTL formulas with a single existential (respectively, universal) scheduler quantifier. Subsequently, we propose an SMT-based encoding to solve the restricted model checking problem. We have implemented our method and analyze it experimentally on three case studies: probabilistic scheduling attacks, side-channel timing attacks, and probabilistic conformance (available at <https://github.com/oreohere/HyperOnMDP>).

It is important to note that the work in [11] (also published in ATVA'20) independently addresses the problem under investigation in this paper. The authors propose the temporal logic PHL. Similar to HyperPCTL, PHL also allows quantification over schedulers, but path quantification of the induced DTMC is achieved by using HyperCTL*. Both papers show that the model checking problem is undecidable for the respective logics. The difference, however, is in our approaches to deal with the undecidability result, which leads two complementary and orthogonal techniques. For both logics the problem is decidable for non-probabilistic memoryless schedulers. We provide an SMT-based verification procedure for HyperPCTL for this class of schedulers. The work in [11] presents two methods for proving and for refuting formulas from a fragment of PHL for general memoryful schedulers. The two papers offer disjoint case studies for evaluation.

Organization. Preliminary concepts are discussed in Section 2. We present the syntax and semantics of HyperPCTL for MDPs and discuss its expressive power in Section 3. Section 4 is dedicated to the applications of HyperPCTL. Sections 5 and 6 present our results on memoryless non-probabilistic schedulers and their evaluation before concluding in Section 7. All proofs are given in the Appendix.

2 Preliminaries

2.1 Discrete-time Markov models

Definition 1. A discrete-time Markov chain (DTMC) is a tuple $\mathcal{D}=(S, \mathbf{P}, \text{AP}, L)$ with the following components:

- S is a nonempty finite set of states;
- $\mathbf{P} : S \times S \rightarrow [0, 1]$ is a transition probability function with $\sum_{s' \in S} \mathbf{P}(s, s') = 1$, for all $s \in S$;
- AP is a finite set of atomic propositions, and
- $L : S \rightarrow 2^{\text{AP}}$ is a labeling function. ■

Fig 1 shows a simple DTMC. An (*infinite*) *path* of \mathcal{D} is an infinite sequence $\pi = s_0 s_1 s_2 \dots \in S^\omega$ of states with $\mathbf{P}(s_i, s_{i+1}) > 0$, for all $i \geq 0$; we write $\pi[i]$ for s_i . Let $\text{Paths}_s^{\mathcal{D}}$ denote the set of all (infinite) paths of \mathcal{D} starting in s , and $f\text{Paths}_s^{\mathcal{D}}$ denote the set of all non-empty finite prefixes of paths from $\text{Paths}_s^{\mathcal{D}}$, which we call *finite paths*. For a finite path $\pi = s_0 \dots s_k \in f\text{Paths}_{s_0}^{\mathcal{D}}$, $k \geq 0$, we define $|\pi| = k$. We will also use the notations $\text{Paths}^{\mathcal{D}} = \cup_{s \in S} \text{Paths}_s^{\mathcal{D}}$ and $f\text{Paths}^{\mathcal{D}} = \cup_{s \in S} f\text{Paths}_s^{\mathcal{D}}$. A state $t \in S$ is *reachable* from a state $s \in S$ in \mathcal{D} if there exists a finite path in $f\text{Paths}_s^{\mathcal{D}}$ with last state t ; we use $f\text{Paths}_{s,T}^{\mathcal{D}}$ to denote the set of all finite paths from $f\text{Paths}_s^{\mathcal{D}}$ with last state in $T \subseteq S$. A state $s \in S$ is *absorbing* if $\mathbf{P}(s, s) = 1$.

The *cylinder set* $\text{Cyl}^{\mathcal{D}}(\pi)$ of a finite path $\pi \in f\text{Paths}_s^{\mathcal{D}}$ is the set of all infinite paths of \mathcal{D} with prefix π . The *probability space for \mathcal{D} and state $s \in S$* is $(\text{Paths}_s^{\mathcal{D}}, \{\cup_{\pi \in R} \text{Cyl}^{\mathcal{D}}(\pi) \mid R \subseteq f\text{Paths}_s^{\mathcal{D}}\}, \text{Pr}_s^{\mathcal{D}})$, where the *probability* of the cylinder set of $\pi \in f\text{Paths}_s^{\mathcal{D}}$ is $\text{Pr}_s^{\mathcal{D}}(\text{Cyl}^{\mathcal{D}}(\pi)) = \prod_{i=1}^{|\pi|} \mathbf{P}(\pi[i-1], \pi[i])$.

Note that the cylinder sets of two finite paths starting in the same state are either disjoint or one is contained in the other. According to the definition of the probability spaces, the total probability for a set of cylinder sets defined by the finite paths $R \subseteq f\text{Paths}_s^{\mathcal{D}}$ is $\text{Pr}^{\mathcal{D}}(R) = \sum_{\pi \in R'} \text{Pr}_s^{\mathcal{D}}(\pi)$ with $R' = \{\pi \in R \mid \text{no } \pi' \in R \setminus \{\pi\} \text{ is a prefix of } \pi\}$. To improve readability, we sometimes omit the DTMC index \mathcal{D} in the notations when it is clear from the context.

Parallel composition formalizes simultaneous runs in different DTMCs.

Definition 2. The parallel composition of two DTMCs $\mathcal{D}_i = (S_i, \mathbf{P}_i, \text{AP}_i, L_i)$, $i = 1, 2$, is the DTMC $\mathcal{D}_1 \times \mathcal{D}_2 = (S, \mathbf{P}, \text{AP}, L)$ with the following components:

- $S = S_1 \times S_2$;
- $\mathbf{P} : S \times S \rightarrow [0, 1]$ with $\mathbf{P}((s_1, s_2), (s'_1, s'_2)) = \mathbf{P}_1(s_1, s'_1) \cdot \mathbf{P}_2(s_2, s'_2)$, for all states $(s_1, s_2), (s'_1, s'_2) \in S$;

- $\text{AP} = \text{AP}_1 \cup \text{AP}_2$, and
- $L : S \rightarrow 2^{\text{AP}}$ with $L((s_1, s_2)) = L_1(s_1) \cup L_2(s_2)$. ■

Markov decision processes extend DTMCs with non-deterministic choices.

Definition 3. A Markov decision process (MDP) is a tuple $\mathcal{M} = (S, \text{Act}, \mathbf{P}, \text{AP}, L)$ with the following components:

- S is a nonempty finite set of states;
- Act is a nonempty finite set of actions;
- $\mathbf{P} : S \times \text{Act} \times S \rightarrow [0, 1]$ is a transition probability function such that for all $s \in S$ the set of enabled actions in s $\text{Act}(s) = \{\alpha \in \text{Act} \mid \sum_{s' \in S} \mathbf{P}(s, \alpha, s') = 1\}$ is not empty and $\sum_{s' \in S} \mathbf{P}(s, \alpha, s') = 0$ for all $\alpha \in \text{Act} \setminus \text{Act}(s)$;
- AP is a finite set of atomic propositions, and
- $L : S \rightarrow 2^{\text{AP}}$ is a labeling function. ■

Fig. 2 shows a simple MDP. Schedulers can be used to eliminate the non-determinism in MDPs, inducing DTMCs with well-defined probability spaces.

Definition 4. A scheduler for an MDP $\mathcal{M} = (S, \text{Act}, \mathbf{P}, \text{AP}, L)$ is a tuple $\sigma = (Q, \text{act}, \text{mode}, \text{init})$, where

- Q is a countable set of modes;
- $\text{act} : Q \times S \times \text{Act} \rightarrow [0, 1]$ is a function for which $\sum_{\alpha \in \text{Act}(s)} \text{act}(q, s, \alpha) = 1$ and $\sum_{\alpha \in \text{Act} \setminus \text{Act}(s)} \text{act}(q, s, \alpha) = 0$ for all $s \in S$ and $q \in Q$;
- $\text{mode} : Q \times S \rightarrow Q$ is a mode transition function, and
- $\text{init} : S \rightarrow Q$ is a function selecting a starting mode for each state of \mathcal{M} . ■

Let $\Sigma^{\mathcal{M}}$ denote the set of all schedulers for the MDP \mathcal{M} . A scheduler is *finite-memory* if Q is finite, *memoryless* if $|Q| = 1$, and *non-probabilistic* if $\text{act}(q, s, \alpha) \in \{0, 1\}$ for all $q \in Q$, $s \in S$ and $\alpha \in \text{Act}$.

Definition 5. Assume an MDP $\mathcal{M} = (S, \text{Act}, \mathbf{P}, \text{AP}, L)$ and a scheduler $\sigma = (Q, \text{act}, \text{mode}, \text{init}) \in \Sigma^{\mathcal{M}}$ for \mathcal{M} . The DTMC induced by \mathcal{M} and σ is defined as $\mathcal{M}^\sigma = (S^\sigma, \mathbf{P}^\sigma, \text{AP}, L^\sigma)$ with $S^\sigma = Q \times S$,

$$\mathbf{P}^\sigma((q, s), (q', s')) = \begin{cases} \sum_{\alpha \in \text{Act}(s)} \text{act}(q, s, \alpha) \cdot \mathbf{P}(s, \alpha, s') & \text{if } q' = \text{mode}(q, s) \\ 0 & \text{otherwise} \end{cases}$$

and $L^\sigma(q, s) = L(s)$ for all $s, s' \in S$ and all $q, q' \in Q$. ■

A state s' is *reachable* from $s \in S$ in MDP \mathcal{M} if there exists a scheduler σ for \mathcal{M} such that s' is reachable from s in \mathcal{M}^σ . A state $s \in S$ is *absorbing* in \mathcal{M} if s is absorbing in \mathcal{M}^σ for all schedulers σ for \mathcal{M} . We sometimes omit the MDP index \mathcal{M} in the notations when it is clear from the context.

3 HyperPCTL for MDPs

In this section we extend HyperPCTL from [7] for DTMCs, to argue also about non-determinism in MDPs.

3.1 HyperPCTL Syntax

HyperPCTL (quantified) state formulas φ^q are inductively defined as follows:

$$\begin{array}{ll}
\text{quantified formula } \varphi^q & ::= \forall \hat{\sigma}.\varphi^q \mid \exists \hat{\sigma}.\varphi^q \mid \forall \hat{s}(\hat{\sigma}).\varphi^q \mid \exists \hat{s}(\hat{\sigma}).\varphi^q \mid \varphi^{nq} \\
\text{non-quantified formula } \varphi^{nq} & ::= \mathbf{true} \mid a_{\hat{s}} \mid \varphi^{nq} \wedge \varphi^{nq} \mid \neg \varphi^{nq} \mid \varphi^{pr} < \varphi^{pr} \\
\text{probability expression } \varphi^{pr} & ::= \mathbb{P}(\varphi^{path}) \mid f(\varphi_1^{pr}, \dots, \varphi_k^{pr}) \\
\text{path formula } \varphi^{path} & ::= \bigcirc \varphi^{nq} \mid \varphi^{nq} \mathcal{U} \varphi^{nq} \mid \varphi^{nq} \mathcal{U}^{[k_1, k_2]} \varphi^{nq}
\end{array}$$

where $\hat{\sigma}$ is a *scheduler variable*¹ from an infinite set $\hat{\Sigma}$, \hat{s} is a *state variable* from an infinite set \hat{S} , φ^{nq} is a quantifier-free state formula, $a \in \text{AP}$ is an atomic proposition, φ^{pr} is a *probability expression*, $f : [0, 1]^k \rightarrow \mathbb{R}$ are k -ary arithmetic operators (binary addition, unary/binary subtraction, binary multiplication) over probabilities, where constants are viewed as 0-ary functions, and φ^{path} is a *path formula*, such that $k_1 \leq k_2 \in \mathbb{N}_{\geq 0}$. The probability operator \mathbb{P} allows the usage of probabilities in arithmetic constraints and relations.

A HyperPCTL construct φ (probability expression φ^{pr} , state formula φ^q , φ^{nq} or path formula φ^{path}) is *well-formed* if each occurrence of any $a_{\hat{s}}$ with $a \in \text{AP}$ and $\hat{s} \in \hat{S}$ is in the scope of a *state quantifier* for $\hat{s}(\hat{\sigma})$ for some $\hat{\sigma} \in \hat{\Sigma}$, and any quantifier for $\hat{s}(\hat{\sigma})$ is in the scope of a *scheduler quantifier* for $\hat{\sigma}$. We restrict ourselves to quantifying first the schedulers then the states, i.e., different state variables can share the same scheduler. One can consider also *local* schedulers when different players cannot explicitly share the same scheduler, or in other words, each scheduler quantifier belongs to exactly one of the quantified states.

HyperPCTL formulas are well-formed HyperPCTL state formulas, where we additionally allow standard syntactic sugar like $\mathbf{false} = \neg \mathbf{true}$, $\varphi_1 \vee \varphi_2 = \neg(\neg \varphi_1 \wedge \neg \varphi_2)$, $\diamond \varphi = \mathbf{true} \mathcal{U} \varphi$, and $\mathbb{P}(\square \varphi) = 1 - \mathbb{P}(\diamond \neg \varphi)$. For example, the HyperPCTL state formula $\forall \hat{\sigma}.\exists \hat{s}(\hat{\sigma}).\mathbb{P}(\bigcirc a_{\hat{s}}) < 0.5$ is a HyperPCTL formula. The HyperPCTL state formula $\mathbb{P}(\bigcirc a_{\hat{s}}) < 0.5$ is not a HyperPCTL formula, but can be extended to such. The HyperPCTL state formula $\forall \hat{s}(\hat{\sigma}).\exists \hat{\sigma}.\mathbb{P}(\bigcirc a_{\hat{s}}) < 0.5$ is not a HyperPCTL formula, and it even cannot be extended to such.

3.2 HyperPCTL Semantics

The semantics of HyperPCTL is based on the n -ary *self-composition* of an MDP.

Definition 6. *The n -ary self-composition of an MDP $\mathcal{M} = (S, \text{Act}, \mathbf{P}, \text{AP}, L)$ for a sequence $\sigma = (\sigma_1, \dots, \sigma_n) \in (\Sigma^{\mathcal{M}})^n$ of schedulers for \mathcal{M} is the DTMC parallel composition $\mathcal{M}^\sigma = \mathcal{M}_1^{\sigma_1} \times \dots \times \mathcal{M}_n^{\sigma_n}$, where $\mathcal{M}_i^{\sigma_i}$ is the DTMC induced by \mathcal{M}_i and σ_i , and where $\mathcal{M}_i = (S, \text{Act}, \mathbf{P}, \text{AP}_i, L_i)$ with $\text{AP}_i = \{a_i \mid a \in \text{AP}\}$ and $L_i(s) = \{a_i \mid a \in L(s)\}$, for all $s \in S$. ■*

HyperPCTL state formulas are evaluated in the context of an MDP $\mathcal{M} = (S, \text{Act}, \mathbf{P}, \text{AP}, L)$, a sequence $\sigma = (\sigma_1, \dots, \sigma_n) \in (\Sigma^{\mathcal{M}})^n$ of schedulers, and a

¹ We use the notation $\hat{\sigma}$ for scheduler variables and σ for schedulers, and analogously \hat{s} for state variables and s for states.

sequence $\mathbf{r} = ((q_1, s_1), \dots, (q_n, s_n))$ of \mathcal{M}^σ states; we use $()$ to denote the empty sequence (of any type) and \circ for concatenation. Intuitively, these sequences store instantiations for scheduler and state variables. The satisfaction of a HyperPCTL quantified formula by \mathcal{M} is defined by

$$\mathcal{M} \models \varphi \quad \text{iff} \quad \mathcal{M}, (), () \models \varphi .$$

The semantics evaluates HyperPCTL formulas by structural recursion. Let in the following $\mathbb{Q}, \mathbb{Q}', \dots$ denote quantifiers from $\{\forall, \exists\}$. When instantiating $\mathbb{Q}\hat{\sigma}.\varphi$ by a scheduler $\sigma \in \Sigma^\mathcal{M}$, we replace in φ each subformula $\mathbb{Q}'\hat{s}(\hat{\sigma}).\varphi'$, that is not in the scope of a quantifier for $\hat{\sigma}$ by $\mathbb{Q}'\hat{s}(\sigma).\varphi'$, and denote the result by $\varphi[\hat{\sigma} \rightsquigarrow \sigma]$. For instantiating a state quantifier $\mathbb{Q}\hat{s}(\sigma).\varphi$ by a state s , we append $\sigma = (Q, act, mode, init)$ and $(init(s), s)$ at the end of the respective sequences, and replace each $a_{\hat{s}}$ in the scope of the given quantifier by a_s , resulting in a formula that we denote by $\varphi[\hat{s} \rightsquigarrow s]$. To evaluate probability expressions, we use the n -ary self-composition of the MDP.

Formally, the semantics judgment rules are as follows:

$$\begin{aligned} \mathcal{M}, \sigma, \mathbf{r} &\models \text{true} \\ \mathcal{M}, \sigma, \mathbf{r} &\models a_i && \text{iff } a_i \in L^\sigma(\mathbf{r}) \\ \mathcal{M}, \sigma, \mathbf{r} &\models \varphi_1 \wedge \varphi_2 && \text{iff } \mathcal{M}, \sigma, \mathbf{r} \models \varphi_1 \text{ and } \mathcal{M}, \sigma, \mathbf{r} \models \varphi_2 \\ \mathcal{M}, \sigma, \mathbf{r} &\models \neg \varphi && \text{iff } \mathcal{M}, \sigma, \mathbf{r} \not\models \varphi \\ \mathcal{M}, \sigma, \mathbf{r} &\models \forall \hat{\sigma}.\varphi && \text{iff } \forall \sigma \in \Sigma^\mathcal{M}. \mathcal{M}, \sigma, \mathbf{r} \models \varphi[\hat{\sigma} \rightsquigarrow \sigma] \\ \mathcal{M}, \sigma, \mathbf{r} &\models \exists \hat{\sigma}.\varphi && \text{iff } \exists \sigma \in \Sigma^\mathcal{M}. \mathcal{M}, \sigma, \mathbf{r} \models \varphi[\hat{\sigma} \rightsquigarrow \sigma] \\ \mathcal{M}, \sigma, \mathbf{r} &\models \forall \hat{s}(\sigma).\varphi && \text{iff } \forall s_{n+1} \in S. \mathcal{M}, \sigma \circ \sigma, \mathbf{r} \circ (init(s_{n+1}), s_{n+1}) \models \varphi[\hat{s} \rightsquigarrow s_{n+1}] \\ \mathcal{M}, \sigma, \mathbf{r} &\models \exists \hat{s}(\sigma).\varphi && \text{iff } \exists s_{n+1} \in S. \mathcal{M}, \sigma \circ \sigma, \mathbf{r} \circ (init(s_{n+1}), s_{n+1}) \models \varphi[\hat{s} \rightsquigarrow s_{n+1}] \\ \mathcal{M}, \sigma, \mathbf{r} &\models \varphi_1^{pr} < \varphi_2^{pr} && \text{iff } \llbracket \varphi_1^{pr} \rrbracket_{\mathcal{M}, \sigma, \mathbf{r}} < \llbracket \varphi_2^{pr} \rrbracket_{\mathcal{M}, \sigma, \mathbf{r}} \\ \llbracket \mathbb{P}(\varphi_{path}) \rrbracket_{\mathcal{M}, \sigma, \mathbf{r}} &= && \Pr^{\mathcal{M}^\sigma}(\{\pi \in Paths_{\mathbf{r}}^{\mathcal{M}^\sigma} \mid \mathcal{M}, \sigma, \pi \models \varphi_{path}\}) \\ \llbracket f(\varphi_1^{pr}, \dots, \varphi_k^{pr}) \rrbracket_{\mathcal{M}, \sigma, \mathbf{r}} &= && f(\llbracket \varphi_1^{pr} \rrbracket_{\mathcal{M}, \sigma, \mathbf{r}}, \dots, \llbracket \varphi_k^{pr} \rrbracket_{\mathcal{M}, \sigma, \mathbf{r}}) \end{aligned}$$

where \mathcal{M} is an MDP; $n \in \mathbb{N}_{\geq 0}$ is non-negative integer; $\sigma \in (\Sigma^\mathcal{M})^n$; \mathbf{r} is a state of \mathcal{M}^σ ; $a \in \text{AP}$ is an atomic proposition and $i \in \{1, \dots, n\}$; $\varphi, \varphi_1, \varphi_2$ are HyperPCTL state formulas; $\sigma = (Q, act, mode, init) \in \Sigma^\mathcal{M}$ is a scheduler for \mathcal{M} ; $\varphi_1^{pr} \dots \varphi_k^{pr}$ are probability expressions, and φ_{path} is a HyperPCTL path formula whose satisfaction relation is as follows:

$$\begin{aligned} \mathcal{M}, \sigma, \pi &\models \bigcirc \varphi && \text{iff } \mathcal{M}, \sigma, \mathbf{r}_1 \models \varphi \\ \mathcal{M}, \sigma, \pi &\models \varphi_1 \mathcal{U} \varphi_2 && \text{iff } \exists j \geq 0. (\mathcal{M}, \sigma, \mathbf{r}_j \models \varphi_2 \wedge \forall i \in [0, j). \mathcal{M}, \sigma, \mathbf{r}_i \models \varphi_1) \\ \mathcal{M}, \sigma, \pi &\models \varphi_1 \mathcal{U}^{[k_1, k_2]} \varphi_2 && \text{iff } \exists j \in [k_1, k_2]. (\mathcal{M}, \sigma, \mathbf{r}_j \models \varphi_2 \wedge \forall i \in [0, j). \mathcal{M}, \sigma, \mathbf{r}_i \models \varphi_1) \end{aligned}$$

where $\pi = \mathbf{r}_0 \mathbf{r}_1 \dots$ with $\mathbf{r}_i = ((q_{i,1}, s_{i,1}), \dots, (q_{i,n}, s_{i,n}))$ is a path of \mathcal{M}^σ ; formulas φ, φ_1 , and φ_2 are HyperPCTL state formulas, and $k_1 \leq k_2 \in \mathbb{N}_{\geq 0}$.

3.3 The Expressiveness Power of HyperPCTL

For MDPs with $|Act(s)| = 1$ for each of its states s , the HyperPCTL semantics reduces to the one proposed in [7] for DTMCs.

For MDPs with non-determinism, the standard PCTL semantics defines that in order to satisfy a PCTL formula $\mathbb{P}_{\sim c}(\psi)$ in a given MDP state s , *all* schedulers should induce a DTMC that satisfies $\mathbb{P}_{\sim c}(\psi)$ in s . Though it should hold for *all* schedulers, it is known that there exist minimal and maximal schedulers that are non-probabilistic and memoryless, therefore it is sufficient to restrict the reasoning to such schedulers. Since for MDPs with finite state and action spaces, the number of such schedulers is finite, PCTL model checking for MDPs is decidable. Given this analogy, one would expect that HyperPCTL model checking should be decidable, but it is not.

Theorem 1. *HyperPCTL model checking for MDPs is in general undecidable.*

What is the source of increased expressiveness that makes HyperPCTL undecidable? State quantification cannot be the source, as the state space is finite and thus there are finitely many possible state quantifier instantiations.

Assume an MDP $\mathcal{M} = (S, Act, \mathbf{P}, AP, L)$ with a state $s \in S$ that is uniquely labelled by the proposition $init \in L(s)$, and let $a, b \in AP$. In PCTL, each probability bound needs to be satisfied under all schedulers. For example:

$$\mathcal{M}, s \models_{\text{PCTL}} \mathbb{P}_{<0.5}(aUb) \text{ iff } \mathcal{M} \models_{\text{HyperPCTL}} \forall \hat{\sigma}. \forall \hat{s}(\hat{\sigma}). (init_{\hat{s}} \rightarrow \mathbb{P}(aUb) < 0.5)$$

Alternatively, we can state:

$$\mathcal{M}, s \models_{\text{PCTL}} \mathbb{P}_{<0.5}(aUb) \text{ iff } \forall \sigma \in \Sigma^{\mathcal{M}}. \mathcal{M}, (\sigma), ((init_{\sigma}(s), s)) \models_{\text{HyperPCTL}} \mathbb{P}(aUb) < 0.5$$

where $init_{\sigma}(s)$ is the starting mode of scheduler σ in state s . Generally, the HyperPCTL fragment which starts with a single universal scheduler quantifier and contains a single bound on a single probability operator is still decidable. However, when a PCTL formula has several probability bounds, its satisfaction requires each bound to be satisfied by all schedulers *independently*. For example,

$$\begin{aligned} \mathcal{M}, s \models_{\text{PCTL}} \left(\mathbb{P}_{<0.5}(aUb) \vee \mathbb{P}_{>0.5}(aUb) \right) \text{ iff} \\ \forall \sigma \in \Sigma^{\mathcal{M}}. \mathcal{M}, (\sigma), ((init_{\sigma}(s), s)) \models_{\text{HyperPCTL}} \mathbb{P}(aUb) < 0.5 \text{ or} \\ \mathcal{M}, (\sigma), ((init_{\sigma}(s), s)) \models_{\text{HyperPCTL}} \mathbb{P}(aUb) > 0.5 \end{aligned}$$

This is *not* equivalent to the HyperPCTL formula

$$\mathcal{M} \models_{\text{HyperPCTL}} \forall \hat{\sigma}. \forall \hat{s}(\hat{\sigma}). (init_{\hat{s}} \rightarrow (\mathbb{P}(a_{\hat{s}}Ub_{\hat{s}}) < 0.5 \vee \mathbb{P}(a_{\hat{s}}Ub_{\hat{s}}) > 0.5))$$

which states that the probability is either less than or larger than 0.5 under all schedulers, which is true if *there exists no scheduler* under which the probability is 0.5 (see also [12]). Thus, even for a fragment restricted to universal scheduler quantification, combinations of probability bounds allows HyperPCTL to express existential *scheduler synthesis* problems.

Finally, consider a scheduler quantifier followed by state quantifiers, whose scope may contain probability expressions. This means we start several “experiments” in parallel, each one represented by a state quantifier. However, we may

use in all experiments the *same scheduler*. Informally, this allows us to express the existence or absence of schedulers with certain probabilistic hyperproperties for the induced DTMCs. It would however also make sense to flip this quantifier order, such that state quantifiers are followed by scheduler quantifiers. This would mean, that we can use different schedulers in the different concurrently running experiments. This would be meaningful e.g. when users can provide input to the system, i.e. when the scheduler choice lies by the “observers” of the individual experiments, and they can adapt their schedulers to observations made in the other concurrently running experiments.

4 Applications of HyperPCTL on MDPs

Side-channel timing leaks open a channel to an attacker to infer the value of a secret by observing the execution time of a function. For example, the heart of the RSA public-key encryption algorithm is the modular exponentiation algorithm that computes $(a^b \bmod n)$, where a is an integer representing the plain-
 text and b is the integer encryption key. A careless implementation can leak b through a probabilistic scheduling channel (see Fig. 4). This program is not secure since the two

```

1 void mexp(){
2   c = 0; d = 1; i = k;
3   while (i >= 0){
4     i = i - 1; c = c * 2;
5     d = (d * d) % n;
6     if (b(i) = 1)
7       c = c + 1;
8     d = (d * a) % n;
9   }
10 }
11 /*****/
12 t = new Thread(mexp());
13 j = 0; m = 2 * k;
14 while (j < m & !t.stop) j++;
15 /*****/
    
```

Fig. 4: Modular exponentiation.

branches of the *if* have different timing behaviors. Under a fair execution scheduler for parallel threads, an attacker thread can infer the value of b by running in parallel to a modular exponentiation thread and iteratively incrementing a counter variable until the other thread terminates (lines 12-14). To model this program by an MDP, we can use two nondeterministic actions for the two branches of the *if* statement, such that the choice of different schedulers corresponds to the choice of different bit configurations $\mathbf{b}(i)$ for the key \mathbf{b} . This algorithm should satisfy the following property: the probability of observing a concrete value in the counter j should be independent of the bit configuration of the secret key \mathbf{b} :

$$\forall \hat{\sigma}_1. \forall \hat{\sigma}_2. \forall \hat{s}(\hat{\sigma}_1). \forall \hat{s}'(\hat{\sigma}_2). \left(\text{init}_{\hat{s}} \wedge \text{init}_{\hat{s}'} \right) \Rightarrow \bigwedge_{l=0}^m \left(\mathbb{P}(\diamond(j=l)_{\hat{s}}) = \mathbb{P}(\diamond(j=l)_{\hat{s}'}) \right)$$

Another example of timing attacks that can be implemented through a probabilistic scheduling side channel is password verification which is typically implemented by comparing an input string with another confidential string (see Fig 5). Also here, an attacker thread can measure the time necessary to break the loop, and use this information to infer the prefix of the input string matching the secret string.

```

1 int str_cmp(char * r){
2   char * s = 'Bg$4\0';
3   i = 0;
4   while (s[i] != '\0'){
5     i++;
6     if (s[i] != r[i]) return 0;
7   }
8   return 1;
9 }
    
```

Fig. 5: String comparison.

Scheduler-specific observational determinism policy (SSODP) [13] is a confidentiality policy in multi-threaded programs that defends against an attacker that chooses an appropriate scheduler to control the set of possible traces. In particular, given any scheduler and two initial states that are indistinguishable with respect to a secret input (i.e., low-equivalent), any two executions from these two states should terminate in low-equivalent states with equal probability. Formally, given a proposition h representing a secret:

$$\forall \hat{\sigma}. \forall \hat{s}(\hat{\sigma}). \forall \hat{s}'(\hat{\sigma}). (h_{\hat{s}} \oplus h_{\hat{s}'}) \Rightarrow \bigwedge_{l \in L} (\mathbb{P}(\diamond l_{\hat{s}}) = \mathbb{P}(\diamond l_{\hat{s}'}))$$

where $l \in L$ are atomic propositions that classify low-equivalent states and \oplus is the exclusive-or operator. A stronger variation of this policy is that the executions are stepwise low-equivalent:

$$\forall \hat{\sigma}. \forall \hat{s}(\hat{\sigma}). \forall \hat{s}'(\hat{\sigma}). (h_{\hat{s}} \oplus h_{\hat{s}'}) \Rightarrow \mathbb{P}\square \left(\bigwedge_{l \in L} ((\mathbb{P} \circ l_{\hat{s}}) = (\mathbb{P} \circ l_{\hat{s}'})) \right) = 1.$$

Probabilistic conformance describes how well a model and an implementation conform with each other with respect to a specification. As an example, consider a 6-sided die. The probability to obtain one possible side of the die is 1/6. We would like to synthesize a protocol that simulates the 6-sided die behavior only by repeatedly tossing a fair coin. We know that such an implementation exists [14], but our aim is to find such a solution automatically by modeling the die as a DTMC and by using an MDP to model all the possible coin-implementations with a given maximum number of states, including 6 absorbing final states to model the outcomes. In the MDP, we associate to each state a set of possible nondeterministic actions, each of them choosing two states as successors with equal probability 1/2. Then, each scheduler corresponds to a particular implementation. Our goal is to check whether there exists a scheduler that induces a DTMC over the MDP, such that repeatedly tossing a coin simulates die-rolling with equal probabilities for the different outcomes:

$$\exists \hat{\sigma}. \forall \hat{s}(\hat{\sigma}). \exists \hat{s}'(\hat{\sigma}). (init_{\hat{s}} \wedge init_{\hat{s}'}) \Rightarrow \bigwedge_{l=1}^6 (\mathbb{P}(\diamond (die = l)_{\hat{s}}) = \mathbb{P}(\diamond (die = l)_{\hat{s}'}))$$

5 HyperPCTL Model Checking for Non-probabilistic Memoryless Schedulers

Due to the undecidability of model checking HyperPCTL formulas for MDPs, we now restrict ourselves to the semantics, where scheduler quantification ranges over non-probabilistic memoryless schedulers only. It is easy to see that this restriction makes the model checking problem decidable, as there are only finitely many such schedulers that can be enumerated. Regarding complexity, we have the following property.

Theorem 2. *The problem to decide for MDPs the truth of HyperPCTL formulas with a single existential (respectively, universal) scheduler quantifier over non-probabilistic memoryless schedulers is NP-complete (respectively, coNP-complete) in the state set size of the given MDP.*

Next we propose an SMT-based technique for solving the model checking problem for non-probabilistic memoryless scheduler domains, and for the simplified case of having a single scheduler quantifier; the general case for an arbitrary number of scheduler quantifiers is similar, but a bit more involved, so the simplified setting might be more suitable for understanding the basic ideas.

The main method listed in Algorithm 1 constructs a formula E that is satisfiable if and only if the input MDP \mathcal{M} satisfies the input HyperPCTL formula with a single scheduler quantifier over the non-probabilistic memoryless scheduler domain. Let us first deal with the case that the scheduler quantifier is *existential*. In line 2 we encode possible instantiations σ for the scheduler variable $\hat{\sigma}$, for which we use a variable σ_s for each MDP state $s \in S$ to encode which action is chosen in that

Algorithm 1: Main SMT encoding algorithm

Input : $\mathcal{M} = (S, Act, \mathbf{P}, AP, L)$: MDP;
 $Q\hat{\sigma}.Q_1\hat{s}_1(\hat{\sigma}) \dots Q_n\hat{s}_n(\hat{\sigma}).\varphi^{nq}$: HyperPCTL formula.

Output: Whether \mathcal{M} satisfies the input formula.

Function $Main(\mathcal{M}, Q\hat{\sigma}.Q_1\hat{s}_1(\hat{\sigma}) \dots Q_n\hat{s}_n(\hat{\sigma}).\varphi^{nq})$

```

2   $E := \bigwedge_{s \in S} (\bigvee_{\alpha \in Act(s)} \sigma_s = \alpha)$  // scheduler choice
3  if  $Q$  is existential then
4       $E := E \wedge \text{Semantics}(\mathcal{M}, \varphi^{nq}, n)$ ;
5       $E := E \wedge \text{Truth}(\mathcal{M}, \exists \hat{\sigma}. Q_1\hat{s}_1(\hat{\sigma}) \dots Q_n\hat{s}_n(\hat{\sigma}).\varphi^{nq})$ ;
6      if  $check(E) = SAT$  then return  $TRUE$ ;
7      else return  $FALSE$ ;
8  else if  $Q$  is universal then
9      //  $\bar{Q}_i$  is  $\forall$  if  $Q_i = \exists$  and  $\exists$  else
10      $E := E \wedge \text{Semantics}(\mathcal{M}, \neg\varphi^{nq}, n)$ ;
11      $E := E \wedge \text{Truth}(\mathcal{M}, \exists \hat{\sigma}. \bar{Q}_1\hat{s}_1(\hat{\sigma}) \dots \bar{Q}_n\hat{s}_n(\hat{\sigma}).\neg\varphi^{nq})$ ;
12     if  $check(E) = SAT$  then return  $FALSE$ ;
13     else return  $TRUE$ ;

```

state. In line 4 we encode the meaning of the quantifier-free inner part φ^{nq} of the input formula, whereas line 5 encodes the meaning of the state quantifiers, i.e. for which sets of composed states φ^{nq} needs to hold in order to satisfy the input formula. In lines 6–7 we check the satisfiability of the encoding and return the corresponding answer. Formulas with a *universal* scheduler quantifier $\forall \hat{\sigma}.\varphi$ are semantically equivalent to $\neg \exists \hat{\sigma}.\neg \varphi$. We make use of this fact in lines 8–12 to check first the satisfaction of an encoding for $\exists \hat{\sigma}.\neg \varphi$ and return the inverted answer.

The Semantics method, shown in Algorithm 2, applies structural recursion to encode the meaning of its quantifier-free input formula. As variables, the encoding uses (1) propositions $holds_{\mathbf{s}, \varphi^{nq}} \in \{\mathbf{true}, \mathbf{false}\}$ to encode the truth of each Boolean sub-formula φ^{nq} of the input formula in each state $\mathbf{s} \in S^n$ of the n -ary self-composition of \mathcal{M} , (2) numeric variables $prob_{\mathbf{s}, \varphi^{pr}} \in [0, 1] \subseteq \mathbb{R}$ to encode the value of each probability expression φ^{pr} in the input formula in the context of each composed state $\mathbf{s} \in S^n$, (3) variables $holdsToInt_{\mathbf{s}, \varphi^{pr}} \in \{0, 1\}$ to encode truth values in a pseudo-Boolean form, i.e. we set $holdsToInt_{\mathbf{s}, \varphi^{pr}} = 1$ for $holds_{\mathbf{s}, \varphi^{nq}} = \mathbf{true}$ and $prob_{\mathbf{s}, \varphi^{pr}} = 0$ else and (4) variables $d_{\mathbf{s}, \varphi}$ to encode the existence of a loop-free path from state \mathbf{s} to a state satisfying φ .

Algorithm 2: SMT encoding for the meaning of the input formula

Input : $\mathcal{M} = (S, Act, \mathbf{P}, AP, L)$: MDP;
 φ : quantifier-free HyperPCTL formula or expression;
 n : number of state variables in φ .

Output : SMT encoding of the meaning of φ in the n -ary self-composition of \mathcal{M} .

1 Function *Semantics*(\mathcal{M}, φ, n)

2 **if** φ *is true* **then** $E := \bigwedge_{s \in S^n} \text{holds}_{s, \varphi}$;

3 **else if** φ *is* a_{s_i} **then**

4 $E := (\bigwedge_{s \in S^n, a \in L(s_i)} (\text{holds}_{s, \varphi})) \wedge (\bigwedge_{s \in S^n, a \notin L(s_i)} (\neg \text{holds}_{s, \varphi}))$;

5 **else if** φ *is* $\varphi_1 \wedge \varphi_2$ **then**

6 $E := \text{Semantics}(\mathcal{M}, \varphi_1, n) \wedge \text{Semantics}(\mathcal{M}, \varphi_2, n) \wedge$

7 $\bigwedge_{s \in S^n} ((\text{holds}_{s, \varphi} \wedge \text{holds}_{s, \varphi_1} \wedge \text{holds}_{s, \varphi_2}) \vee (\neg \text{holds}_{s, \varphi} \wedge (\neg \text{holds}_{s, \varphi_1} \vee \neg \text{holds}_{s, \varphi_2})))$;

8 **else if** φ *is* $\neg \varphi'$ **then**

9 $E := \text{Semantics}(\mathcal{M}, \varphi', n) \wedge \bigwedge_{s \in S^n} (\text{holds}_{s, \varphi} \oplus \text{holds}_{s, \varphi'})$;

10 **else if** φ *is* $\varphi_1 < \varphi_2$ **then**

11 $E := \text{Semantics}(\mathcal{M}, \varphi_1, n) \wedge \text{Semantics}(\mathcal{M}, \varphi_2, n) \wedge$

12 $\bigwedge_{s \in S^n} ((\text{holds}_{s, \varphi} \wedge \text{prob}_{s, \varphi_1} < \text{prob}_{s, \varphi_2}) \vee (\neg \text{holds}_{s, \varphi} \wedge \text{prob}_{s, \varphi_1} \geq \text{prob}_{s, \varphi_2}))$;

13 **else if** φ *is* $\mathbb{P}(\bigcirc \varphi')$ **then**

14 $E := \text{Semantics}(\mathcal{M}, \varphi', n) \wedge$

15 $\bigwedge_{s \in S^n} ((\text{holdsToInt}_{s, \varphi'} = 1 \wedge \text{holds}_{s, \varphi'}) \vee (\text{holdsToInt}_{s, \varphi'} = 0 \wedge \neg \text{holds}_{s, \varphi'}))$;

16 **foreach** $\mathbf{s} = (s_1, \dots, s_n) \in S^n$ **do**

17 **foreach** $\alpha = (\alpha_1, \dots, \alpha_n) \in Act(s_1) \times \dots \times Act(s_n)$ **do**

18 $E := E \wedge ([\bigwedge_{i=1}^n \sigma_{s_i} = \alpha_i] \rightarrow [\text{prob}_{s, \varphi} =$

19 $\sum_{s' \in \text{supp}(\alpha_1) \times \dots \times \text{supp}(\alpha_n)} ((\prod_{i=1}^n \mathbf{P}(s_i, \alpha_i, s'_i)) \cdot \text{holdsToInt}_{s', \varphi'})])$;

20 **else if** φ *is* $\mathbb{P}(\varphi_1 \mathcal{U} \varphi_2)$ **then** $E := \text{SemanticsUnboundedUntil}(\mathcal{M}, \varphi, n)$;

21 **else if** φ *is* $\mathbb{P}(\varphi_1 \mathcal{U}^{[k_1, k_2]} \varphi_2)$ **then** $E := \text{SemanticsBoundedUntil}(\mathcal{M}, \varphi, n)$;

22 **else if** φ *is* c **then** $E := \bigwedge_{s \in S^n} (\text{prob}_{s, \varphi} = c)$;

23 **else if** φ *is* $\varphi_1 \text{ op } \varphi_2$ **/*** $\text{op} \in \{+, -, *\}$ ***/** **then**

24 $E := \text{Semantics}(\mathcal{M}, \varphi_1, n) \wedge \text{Semantics}(\mathcal{M}, \varphi_2, n) \wedge$

$\bigwedge_{s \in S^n} (\text{prob}_{s, \varphi} = (\text{prob}_{s, \varphi_1} \text{ op } \text{prob}_{s, \varphi_2}))$;

25 **return** E ;

There are two base cases: the Boolean constant **true** holds in all states (line 2), whereas atomic propositions hold in exactly those states that are labelled by them (line 3). For conjunction (line 5) we recursively encode the truth values of the operands and state that the conjunction is true iff both operands are true. For negation (line 8) we again encode the meaning of the operand recursively and flip its truth value. For the comparison of two probability expressions (line 10), we recursively encode the probability values of the operands and state the respective relation between them for the satisfaction of the comparison.

The remaining cases encode the semantics of probability expressions. The cases for constants (line 22) and arithmetic operations (line 23) are straightforward. For the probability $\mathbb{P}(\bigcirc \varphi')$ (line 13), we encode the Boolean value of φ' in the variables $\text{holds}_{s, \varphi'}$ (line 14), turn them into pseudo-Boolean values $\text{holdsToInt}_{s, \varphi'}$ (1 for true and 0 for false, line 15), and state that for each composed state, the probability value of $\mathbb{P}(\bigcirc \varphi')$ is the sum of the probabilities to

Algorithm 3: SMT encoding for the meaning of unbounded until formulas

Input : $\mathcal{M} = (S, Act, \mathbf{P}, AP, L)$: MDP; φ : HyperPCTL unbounded until formula of the form $\mathbb{P}(\varphi_1 \mathcal{U} \varphi_2)$; n : number of state variables in φ .
Output : SMT encoding of φ 's meaning in the n -ary self-composition of \mathcal{M} .

```

1 Function SemanticsUnboundedUntil( $\mathcal{M}, \varphi = \mathbb{P}(\varphi_1 \mathcal{U} \varphi_2), n$ )
2    $E := \text{Semantics}(\mathcal{M}, \varphi_1, n) \wedge \text{Semantics}(\mathcal{M}, \varphi_2, n)$ ;
3   foreach  $\mathbf{s} = (s_1, \dots, s_n) \in S^n$  do
4      $E := E \wedge (\text{holds}_{\mathbf{s}, \varphi_2} \rightarrow \text{prob}_{\mathbf{s}, \varphi} = 1) \wedge ((\neg \text{holds}_{\mathbf{s}, \varphi_1} \wedge \neg \text{holds}_{\mathbf{s}, \varphi_2}) \rightarrow \text{prob}_{\mathbf{s}, \varphi} = 0)$ ;
5     foreach  $\alpha = (\alpha_1, \dots, \alpha_n) \in Act(s_1) \times \dots \times Act(s_n)$  do
6        $E := E \wedge ([\text{holds}_{\mathbf{s}, \varphi_1} \wedge \neg \text{holds}_{\mathbf{s}, \varphi_2} \wedge \bigwedge_{i=1}^n \sigma_{s_i} = \alpha_i] \rightarrow$ 
7          $[\text{prob}_{\mathbf{s}, \varphi} = \sum_{\mathbf{s}' \in \text{supp}(\alpha_1) \times \dots \times \text{supp}(\alpha_n)} ((\prod_{i=1}^n \mathbf{P}(s_i, \alpha_i, s'_i)) \cdot \text{prob}_{\mathbf{s}', \varphi}) \wedge$ 
8          $(\text{prob}_{\mathbf{s}, \varphi} > 0 \rightarrow (\bigvee_{\mathbf{s}' \in \text{supp}(\alpha_1) \times \dots \times \text{supp}(\alpha_n)} (\text{holds}_{\mathbf{s}', \varphi_2} \vee d_{\mathbf{s}, \varphi_2} > d_{\mathbf{s}', \varphi_2})))]]$ );
9   return  $E$ ;
    
```

get to a successor state where the operand φ' holds; since the successors and their probabilities are scheduler-dependent, we need to iterate over all scheduler choices and use $\text{supp}(\alpha_i)$ to denote the support $\{s \in S \mid \alpha_i(s) > 0\}$ of the distribution α_i (line 17). The encodings for the probabilities of unbounded until formulas (line 20) and bounded until formulas (line 21) are listed in Algorithm 3 and 4, respectively.

For the probabilities $\mathbb{P}(\varphi_1 \mathcal{U} \varphi_2)$ to satisfy an unbounded until formula, the method *SemanticsUnboundedUntil* shown in Algorithm 3 first encodes the meaning of the until operands (line 2). For each composed state $\mathbf{s} \in S^n$, the probability of satisfying the until formula in \mathbf{s} is encoded in the variable $\text{prob}_{\mathbf{s}, \mathbb{P}(\varphi_1 \mathcal{U} \varphi_2)}$. If the second until-operand φ_2 holds in \mathbf{s} then this probability is 1 and if none of the operands are true in \mathbf{s} then it is 0 (line 4). Otherwise, depending on the scheduler σ of \mathcal{M} (line 5), the value of $\text{prob}_{\mathbf{s}, \mathbb{P}(\varphi_1 \mathcal{U} \varphi_2)}$ is a sum, adding up for each successor state \mathbf{s}' of \mathbf{s} the probability to get from \mathbf{s} to \mathbf{s}' in one step times the probability to satisfy the until-formula on paths starting in \mathbf{s}' (line 7). However, these encodings work only when at least one state satisfying φ_2 is reachable from \mathbf{s} with a positive probability: for any bottom SCC whose states all violate φ_2 , the probability $\mathbb{P}(\varphi_1 \mathcal{U} \varphi_2)$ is obviously 0, however, assigning any fixed value from $[0, 1]$ to all states of this bottom SCC would yield a fixed-point for the underlying equation system. To assure correctness, in line 8 we enforce smallest fixed-points by requiring that if $\text{prob}_{\mathbf{s}, \mathbb{P}(\varphi_1 \mathcal{U} \varphi_2)}$ is positive then there exists a loop-free path from \mathbf{s} to any state satisfying φ_2 . In the encoding of this property we use fresh variables $d_{\mathbf{s}, \varphi_2}$ and require a path over states with strong monotonically decreasing $d_{\mathbf{s}, \varphi_2}$ -values to a φ_2 -state (where the decreasing property serves to exclude loops). The domain of the distance-variables $d_{\mathbf{s}, \varphi_2}$ can be e.g. integers, rationals or reals; the only restriction is that it should contain at least $|S|^n$ ordered values. Especially, it does not need to be lower bounded (note that each solution assigns to each $d_{\mathbf{s}, \varphi_2}$ a fixed value, leading a finite number of distance values).

Algorithm 4: SMT encoding for the meaning of bounded until formulas

Input : $\mathcal{M} = (S, Act, \mathbf{P}, AP, L)$: MDP; φ : HyperPCTL bounded until formula of the form $\mathbb{P}(\varphi_1 \mathcal{U}^{[k_1, k_2]} \varphi_2)$; n : number of state variables in φ .

Output : SMT encoding of φ 's meaning in the n -ary self-composition of \mathcal{M} .

1 Function *SemanticsBoundedUntil*($\mathcal{M}, \varphi = \mathbb{P}(\varphi_1 \mathcal{U}^{[k_1, k_2]} \varphi_2), n$)

2 **if** $k_2 = 0$ **then**

3 $E := \text{Semantics}(\mathcal{M}, \varphi_1, n) \wedge \text{Semantics}(\mathcal{M}, \varphi_2, n)$;

4 **foreach** $\mathbf{s} = (s_1, \dots, s_n) \in S^n$ **do**

5 $E := E \wedge (\text{holds}_{\mathbf{s}, \varphi_2} \rightarrow \text{prob}_{\mathbf{s}, \varphi} = 1) \wedge (\neg \text{holds}_{\mathbf{s}, \varphi_2} \rightarrow \text{prob}_{\mathbf{s}, \varphi} = 0)$;

6 **else if** $k_1 = 0$ **then**

7 $E := \text{SemanticsBoundedUntil}(\mathcal{M}, \mathbb{P}(\varphi_1 \mathcal{U}^{[0, k_2-1]} \varphi_2), n)$;

8 **foreach** $\mathbf{s} = (s_1, \dots, s_n) \in S^n$ **do**

9 $E := E \wedge (\text{holds}_{\mathbf{s}, \varphi_2} \rightarrow \text{prob}_{\mathbf{s}, \varphi} = 1) \wedge ((\neg \text{holds}_{\mathbf{s}, \varphi_1} \wedge \neg \text{holds}_{\mathbf{s}, \varphi_2}) \rightarrow \text{prob}_{\mathbf{s}, \varphi} = 0)$;

10 **foreach** $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_n) \in Act(s_1) \times \dots \times Act(s_n)$ **do**

11 $E := E \wedge ([\text{holds}_{\mathbf{s}, \varphi_1} \wedge \neg \text{holds}_{\mathbf{s}, \varphi_2} \wedge \bigwedge_{i=1}^n \sigma_{s_i} = \alpha_i] \rightarrow [\text{prob}_{\mathbf{s}, \varphi} =$

12 $\sum_{\mathbf{s}' \in \text{supp}(\alpha_1) \times \dots \times \text{supp}(\alpha_n)} ((\prod_{i=1}^n \mathbf{P}(s_i, \alpha_i, s'_i)) \cdot \text{prob}_{\mathbf{s}', \mathbb{P}(\varphi_1 \mathcal{U}^{[0, k_2-1]} \varphi_2)})])$;

13 **else if** $k_1 > 0$ **then**

14 $E := \text{SemanticsBoundedUntil}(\mathcal{M}, \mathbb{P}(\varphi_1 \mathcal{U}^{[k_1-1, k_2-1]} \varphi_2), n)$;

15 **foreach** $\mathbf{s} = (s_1, \dots, s_n) \in S^n$ **do**

16 $E := E \wedge (\neg \text{holds}_{\mathbf{s}, \varphi_1} \rightarrow \text{prob}_{\mathbf{s}, \varphi} = 0)$;

17 **foreach** $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_n) \in Act(s_1) \times \dots \times Act(s_n)$ **do**

18 $E := E \wedge ([\text{holds}_{\mathbf{s}, \varphi_1} \wedge \bigwedge_{i=1}^n \sigma_{s_i} = \alpha_i] \rightarrow [\text{prob}_{\mathbf{s}, \varphi} =$

19 $\sum_{\mathbf{s}' \in \text{supp}(\alpha_1) \times \dots \times \text{supp}(\alpha_n)} ((\prod_{i=1}^n \mathbf{P}(s_i, \alpha_i, s'_i)) \cdot$

20 $\text{prob}_{\mathbf{s}', \mathbb{P}(\varphi_1 \mathcal{U}^{[k_1-1, k_2-1]} \varphi_2)})])$;

21 **return** E ;

The *SemanticsBoundedUntil* method, listed in Algorithm 4, encodes the probability $\mathbb{P}(\varphi_1 \mathcal{U}^{[k_1, k_2]} \varphi_2)$ of a bounded until formula in the numeric variables $\text{prob}_{\mathbf{s}, \mathbb{P}(\varphi_1 \mathcal{U}^{[k_1, k_2]} \varphi_2)}$ for all (composed) states $\mathbf{s} \in S^n$ and recursively

reduced time bounds. There are three main cases: (i) the satisfaction of $\varphi_1 \mathcal{U}^{[0, k_2-1]} \varphi_2$ requires to satisfy φ_2 immediately (lines 2–5); (ii) $\varphi_1 \mathcal{U}^{[0, k_2-1]} \varphi_2$ can be satisfied by either satisfying φ_2 immediately or satisfying it later, but in the latter case φ_1 needs to hold currently (lines 6–12); (iii) φ_1 has to hold and φ_2 needs to be satisfied some time later (lines 13–19). To avoid the repeated encoding of the semantics of the operands, we do it only when we reach case (i) where recursion stops (line 3). For the other cases, we recursively encode the probability to reach a φ_2 -state over φ_1 states where the deadlines are reduced

Algorithm 5: SMT encoding of the truth of the input formula

Input : $\mathcal{M} = (S, Act, \mathbf{P}, AP, L)$: MDP;
 $\exists \hat{\sigma}. Q_1 \hat{s}_1(\hat{\sigma}) \dots Q_n \hat{s}_n(\hat{\sigma}). \varphi^{nq}$: HyperPCTL formula.

Output: Encoding of the truth of the input formula in \mathcal{M} .

1 Function *Truth*($\mathcal{M}, \exists \hat{\sigma}. Q_1 \hat{s}_1(\hat{\sigma}) \dots Q_n \hat{s}_n(\hat{\sigma}). \varphi^{nq}$)

2 **foreach** $i = 1, \dots, n$ **do**

3 **if** $Q_i = \forall$ **then** $B_i := \bigwedge_{s_i \in S} \text{”}$;

4 **else** $B_i := \bigvee_{s_i \in S} \text{”}$;

5 **return** $B_1 \dots B_n \text{ holds}_{(s_1, \dots, s_n), \varphi^{nq}}$;

with one step (lines 7 resp. 14) and use these to fix the values of the variables $prob_{\mathbf{s}, \mathbb{P}(\varphi_1 \mathcal{U}^{[k_1, k_2]} \varphi_2)}$, similarly to the unbounded case but under additional consideration of time bounds.

Finally, the Truth method listed in Algorithm 5 encodes the meaning of the state quantification: it states for each universal quantifier that instantiating it with any MDP state should satisfy the formula (conjunction over all states in line 3), and for each existential state quantification that at least one state should lead to satisfaction (disjunction in line 4).

Theorem 3. *Algorithm 1 returns a formula that is true iff its input HyperPCTL formula is satisfied by the input MDP.*

We note that the satisfiability of the generated SMT encoding for a formula with an existential scheduler quantifier does not only prove the truth of the formula but provides also a scheduler as witness, encoded in the solution of the SMT encoding. Conversely, unsatisfiability of the SMT encoding for a formula with a universal scheduler quantifier provides a counterexample scheduler.

6 Evaluation

We developed a prototypical implementation of our algorithm in python, with the help of several libraries. There is an extensive use of STORMPY [15,16], which provides efficient solution to parsing, building, and storage of MDPs. We used the SMT-solver Z3 [17] to solve the logical encoding generated by Algorithm 1. All of our experiments were run on a MacBook Pro laptop with a 2.3GHz i7 processor with 32GB of RAM. The results are presented in Table 1.

As the first case study, we model and analyze information leakage in the modular exponentiation algorithm (function `modexp` in Fig. 4); the corresponding results in Table 1 are marked by **TA**. We experimented with 1, 2, and 3 bits for the encryption key (hence, $m \in \{2, 4, 6\}$). The specification checks whether there is a timing channel for all possible schedulers, which is the case for the implementation in `modexp`.

Our second case study is verification of password leakage thorough the string comparison algorithm (function `str_cmp` in Fig 5). Here, we also experimented with $m \in \{2, 4, 6\}$; results in Table 1 are denoted by **PW**.

In our third case study, we assume two concurrent processes. The first process decrements the value of a secret h by 1 as long as the value is still positive, and after this it sets a low variable l to 1. A second process just sets the value of the same low variable l to 2. The two threads run in parallel; as long as none of them terminated, a fair scheduler chooses for each CPU cycle the next executing thread. As discussed in Section 1, this MDP opens a probabilistic thread scheduling channel and leaks the value of h . We denote this case study by **TS** in Table 1, and compare observations for executions with different secret values h_1 and h_2 (denoted as $h = (h_1, h_2)$ in the table). There is an interesting relation between the execution times for **TA** and **TS**. For example, although the MDP for **TA** with $m = 4$ has 60 reachable states and the MDP for **TS**

Case study	Running time (s)			#SMT variables	#subformulas	#states	#transitions	
	SMT encoding	SMT solving	Total					
TA	$m = 2$	5.43	0.31	5.74	8088	50654	24	46
	$m = 4$	114	20	134	50460	368062	60	136
	$m = 6$	1721	865	2586	175728	1381118	112	274
PW	$m = 2$	5.14	0.3	8.14	8088	43432	24	46
	$m = 4$	207	40	247	68670	397852	70	146
	$m = 6$	3980	1099	5079	274540	1641200	140	302
TS	$h = (0, 1)$	0.83	0.07	0.9	1379	7913	7	13
	$h = (0, 15)$	60	1607	1667	34335	251737	35	83
	$h = (4, 8)$	11.86	17.02	28.88	12369	87097	21	48
	$h = (8, 15)$	60	1606	1666	34335	251737	35	83
PC	$s=(0)$	277	1996	2273	21220	1859004	20	158
	$s=(0,1)$	822	5808	6630	21220	5349205	20	280
	$s=(0,1,2)$	1690	58095	59785	21220	11006581	20	404

Table 1: Experimental results. **TA**: Timing attack. **PW**: Password leakage. **TS**: Thread scheduling. **PC**: Probabilistic conformance.

comparing executions for $h = (0, 15)$ has 35 reachable states, verification of **TS** takes 20 times more than **TA**. We believe this is because the MDP of **TS** is twice deeper than the MDP of **TA**, making the SMT constraints more complex.

Our last case study is on probabilistic conformance, denoted **PC**. The input is a DTMC that encodes the behavior of a 6-sided die as well as a structure of actions having probability distributions with two successor states each; these transitions can be pruned using a scheduler to obtain a DTMC which simulates the die outcomes using a fair coin. Given a fixed state space, we experiment with different numbers of transitions. In particular, we started from the implementation in [14] and then we added all the possible nondeterministic transitions from the first state to all the other states ($s=0$), from the first and second states to all the others ($s=0,1$), and from the first, second, and third states to all the others ($s=0,1,2$). Each time we were able not only to satisfy the formula, but also to obtain the witness corresponding to the scheduler satisfying the property.

Regarding the running times listed in Table 1, we note that our implementation is only prototypical and there are possibilities for numerous optimizations. Most importantly, for purely existentially or purely universally quantified formulas, we could define a more efficient encoding with much less variables. However, it is clear that the running times for even relatively small MPDs are large. This is simply because of the high complexity of the verification of hyperproperties. In addition, the HyperPCTL formulas in our case studies have multiple scheduler and/or state quantifiers, making the problem significantly more difficult.

7 Conclusion and Future Work

We investigated the problem of specifying and model checking probabilistic hyperproperties of Markov decision processes (MDPs). Our study is motivated by the fact that many systems have probabilistic nature and are influenced by

nondeterministic actions of their environment. We extended the temporal logic HyperPCTL for DTMCs [7] to the context of MDPs by allowing formulas to quantify over schedulers. This additional expressive power leads to undecidability of the HyperPCTL model checking problem on MDPs, but we also showed that the undecidable fragment becomes decidable for non-probabilistic memoryless schedulers. Indeed, all applications discussed in this paper only require this type of schedulers.

Due to the high complexity of the problem, more efficient model checking algorithms are greatly needed. An orthogonal solution is to design less accurate and/or approximate algorithms such as statistical model checking that scale better and provide certain probabilistic guarantees about the correctness of verification. Another interesting direction is using counterexample-guided techniques to manage the size of the state space.

References

1. Clarkson, M.R., Schneider, F.B.: Hyperproperties. *Journal of Computer Security* **18**(6) (2010) 1157–1210
2. Alpern, B., Schneider, F.B.: Defining liveness. *Information Processing Letters* **21** (1985) 181–185
3. Goguen, J.A., Meseguer, J.: Security policies and security models. In: *IEEE Symp. on Security and Privacy*. (1982) 11–20
4. Zdancewic, S., Myers, A.C.: Observational determinism for concurrent program security. In: *Proc. of CSFW’03*. (2003) 29
5. Wang, Y., Zarei, M., Bonakdarpour, B., Pajic, M.: Statistical verification of hyperproperties for cyber-physical systems. *ACM Transactions on Embedded Computing Systems (TECS)* **18**(5s) (2019) 92:1–92:23
6. Bonakdarpour, B., Sánchez, C., Schneider, G.: Monitoring hyperproperties by combining static analysis and runtime verification. In: *Proc. of ISoLA’18*. (2018) 8–27
7. Abraham, E., Bonakdarpour, B.: HyperPCTL: A temporal logic for probabilistic hyperproperties. In: *Proc. of QEST’18*. (2018) 20–35
8. Abraham, E., Bartocci, E., Bonakdarpour, B., Dobe, O.: Parameter synthesis for probabilistic hyperproperties. In: *Proc. of LPAR-23: the 23rd International Conference on Logic for Programming, Artificial Intelligence and Reasoning*. Volume 73 of *EPiC Series in Computing*, EasyChair (2020) 12–31
9. Wang, Y., Nalluri, S., Bonakdarpour, B., Pajic, M.: Statistical model checking for hyperproperties. In: *Proceedings of the IEEE 34th Computer Security Foundations (CSF)*. (2021) To appear.
10. Guarnieri, M., Marinovic, S., Basin, D.: Securing databases from probabilistic inference. In: *Proc. of CSF’17*. (2017) 343–359
11. Dimitrova, R., Finkbeiner, B., Torfah, H.: Probabilistic hyperproperties of markov decision processes. In: *Proceedings of the 18th Symposium on Automated Technology for Verification and Analysis (ATVA)*. (2020) To appear.
12. Baier, C., Brázdil, T., Größer, M., Kucera, A.: Stochastic game logic. *Acta Informatica* **49**(4) (2012) 203–224
13. Ngo, T.M., Stoelinga, M., Huisman, M.: Confidentiality for probabilistic multi-threaded programs and its verification. In: *Proc. of ESSoS’13*. (2013) 107–122

14. Knuth, D., Yao, A.: The complexity of nonuniform random number generation. In: Algorithms and Complexity: New Directions and Recent Results. Academic Press (1976)
15. : STORMPY. <https://moves-rwth.github.io/stormpy/>
16. Dehnert, C., Junges, S., Katoen, J., Volk, M.: A Storm is coming: A modern probabilistic model checker. In: Proc. of CAV'17. (2017) 592–600
17. de Moura, L.M., Bjørner, N.: Z3: An efficient SMT solver. In: Proc. of TACAS'08. (2008) 337–340
18. Baier, C., Bertrand, N., Größer, M.: On decision problems for probabilistic Büchi automata. In: Proc. of FOSSACS'08. (2008) 287–301
19. Bonakdarpour, B., Finkbeiner, B.: The complexity of monitoring hyperproperties. In: Proc. of CSF'18. (2018) 162–174
20. Baier, C., Katoen, J.P.: Principles of Model Checking. The MIT Press (2008)

A Proof of Theorem 1

Theorem 1. HyperPCTL model checking for MDPs is in general undecidable.

Proof. We reduce the *emptiness* problem in *probabilistic Büchi automata* (PBA), which is known to be undecidable [18], to our problem.

A.1 Probabilistic Büchi Automata

PBA can be viewed as nondeterministic Büchi automata where the nondeterminism is resolved by a probabilistic choice. That is, for any state q and letter a in alphabet Σ , either q does not have any a -successor or there is a probability distribution for the a -successors of q .

Definition 7. A probabilistic Büchi automaton (PBA) over a finite alphabet Σ is a tuple $\mathcal{P} = (Q, \delta, \Sigma, F)$, where Q is a finite state space, $\delta : Q \times \Sigma \times Q \rightarrow [0, 1]$ is the transition probability function, such that for all $q \in Q$ and $a \in \Sigma$:

$$\sum_{q' \in Q} \delta(q, a, q') \in \{0, 1\}$$

and $F \subseteq Q$ is the set of accepting states.

A *run* for an infinite word $w = a_1 a_2 \dots \in \Sigma^\omega$ is an infinite sequence $\pi = q_0 q_1 q_2 \dots$ of states in Q , such that $q_{i+1} \in \delta(q_i, a_{i+1}) = \{q' \mid \delta(q_i, a_{i+1}, q') > 0\}$ for all $i \geq 0$. Let $\text{Inf}(\pi)$ denote the set of states that are visited infinitely often in π . Run π is called *accepting* if $\text{Inf}(\pi) \cap F \neq \emptyset$. Given an infinite input word $w \in \Sigma^\omega$, the behavior of \mathcal{P} is given by the infinite Markov chain that is obtained by unfolding \mathcal{P} into a tree using w . This is similar to an induced Markov chain from an MDP by a scheduler. Hence, standard concepts for Markov chains can be applied to define the acceptance probability of w in \mathcal{P} , denoted by $\text{Pr}_{\mathcal{P}}(w)$ or briefly $\text{Pr}(w)$, by the probability measure of the set of accepting runs for w in \mathcal{P} . We define the accepted language of \mathcal{P} as:

$$\mathcal{L}(\mathcal{P}) = \{w \in \Sigma^\omega \mid \text{Pr}_{\mathcal{P}}(w) > 0\}.$$

The *emptiness problem* is to decide whether or not $\mathcal{L}(\mathcal{P}) = \emptyset$ for a given input \mathcal{P} .

A.2 Mapping

Our idea of mapping the emptiness problem in PBA to HyperPCTL model checking for MDPs is as follows. We map a PBA to an MDP such that the words of the PBA are mimicked by the runs of the MDP. In other words, letters of the words in the PBA appear as propositions on states of the MDP. This way, the existence of a word in the language of the PBA corresponds to the existence of a scheduler that produces a satisfying computation tree in the induced Markov

chain of the MDP.

MDP $\mathcal{M} = (S, Act, \mathbf{P}, AP, L)$: Let $\mathcal{P} = (Q, \delta, \Sigma, F)$ be a PBA with alphabet Σ . We obtain an MDP $\mathcal{M} = (S, Act, \mathbf{P}, AP, L)$ as follows:

- The set of states is $S = Q \times \Sigma$.
- The set of actions is $Act = \Sigma$.
- The transition probability function \mathbf{P} is defined as follows:

$$\mathbf{P}\left((q, a), b, (q', a')\right) = \begin{cases} \delta(q, b, q') & \text{if } a' = b \\ 0 & \text{otherwise} \end{cases}$$

- The set of atomic propositions is $AP = \Sigma \cup \{f\}$, where $f \notin \Sigma$ (we use f to label the accepting states).
- The labeling function L is defined as follows. For each $a \in \Sigma$ and $q \in Q$, we have:

$$L(q, a) = \begin{cases} \{a, f\} & \text{if } q \in F \\ \{a\} & \text{otherwise} \end{cases}$$

HyperPCTL formula: The HyperPCTL formula in our mapping is the following:

$$\varphi_{\text{map}} = \exists \hat{\sigma}. \exists \hat{s}(\hat{\sigma}). \forall \hat{s}'(\hat{\sigma}). \left(\mathbb{P}\left(\square \bigwedge_{a \in AP \setminus \{f\}} (a_{\hat{s}} \leftrightarrow a_{\hat{s}'})\right) = 1 \right) \wedge \left(\mathbb{P}\left(\diamond \mathbb{P}\left(\square \mathbb{P}\left(\diamond f_{\hat{s}}\right) = 1\right) = 1\right) > 0 \right)$$

Intuitively, the formula establishes connection between the PBA emptiness problem and HyperPCTL model checking MDPs. In particular:

- The existence of scheduler $\hat{\sigma}$ in φ_{map} corresponds to the existence of a word w in $\mathcal{L}(\mathcal{P})$;
- the state quantifiers and the left conjunct ensure that the path in the induced Markov chain and the PBA follow the sequence of actions (respectively, letters) in the witness to $\hat{\sigma}$ (respectively, w), and
- the right conjunct mimics that a state in F is visited with non-zero probability if and only if a state labeled by proposition f is visited infinitely often in the MDP with non-zero probability.

A.3 Reduction

We now show that $\mathcal{L}(\mathcal{P}) \neq \emptyset$ if and only if $\mathcal{M} \models \varphi_{\text{map}}$. We distinguish two cases:

- (\Rightarrow) Suppose we have $\mathcal{L}(\mathcal{P}) \neq \emptyset$. This means there exists a word $w \in \Sigma^\omega$, such that $\Pr_{\mathcal{P}}(w) > 0$. We use w to eliminate the existential scheduler quantifier and instantiate $\hat{\sigma}$ in formula φ_{map} . This induces a DTMC and now, we

show that the induced DTMC satisfies the following HyperPCTL formula as prescribed in [7]:

$$\exists \hat{s}. \forall \hat{s}'. \left(\mathbb{P} \left(\Box \bigwedge_{a \in \text{AP} \setminus \{f\}} (a_{\hat{s}} \leftrightarrow a_{\hat{s}'}) \right) = 1 \right) \wedge \left(\mathbb{P} \left(\Diamond \mathbb{P} \left(\Box \mathbb{P} \left(\Diamond f_{\hat{s}} \right) = 1 \right) = 1 \right) > 0 \right)$$

To this end, observe that the right conjunct is trivially satisfied due to the fact that $\Pr_{\mathcal{P}}(w) > 0$. That is, since a state in F is visited infinitely often with non-zero probability in \mathcal{P} , a state labeled by f in \mathcal{M} is also visited infinitely often with non-zero probability. The left conjunct is also satisfied by construction of the mapped MDP, since the sequence of letters in w appear in all paths of the induced DTMC as propositions.

- (\Leftarrow) The reverse direction is pretty similar. Since the answer to the model checking problem is affirmative, a witness to scheduler quantifier $\hat{\sigma}$ exists. This scheduler induces a DTMC whose paths follow the same sequence of propositions. This sequence indeed provides us with the word w for \mathcal{P} . Finally, since the right conjunct in φ_{map} is satisfied by the MDP, we are guaranteed that w reaches an accepting state in F infinitely often with non-zero probability.

And this concludes the proof. ■

B Proof of Theorem 2

Theorem. 2 The problem to decide for MDPs the truth of HyperPCTL formulas with a single existential (respectively, universal) scheduler quantifier over non-probabilistic memoryless schedulers is NP-complete (respectively, coNP-complete) in the state set size of the given MDP.

Proof. In order to show membership to NP, let \mathcal{M} be an MDP and $\varphi = \exists \hat{\sigma}. \varphi'$ be a HyperPCTL formula. We show that given a solution to the problem, we can verify the solution in polynomial time. Observe that given a non-probabilistic memoryless scheduler as a witness to the existential quantifier $\exists \hat{\sigma}$, one can compute the induced DTMC and then verify the DTMC against the resulting HyperPCTL formula in polynomial time in the size of the induced DTMC [7].

Inspired by the proof technique introduced in [19], for the lower bound, we reduce the SAT problem to our model checking problem.

B.1 The Satisfiability Problem

The SAT problem is as follows:

Let $\{x_1, x_2, \dots, x_n\}$ be a set of propositional variables. Given is a Boolean formula $y = y_1 \wedge y_2 \wedge \dots \wedge y_m$, where each y_j , for $j \in [1, m]$, is a disjunction of at least three literals. Is y satisfiable? That is, does there exist an assignment of truth values to x_1, x_2, \dots, x_n , such that y evaluates to true?

B.2 Mapping

We now present a mapping from an arbitrary instance of SAT to the model checking problem of an MDP and a HyperPCTL formula of the form $\exists\hat{\sigma}.\exists\hat{s}(\hat{\sigma}).\forall\hat{s}'(\hat{\sigma}).\psi$. Then, we show that the MDP satisfies this formula if and only if the answer to the SAT problem is affirmative. Figure 6 shows an example.

MDP $\mathcal{M} = (S, Act, \mathbf{P}, AP, L)$:

- (*Atomic propositions AP*) We include four atomic propositions: p and \bar{p} to mark the positive and negative literals in each clause and c and \bar{c} to mark paths that correspond to clauses of the SAT formula. Thus,

$$AP = \{p, \bar{p}, c, \bar{c}\}.$$

- (*Set of states S*) We now identify the members of S :
 - For each clause y_j , where $j \in [1, m]$, we include a state r_j , labeled by proposition c . We also include a state r_0 labeled by \bar{c}
 - For each clause y_j , where $j \in [1, m]$, we introduce the following n states:

$$\{v_i^j \mid i \in [1, n]\}.$$

Each state v_i^j is labeled with proposition p if x_i is a literal in y_j , or with \bar{p} if $\neg x_i$ is a literal in y_j .

- For each Boolean variable x_i , where $i \in [1, n]$, we include two states s_i and \bar{s}_i . Each state s_i (respectively, \bar{s}_i) is labeled by p (respectively, \bar{p}).
- (*Set of actions Act*) The set of actions is $Act = \{\alpha, \beta, \gamma\}$. Intuitively, the scheduler chooses action α (respectively, β) at a state s_i or \bar{s}_i to assign true (respectively, false) to variable x_{i+1} . Action γ is the sole action available at all other states.
- (*Transition probability function P*) We now identify the members of P . All transitions have probability 1, so we only discuss the actions.
 - We add transitions (r_j, γ, v_1^j) for each $j \in [1, m]$, where from r_j , the probability of reaching v_1^j is 1.
 - For each $i \in [1, n]$, we include four transitions (s_i, α, s_{i+1}) , $(s_i, \beta, \bar{s}_{i+1})$, $(\bar{s}_i, \alpha, s_{i+1})$, and $(\bar{s}_i, \beta, \bar{s}_{i+1})$. The intuition here is that when the scheduler chooses action α at state s_i or \bar{s}_i , variable x_{i+1} evaluates to true and when the scheduler chooses action β at state s_i or \bar{s}_i , variable x_{i+1} evaluates to false in the SAT instance. We also include two transitions (r_0, α, s_1) and (r_0, β, \bar{s}_1) with the same intended meaning.
 - Finally, we include self-loops (s_n, γ, s_n) , $(\bar{s}_n, \gamma, \bar{s}_n)$, and (v_n^j, γ, v_n^j) , for each $j \in [1, m]$.

HyperPCTL formula: The HyperPCTL formula in our mapping is the following:

$$\varphi_{\text{map}} = \exists\hat{\sigma}.\exists\hat{s}(\hat{\sigma}).\forall\hat{s}'(\hat{\sigma}).\bar{c}_{\hat{s}(\hat{\sigma})} \wedge \left(c_{\hat{s}'(\hat{\sigma})} \Rightarrow \right.$$

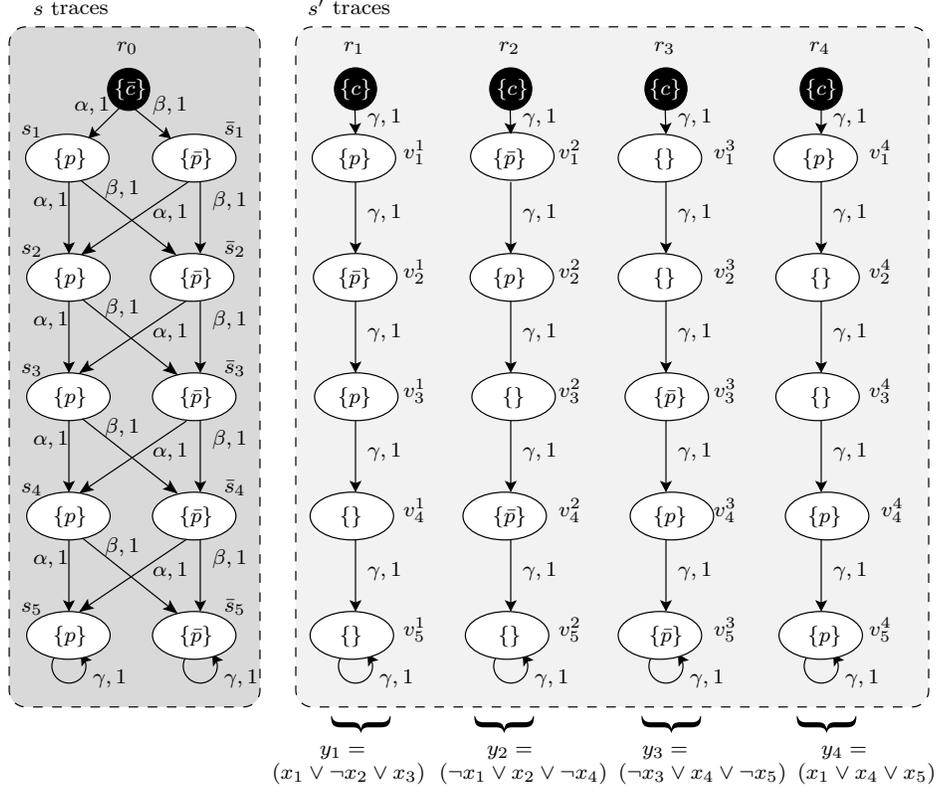


Fig. 6: Example of mapping SAT to HyperPCTL model checking.

$$\mathbb{P}\left(\diamond\left(\left(p_{\hat{s}(\hat{\sigma})} \wedge p_{\hat{s}'(\hat{\sigma})}\right) \vee \left(\bar{p}_{\hat{s}(\hat{\sigma})} \wedge \bar{p}_{\hat{s}'(\hat{\sigma})}\right)\right)\right) = 1$$

The intended meaning of the formula is that if there exists a scheduler that makes the formula true by choosing the α and β actions, this scheduler gives us the assignment to the Boolean variables in the SAT instance. This is achieved by making all clauses true, hence, the $\forall \hat{s}'(\hat{\sigma})$ subformula.

B.3 Reduction

We now show that the given SAT formula is satisfiable if and only if the MDP obtained by our mapping satisfies the HyperPCTL formula φ_{map} .

(\Rightarrow) Suppose that y is satisfiable. Then, there is an assignment that makes each clause y_j , where $j \in [1, m]$, true. We now use this assignment to instantiate a scheduler for the formula φ_{map} . If $x_i = \mathbf{true}$, then we instantiate scheduler $\hat{\sigma}$ such that in state s_{i-1} or \bar{s}_{i-1} , it chooses action α . Likewise, if $x_i = \mathbf{false}$, then we instantiate scheduler $\hat{\sigma}$, such that in state s_{i-1} or \bar{s}_{i-1} , it chooses

action β . We now show that this scheduler instantiation evaluates formula φ_{map} to true. First observe that $\hat{s}(\hat{\sigma})$ can only be instantiated with state r_0 and $\hat{s}'(\hat{\sigma})$ can only be instantiated with states r_j , where $j \in [1, m]$. Otherwise, the left side of the implication in φ_{map} becomes false, making the formula vacuously true. Since each y_j is true, there is at least one literal in y_j that is true. If this literal is of the form x_i , then we have $x_i = \text{true}$ and the path that starts from r_0 will include s_i , which is labeled by p . Hence, the values of p , in both paths that start from $\hat{s}(\hat{\sigma})$ and $\hat{s}'(\hat{\sigma})$ are eventually equal. If the literal in y_j is of the form $\neg x_i$, then $x_i = \text{false}$ and the path that starts from $\hat{s}(\hat{\sigma})$ will include \bar{s}_i . Again, the values of \bar{p} are eventually equal. Finally, since all clauses are true, all paths that start from $\hat{s}'(\hat{\sigma})$ reach a state where the right side of the implication becomes true.

(\Leftarrow) Suppose our mapped MDP satisfies formula φ_{map} . This means that there exists a scheduler and state $\hat{s}(\hat{\sigma})$ that makes the subformula $\forall \hat{s}'(\hat{\sigma})$ true, i.e., the path that starts from r_0 results in making the inner PCTL formula true for all paths that start from r_j . We obtain the truth assignment to the SAT problem as follows. If the scheduler chooses action α to state s_i , then we assign $x_i = \text{true}$. Likewise, if the scheduler chooses action β to state \bar{s}_i , then we assign $x_i = \text{false}$. Observe that since in no state p and \bar{p} are simultaneously true and no path includes both s_i and \bar{s}_i , variable x_i will have only one truth value. Similar to the forward direction, it is straightforward to see that this valuation makes every clause y_j of the SAT instance true.

And this concludes the proof. \blacksquare

C Proof of Theorem 3

Theorem. 3 Algorithm 1 returns a formula that is true iff its input HyperPCTL formula is satisfied by the input MDP.

Proof. The proof is by structural induction over the formula type. The cases for constants, atomic propositions, Boolean combinations and arithmetic expressions are straightforward. The remaining cases are the probabilistic temporal operators.

For the next-operator $\mathbb{P}(\bigcirc \varphi')$ (line 13 in Algorithm 2), we first encode the meaning of the operand φ' (line 14); by induction assumption this encoding is sound. Since the operand is a state formula, its value is Boolean. The probability that the operand is true after one step is the sum of the probabilities to get to a state where the operand is true; we express this in pseudo-arithmetic by setting for each composed state \mathbf{s} the value of $\text{prob}_{\mathbf{s}, \varphi'}$ to 1 if φ' holds there and to 0 otherwise. Using these pseudo-Boolean values, we express for each composed state $\mathbf{s} = (s_1, \dots, s_n)$ and each scheduler $\boldsymbol{\sigma} = (\sigma_1, \dots, \sigma_n)$ the probability that φ' holds after one step by summing up for each possible successor $\mathbf{s}' = (s'_1, \dots, s'_n)$ the probability $\prod_{i=1}^n \mathbf{P}(s_i, \alpha_i, s'_i)$ to get there in one step times the pseudo-Boolean value of $\text{prob}_{\mathbf{s}', \varphi'}$ (lines 16–19).

The semantics of unbounded until $\mathbb{P}(\varphi_1 \mathcal{U} \varphi_2)$ is encoded in Algorithm 3. Similarly to the next-operator, we recursively encode the truth of the Boolean-valued operands (line 2). The remaining encoding follows for each composed state \mathbf{s} and each scheduler $\sigma = (\sigma_1, \dots, \sigma_n)$ the standard fixedpoint-encoding of the probability to satisfy the until formula in the induced DTMC [20]. This probability is 1 for all states satisfying φ_2 and 0 for all states that do not satisfy any of the operands (line 4). Furthermore, the probability is also 0 for all states from which no φ_2 -states are reachable; we assure this by requiring for all positive probabilities the existence of a finite loop-free path to a φ_2 -state using decreasing sequences of the arithmetic variables $d_{\mathbf{s}, \varphi_2}$ (line 8). For all other cases, the encoding is similar to the next-operator, summing up for all possible successor states the probabilities to get there times the probability to satisfy the until formula along paths starting from there (line 7).

The case for unbounded until is based on a technically rather complex case distinction, but it is just a direct encoding the semantics of bounded until.