

Model-Checking Markov Chains in the Presence of Uncertainties

Koushik Sen, Mahesh Viswanathan, and Gul Agha

Department of Computer Science,
University of Illinois at Urbana-Champaign
{ksen, vmahesh, agha}@uiuc.edu

Abstract. We investigate the problem of model checking Interval-valued Discrete-time Markov Chains (IDTMC). IDTMCs are discrete-time finite Markov Chains for which the exact transition probabilities are not known. Instead in IDTMCs, each transition is associated with an interval in which the actual transition probability must lie. We consider two semantic interpretations for the uncertainty in the transition probabilities of an IDTMC. In the first interpretation, we think of an IDTMC as representing a (possibly uncountable) family of (classical) discrete-time Markov Chains, where each member of the family is a Markov Chain whose transition probabilities lie within the interval range given in the IDTMC. This semantic interpretation we call Uncertain Markov Chains (UMC). In the second semantics for an IDTMC, which we call Interval Markov Decision Process (IMDP), we view the uncertainty as being resolved through non-determinism. In other words, each time a state is visited, we adversarially pick a transition distribution that respects the interval constraints, and take a probabilistic step according to the chosen distribution. We show that the PCTL model checking problem for both Uncertain Markov Chain semantics and Interval Markov Decision Process semantics is decidable in PSPACE. We also prove lower bounds for these model checking problems.

1 Introduction

Discrete time stochastic models such as *Discrete Time Markov Chains* (DTMCs) have been used to analyze the correctness, reliability, and performance of systems [8, 11, 19, 13]. In a DTMC, the system is assumed to have finitely many states, and the system's future behavior is completely determined by its current state. From each state of the system, the probability of transitioning to any other given state at the next step is fixed and is given by the transition probability matrix of the DTMC.

The assumption that the system makes transitions according to a fixed distribution at each step and that this distribution is precisely known when modeling, is a strong assumption that may often not hold in practice [12, 15, 26, 14]. If the system being modeled is an open system, i.e., interacts with an environment, then uncertainty in the transitions may arise due to imperfect information about the environment. For example, consider a system that interacts with an imperfect communication medium that may lose messages. The probability of message loss may either depend on choice of the communication medium or on a complicated, time-varying dependence on events that are not

precisely understood at the time of modeling the system. Another source of imprecision is that the transition probabilities in the system model are often estimated through statistical experiments, which only provide bounds on the transition probabilities.

In order to faithfully capture these system uncertainties in stochastic models, the model of *Interval-valued Discrete-time Markov Chains (IDTMC)* has been introduced [12, 14]. These are DTMC models where the exact probability of taking a state transition is not known, and instead the transition probability is assumed to lie within a range associated with the transition. Two semantic interpretations have been suggested for such models. *Uncertain Markov Chains (UMC)* [12] is an interpretation of an IDTMC as a family of (possibly uncountably many) DTMCs, where each member of the family is a DTMC whose transition probabilities lie within the interval range given in the IDTMC. In the second interpretation, called *Interval Markov Decision Process (IMDP)*, we view the uncertainty as being resolved through non-determinism. In other words, each time a state is visited, we adversarially pick a transition distribution that respects the interval constraints, and take a probabilistic step according to the chosen distribution. Thus, IMDPs allow the possibility of modeling a non-deterministic choice made from a set of (possibly) uncountably many choices. An IMDP can be seen as a generalization of Markov Decision Processes (MDPs) [17, 3, 21].

We investigate the problem of model checking PCTL specifications for IDTMC. The two semantic interpretations of IDTMCs yield very different model checking results (whenever the property has at least two probabilistic operators, not necessarily nested; see example in Figure 1) and require different algorithmic techniques. For the case of UMCs, we show that PCTL model checking problem can be reduced to finding feasible solutions to inequality constraints, much like in the case of DTMC and MDP [8, 4, 3, 19, 7]. However, there is one important difference. The constraints to be solved in the case of UMCs are polynomial and not just linear (as for DTMCs and MDPs). Since the *existential theory of reals* is decidable in PSPACE [18, 6], the feasibility of the polynomial constraints arising in model checking, can be determined by making a “query” to the existential theory of reals. Thus, the PCTL model checking problem for UMCs is in PSPACE. In practice, however, this algorithm may not be the most efficient. The constraints we obtain during model checking all take a special form: the polynomials are *bilinear*¹. Therefore, it might be more efficient to instead use algorithms for solving *bilinear matrix inequalities (BMIs)* [10, 9] or tools developed for this purpose [16]. Checking feasibility of BMIs is known to be NP-hard [24], but the exact complexity, which is lower than PSPACE, is unknown. On the other hand, in the case of IMDPs, we show that the model checking problem can be reduced to model checking an MDP of exponential size. We then use known results for MDPs to show that IMDPs can be model checked in PSPACE. We also present an iterative model checking algorithm for IMDPs which may prove to be more efficient in practice.

In addition to demonstrating the decidability of the model checking problem, we also prove lower bounds on the complexity of the model checking problem. We show that the model checking problem for UMCs is NP-hard and co-NP-hard; thus, for UMCs the problem is unlikely to be in P. A straightforward corollary of our results is that solving

¹ The highest power of any variable in the polynomial is 1, and any term is the product of at most two variables.

BMIs is also co-NP-hard. For IMDPs, we can only show P-hardness; in fact, even this is a consequence of the P-hardness of (classical) DTMC model checking.

The rest of the paper is organized as follows. We briefly discuss related work next. In Section 2 we formally define IDTMC and give its semantics as UMC and IMDP. PCTL and the model checking problem is introduced in Section 3. We then revisit the model checking algorithm for DTMC (Section 4) and present a modified version of the classical algorithm. The ideas in the section play a key role in our UMC model checking algorithm. Section 5 (UMC) and Section 6 (IMDP) contain our main results about the model checking problem, providing both upper and lower bounds. Finally we present our conclusions in Section 7. Due to lack of space, we do not present any proofs here; all proofs including motivating examples of UMCs and IMDPs can be found in [23].

Related Work. The model of IDTMCs has been introduced independently by Jonsson and Larsen [12] and Kozine and Utkin [14] under the names *interval specification systems* and *interval-valued finite Markov chains*, respectively. However, they consider different semantic interpretations. Jonsson and Larsen consider the UMC interpretation and study bisimulation and simulation preorders for such an interpretation. Kozine and Utkin, on the other hand, take the IMDP interpretation and present algorithms to compute the probability distribution on the states after t steps. Neither of these papers investigate the PCTL model checking problem which is the focus of this paper. We introduce new names to emphasize the subtle semantic difference in the two interpretations. A more general model called *generalized Markov processes* for describing infinite families of Markov Chains was introduced in [1]. In that paper, they showed that model checking such models with respect to PCTL* (a more general logic than PCTL) is decidable and has elementary complexity. PCTL model checking for classical DTMC and MDP models has been considered in [8, 4, 3, 19, 7].

2 Formal Models

Definition 1. A discrete-time Markov chain (DTMC) is a 4-tuple $\mathcal{M} = (S, s_I, \mathbf{P}, L)$, where

1. S is a finite set of states,
2. $s_I \in S$ is the initial state,
3. $\mathbf{P}: S \times S \rightarrow [0, 1]$ is a transition probability matrix, such that $\sum_{s' \in S} \mathbf{P}(s, s') = 1$, and
4. $L: S \rightarrow 2^{\text{AP}}$ is a labeling function that maps states to sets of atomic propositions from a set AP.

A non-empty sequence $\pi = s_0 s_1 s_2 \dots$ is called a *path* of \mathcal{M} , if each $s_i \in S$ and $\mathbf{P}(s_i, s_{i+1}) > 0$ for all $i \geq 0$. We denote the i^{th} state in a path π by $\pi[i] = s_i$. We let $\text{Path}(s)$ be the set of paths starting at state s . A probability measure on paths is induced by the matrix \mathbf{P} as follows.

Let $s_0, s_1, \dots, s_k \in S$ with $\mathbf{P}(s_i, s_{i+1}) > 0$ for all $0 \leq i < k$. Then $C(s_0 s_1 \dots s_k)$ denotes a *cylinder set* consisting of all paths $\pi \in \text{Path}(s_0)$ such that $\pi[i] = s_i$ (for $0 \leq i \leq k$). Let \mathcal{B} be the smallest σ -algebra on $\text{Path}(s_0)$ which contains all the cylinders $C(s_0 s_1 \dots s_k)$. The measure μ on cylinder sets can be defined as follows

$$\mu(C(s_0 s_1 \dots s_k)) = \begin{cases} 1 & \text{if } k = 0 \\ \mathbf{P}(s_0, s_1) \cdots \mathbf{P}(s_{k-1}, s_k) & \text{otherwise} \end{cases}$$

The *probability measure* on \mathcal{B} is then defined as the unique measure that agrees with μ (as defined above) on the cylinder sets.

Definition 2. An Interval-valued Discrete-time Markov chain (IDTMC) is a 5-tuple $\mathcal{I} = (S, s_I, \check{\mathbf{P}}, \hat{\mathbf{P}}, L)$, where

1. S is a finite set of states,
2. $s_I \in S$ is the initial state,
3. $\check{\mathbf{P}}: S \times S \rightarrow [0, 1]$ is a transition probability matrix, where each $\check{\mathbf{P}}(s, s')$ gives the lower bound of the transition probability from the state s to the state s' ,
4. $\hat{\mathbf{P}}: S \times S \rightarrow [0, 1]$ is a transition probability matrix, where each $\hat{\mathbf{P}}(s, s')$ gives the upper bound of the transition probability from the state s to the state s' ,
5. $L: S \rightarrow 2^{\text{AP}}$ is a labeling function that maps states to sets of atomic propositions from a set AP.

We consider two semantics interpretations of an IDTMC model, namely Uncertain Markov Chains (UMC) and Interval Markov Decision Processes (IMDP).

Uncertain Markov Chains. An IDTMC \mathcal{I} may represent an infinite set of DTMCs, denoted by $[\mathcal{I}]$, where for each DTMC $(S, s_I, \mathbf{P}, L) \in [\mathcal{I}]$ the following is true,

- $\check{\mathbf{P}}(s, s') \leq \mathbf{P}(s, s') \leq \hat{\mathbf{P}}(s, s')$ for all pairs of states s and s' in S

In the Uncertain Markov Chains semantics, or simply, in the UMCs, we assume that the external environment non-deterministically picks a DTMC from the set $[\mathcal{I}]$ at the beginning and then all the transitions take place according to the chosen DTMC. Note that in this semantics, the external environment makes only one non-deterministic choice. Henceforth, we will use the term UMC to denote an IDTMC interpreted according to the Uncertain Markov Chains semantics.

Interval Markov Decision Processes. In the Interval Markov Decision Processes semantics, or simply, in the IMDPs, we assume that before every transition the external environment non-deterministically picks a DTMC from the set $[\mathcal{I}]$ and then takes a one-step transition according to the probability distribution of the chosen DTMC. Note that in this semantics, the external environment makes a non-deterministic choice before every transition. Henceforth, we will use the term IMDP to denote an IDTMC interpreted according to the Interval Markov Decision Processes semantics. We now formally define this semantics.

Let $\text{Steps}(s)$ be the set of probability density functions over S defined as follows:

$$\text{Steps}(s) = \{ \mu: S \rightarrow \mathbb{R}^{\geq 0} \mid \sum_{s' \in S} \mu(s') = 1 \text{ and } \check{\mathbf{P}}(s, s') \leq \mu(s') \leq \hat{\mathbf{P}}(s, s') \text{ for all } s' \in S \}$$

In an IMDP, at every state $s \in S$, a probability density function μ is chosen non-deterministically from the set $\text{Steps}(s)$. A successor state s' is then chosen according to the probability distribution μ over S .

A path π in an IMDP $\mathcal{I} = (S, s_I, \check{\mathbf{P}}, \hat{\mathbf{P}}, L)$ is a non-empty sequence of the form $s_0 \xrightarrow{\mu_1^1} s_1 \xrightarrow{\mu_2^2} \dots$, where $s_i \in S$, $\mu_{i+1} \in \text{Steps}(s_i)$, and $\mu_{i+1}(s_{i+1}) > 0$ for all $i \geq 0$. A path can be either finite or infinite. We use π_{fin} to denote a finite path. Let $\text{last}(\pi_{\text{fin}})$ be the last state in the finite path π_{fin} . As in DTMC, we denote the i^{th} state in a path π by $\pi[i] = s_i$. We let $\text{Path}(s)$ and $\text{Path}_{\text{fin}}(s)$ be the set of all infinite and finite paths, respectively, starting at state s . To associate a probability measure with the paths, we resolve the non-deterministic choices by an *adversary*, which is defined as follows:

Definition 3. An adversary A of an IMDP \mathcal{I} is a function mapping every finite path π_{fin} of \mathcal{I} onto an element of the set $\text{Steps}(\text{last}(\pi_{\text{fin}}))$. Let $\mathcal{A}_{\mathcal{I}}$ denote the set of all possible adversaries of the IMDP \mathcal{I} . Let $\text{Path}^A(s)$ denote the subset of $\text{Path}(s)$ which corresponds to A .

The behavior of an IMDP $\mathcal{I} = (S, s_I, \check{\mathbf{P}}, \hat{\mathbf{P}}, L)$ under a given adversary A is purely deterministic. The behavior of a IMDP \mathcal{I} from a state s can be described by an infinite-state DTMC $\mathcal{M}^A = (S^A, s_I^A, \mathbf{P}^A, L^A)$ where

- $S^A = \text{Path}_{\text{fin}}(s)$,
- $s_I^A = s$, and
- $\mathbf{P}^A(\pi_{\text{fin}}, \pi'_{\text{fin}}) = \begin{cases} A(\pi_{\text{fin}})(s') & \text{if } \pi'_{\text{fin}} \text{ is of the form } \pi_{\text{fin}} \xrightarrow{A(\pi_{\text{fin}})} s' \\ 0 & \text{otherwise} \end{cases}$

There is a one-to-one correspondence between the paths of \mathcal{M}^A and $\text{Path}^A(s)$ of \mathcal{I} . Therefore, we can define a probability measure Prob_s^A over the set of paths $\text{Path}^A(s)$ using the probability measure of the DTMC \mathcal{M}^A .

3 Probabilistic Computation Tree Logic (PCTL)

In this paper we consider a sub-logic of PCTL that excludes the steady-state probabilistic operators. The formal syntax and semantics of this logic is as follows.

PCTL Syntax

$$\begin{aligned} \phi &::= \text{true} \mid a \mid \neg\phi \mid \phi \wedge \phi \mid \mathcal{P}_{\triangleright p}(\psi) \\ \psi &::= \phi \mathcal{U} \phi \mid \mathbf{X}\phi \end{aligned}$$


where $a \in \text{AP}$ is an atomic propositions, $\triangleright \in \{<, \leq, >, \geq\}$, $p \in [0, 1]$, and $k \in \mathbb{N}$. Here ϕ represents a *state* formula and ψ represents a *path* formula.

PCTL Semantics for DTMC

The notion that a state s (or a path π) *satisfies* a formula ϕ in a DTMC \mathcal{M} is denoted by $s \models_{\mathcal{M}} \phi$ (or $\pi \models_{\mathcal{M}} \phi$), and is defined inductively as follows:

$$\begin{aligned} s &\models_{\mathcal{M}} \text{true} \\ s &\models_{\mathcal{M}} a && \text{iff } a \in L(s) \\ s &\models_{\mathcal{M}} \neg\phi && \text{iff } s \not\models_{\mathcal{M}} \phi \\ s &\models_{\mathcal{M}} \phi_1 \wedge \phi_2 && \text{iff } s \models_{\mathcal{M}} \phi_1 \text{ and } s \models_{\mathcal{M}} \phi_2 \\ s &\models_{\mathcal{M}} \mathcal{P}_{\triangleright p}(\psi) && \text{iff } \text{Prob}\{\pi \in \text{Path}(s) \mid \pi \models_{\mathcal{M}} \psi\} \triangleright p \\ \pi &\models_{\mathcal{M}} \mathbf{X}\phi && \text{iff } \pi[1] \models_{\mathcal{M}} \phi \\ \pi &\models_{\mathcal{M}} \phi_1 \mathcal{U} \phi_2 && \text{iff } \exists i \geq 0 (\pi[i] \models_{\mathcal{M}} \phi_2 \text{ and } \forall j < i. \pi[j] \models_{\mathcal{M}} \phi_1) \end{aligned}$$

$$\begin{aligned}
 0 \leq p \leq 1 & \quad L(s_1) = \{\} \\
 0 \leq q \leq 1 & \quad L(s_2) = \{a\} \\
 & \quad L(s_3) = \{b\}
 \end{aligned}$$



$$\phi = P_{>0.4}(Xa) \vee P_{>0.4}(Xb)$$

Fig. 1. Example IDTMC and PCTL formula ϕ . The UMC interpretation of the IDTMC satisfies ϕ , whereas the IMDP interpretation of the IDTMC violates ϕ .

$$\begin{aligned}
 s &\models true \\
 s &\models a && \text{iff } a \in AP(s) \\
 s &\models \neg\phi && \text{iff } s \not\models \phi \\
 s &\models \phi_1 \wedge \phi_2 && \text{iff } s \models \phi_1 \text{ and } s \models \phi_2 \\
 s &\models \mathcal{P}_{\bowtie p}(\psi) && \text{iff } Prob_s^A(\{\pi \in Path^A(s) \mid \pi \models \psi\}) \bowtie p \\
 &&& \text{for all } A \in \mathcal{A} \\
 \pi &\models \mathbf{X}\phi && \text{iff } \pi[1] \models \phi \\
 \pi &\models \phi_1 \mathcal{U} \phi_2 && \text{iff } \exists i \geq 0 (\pi[i] \models \phi_2 \text{ and } \forall j < i. \pi[j] \models \phi_1)
 \end{aligned}$$

Fig. 2. PCTL semantics for IMDP

It can be shown that for any path formula ψ and any state s , the set $\{\pi \in Path(s) \mid \pi \models_{\mathcal{M}} \psi\}$ is measurable [25]. A formula $\mathcal{P}_{\bowtie p}(\psi)$ is satisfied by a state s if $Prob[\text{path starting at } s \text{ satisfies } \psi] \bowtie p$. The path formula $\mathbf{X}\phi$ holds over a path if ϕ holds at the second state on the path. The formula $\phi_1 \mathcal{U} \phi_2$ is true over a path π if ϕ_2 holds in some state along π , and ϕ_1 holds along all prior states along π .

Given a DTMC \mathcal{M} and a PCTL state formula ϕ , $\mathcal{M} \models \phi$ iff $s_I \models_{\mathcal{M}} \phi$.

PCTL Semantics for UMC

Given a UMC \mathcal{I} and a PCTL state formula ϕ , we say $\mathcal{I} \models \phi$ iff, for all $\mathcal{M} \in [\mathcal{I}]$, $\mathcal{M} \models \phi$. Note that $\mathcal{I} \not\models \phi$ does not imply that $\mathcal{I} \models \neg\phi$. This is because if $\mathcal{I} \not\models \phi$, there may exist $\mathcal{M}, \mathcal{M}' \in [\mathcal{I}]$ such that $\mathcal{M} \models \phi$ and $\mathcal{M}' \not\models \phi$.

PCTL Semantics for IMDP

The interpretation of a state formula and a path formula of PCTL for IMDPs is the same as for DTMCs except for the state formulas of the form $\mathcal{P}_{\bowtie p}(\psi)$.

The notion that a state s (or a path π) satisfies a formula ϕ in an IMDP \mathcal{I} is denoted by $s \models \phi$ (or $\pi \models \phi$), and is defined inductively in Figure 2.

The model checking of IDTMC with respect to the two semantics can give different results. For example, consider the IDTMC in Figure 1 and the PCTL formula ϕ . The UMC semantics of this IDTMC satisfies ϕ , while the IMDP semantics violates ϕ .

4 Revisiting DTMC Model-Checking

In this section we outline the basic model checking algorithm for (classical) DTMCs. The algorithm that we outline here for DTMCs is not the most efficient (like the one presented in [8]); however the main ideas presented here will form the crux of our model checking algorithm for UMCs.

The algorithm for model checking DTMCs will reduce the problem to checking the feasibility of simultaneously satisfying a finite set of polynomial inequalities. This feasibility test can be done by checking if a first-order formula with existential quantifiers about the real numbers is true. More precisely, we need to check if a formula of the form $\exists x_1, \dots, x_n P(x_1, \dots, x_n)$ is valid over the reals, where P is a boolean function of atomic predicates of the form $f_i(x_1, \dots, x_n) \bowtie 0$, where f_i is a multivariate polynomial and $\bowtie \in \{=, \neq, \leq, \geq, <, >\}$. It is well-known that this problem can be decided in PSPACE [18, 6]².

The model checking algorithm for DTMC takes a DTMC $\mathcal{M} = (S, s_I, \mathbf{P}, L)$ and a PCTL formula ϕ as input. The output is the set $\text{Sat}(\phi) = \{s \in S \mid s \models_{\mathcal{M}} \phi\}$, i.e., the set of all states of the model that satisfy ϕ . We say $\mathcal{M} \models \phi$ iff $s_I \in \text{Sat}(\phi)$.

The algorithm works by recursively computing the set $\text{Sat}(\phi')$ for each sub-formula ϕ' of ϕ as follows.

$$\begin{aligned} \text{Sat}(true) &= S & \text{Sat}(a) &= \{s \mid a \in L(S)\} \\ \text{Sat}(\neg\phi) &= S \setminus \text{Sat}(\phi) & \text{Sat}(\phi_1 \wedge \phi_2) &= \text{Sat}(\phi_1) \cap \text{Sat}(\phi_2) \\ \text{Sat}(\mathcal{P}_{\bowtie p}(\psi)) &= \{s \in S \mid p_s(\psi) \bowtie p\} \end{aligned}$$

where $p_s(\psi) = \text{Prob}\{\pi \in \text{Path}(s) \mid \pi \models_{\mathcal{M}} \psi\}$. The computation of the set $\text{Sat}(\mathcal{P}_{\bowtie p}(\psi))$ requires the computation of $p_s(\psi)$ at every state $s \in S$.

If $\psi = \mathbf{X}\phi$, then $p_s(\psi) = \sum_{s' \in \text{Sat}(\phi)} \mathbf{P}(s, s')$.

To compute $p_s(\phi_1 \mathcal{U} \phi_2)$, we first split the set of states S into three disjoint subsets, S^{no} , S^{yes} , and $S^?$ where $S^{\text{no}} = \text{Sat}(\neg\phi_1 \wedge \neg\phi_2)$, $S^{\text{yes}} = \text{Sat}(\phi_2)$, and $S^? = S \setminus (S^{\text{no}} \cup S^{\text{yes}})$. Moreover, let $S^{? \text{no}}$ be the set $\{s \mid p_s(\phi_1 \mathcal{U} \phi_2) = 0\} \setminus S^{\text{no}}$ and $S^{>0}$ be the set $\{s \mid p_s(\phi_1 \mathcal{U} \phi_2) > 0\}$. Note that $S = S^{>0} \cup S^{? \text{no}} \cup S^{\text{no}}$. By [8], $\{x_s = p_s(\phi_1 \mathcal{U} \phi_2) \mid s \in S\}$ is a solution of the following linear equation system.

$$x_s = \begin{cases} 0 & \text{if } s \in S^{\text{no}} \\ 1 & \text{if } s \in S^{\text{yes}} \\ \sum_{s' \in S} \mathbf{P}(s, s')x_{s'} & \text{if } s \in S^? \end{cases} \quad (1)$$

Note that the equation system (1) can have infinite number of solutions. For example, consider the formula $true \mathcal{U} a$, where a is an atomic proposition and the DTMC $\mathcal{M} = (\{s\}, s, \mathbf{P}, L)$, where $\mathbf{P}(s, s) = 1$ and $L(s) = \emptyset$. Note that $s \in S^{? \text{no}}$. The linear equation system (1) that is instantiated for computing $p_s(true \mathcal{U} a)$ for \mathcal{M} is $x_s = x_s$. The system has infinite number of solutions.

We can ensure that $\{x_s = p_s(\phi_1 \mathcal{U} \phi_2) \mid s \in S\}$ is a unique solution of a system of equations as follows. Fix a γ such that $0 < \gamma < 1$. Consider the following linear equation system.

$$x'_s = \begin{cases} 0 & \text{if } s \in S^{\text{no}} \\ 1 & \text{if } s \in S^{\text{yes}} \\ \sum_{s' \in S} \gamma \mathbf{P}(s, s')x'_{s'} & \text{if } s \in S^? \end{cases} \quad (2)$$

² If one takes the computational model to be Turing machines, then this result holds when the coefficients of the polynomials are rationals. On the other hand, if one considers a model of computation that is appropriate for real number computation, like the one proposed by Blum, Shub, and Smale [5], then the algorithm can handle even real coefficients.

Lemma 1. $x'_s > 0$ iff $s \in S^{>0}$.

Lemma 2. The system of linear equations in (2) has a unique solution.

Lemma 3. $x'_s = 0$ iff $s \in S^{?no} \cup S^{no}$.

Consider the following system of constraints.

$$x'_s = 0 \text{ iff } x_s = 0 \text{ for all } s \in S \quad (3)$$

where x'_s are variables of (2) and x_s are variables of (1).

Lemma 4. The system of linear equations in (1) and (2) has a unique solution given that the constraints in (3) hold. Moreover, for this unique solution $x_s = p_s(\phi_1 \mathcal{U} \phi_2)$, for all $s \in S$.

Note that the set of constraints (1), (2), and (3) can be written compactly as follows.

$$x_s = \begin{cases} 0 & \text{if } s \in S^{no} \\ 1 & \text{if } s \in S^{yes} \\ \sum_{s' \in S} \mathbf{P}(s, s') x_{s'} & \text{if } s \in S^? \end{cases} \quad x'_s = \begin{cases} 0 & \text{if } s \in S^{no} \\ 1 & \text{if } s \in S^{yes} \\ \sum_{s' \in S} \gamma \mathbf{P}(s, s') x'_{s'} & \text{if } s \in S^? \end{cases} \quad (4)$$

$$\delta_s > 0 \quad x_s = \delta_s x'_s$$

where for each $s \in S$, we introduce the variable δ_s , such that we can impose the constraint that $x_s = 0$ iff $x'_s = 0$. The satisfiability of the set of constraints (4) can be easily reduced to checking if a formula with existential quantifiers belongs to the theory of reals. The constructed formula is linear in the size of the DTMC.

5 Model Checking UMC

In this section, we reduce the problem of model-checking a UMC to checking the feasibility of a bilinear matrix inequality. (More details about bilinear matrix inequality can be found in [23].) In the non-trivial reduction, we introduce a number of auxiliary variables to achieve the goal. Note that a simpler PSPACE algorithm, which avoids the extra auxiliary variables by guessing their values non-deterministically, is possible and is easy to come up from our reduction. However, we believe that the following reduction is important from the perspective of implementation in practice using algorithms to solve bilinear matrix inequalities (BMIs).

Given a UMC \mathcal{I} and a PCTL state formula ϕ , our goal is to check whether $\mathcal{I} \models \phi$. In other words, for every $\mathcal{M} \in [\mathcal{I}]$, $\mathcal{M} \models \phi$. Thus, to check whether $\mathcal{I} \models \phi$, we check if there exists some $\mathcal{M} \in [\mathcal{I}]$ such that $\mathcal{M} \models \neg\phi$. If such an \mathcal{M} does not exist, we conclude that $\mathcal{I} \models \phi$. We will view the problem of discovering whether a $\mathcal{M} \in [\mathcal{I}]$ satisfies $\neg\phi$ as problem of checking the feasibility of a set of bilinear inequality constraints as follows. Each transition probability of the DTMC \mathcal{M} that we are searching for, will be a variable taking a value within the bounds. We will also have variables denoting the satisfaction (or non-satisfaction) of each subformula at each state, and variables denoting the probability of a path subformula being satisfied at each state. Inequality constraints on these variables will ensure that they all have consistent values. We now describe this construction formally.

Let us fix an UMC $\mathcal{I} = (S, s_I, \check{\mathbf{P}}, \hat{\mathbf{P}}, L)$ and a PCTL formula ϕ . Let $\mathcal{M} = (S, s_I, \mathbf{P}, L)$ be an arbitrary Markov chain in $[\mathcal{I}]$.

For every pair of states $s, s' \in S$, let the variable $p_{ss'}$ denote the transition probability from s to s' in \mathcal{M} , i.e., $p_{ss'}$ denotes $\mathbf{P}(s, s')$. Since \mathcal{M} is an arbitrary DTMC in $[\mathcal{I}]$, by the definition of UMC, the following constraints hold: For every state $s \in S$, $\sum_{s' \in S} p_{ss'} = 1$ and for every pair of states $s, s' \in S$, $\check{\mathbf{P}}(s, s') \leq p_{ss'} \leq \hat{\mathbf{P}}(s, s')$.

Given any PCTL formula ϕ , let us define the set $\text{subfS}(\phi)$ (of state sub-formulas) recursively as follows:

$$\begin{aligned} \text{subfS}(a) &= \{a\} & \text{subfS}(\neg\phi) &= \{\neg\phi\} \cup \text{subfS}(\phi) \\ \text{subfS}(\phi_1 \wedge \phi_2) &= \{\phi_1 \wedge \phi_2\} \cup \text{subfS}(\phi_1) \cup \text{subfS}(\phi_2) & \text{subfS}(\mathcal{P}_{\triangleright p}(\psi)) &= \{\mathcal{P}_{\triangleright p}(\psi)\} \cup \text{subfS}(\psi) \\ \text{subfS}(\phi_1 \mathcal{U} \phi_2) &= \text{subfS}(\phi_1 \wedge \neg\phi_2) & \text{subfS}(\mathbf{X}\phi) &= \text{subfS}(\phi) \end{aligned}$$

Given a state $s \in S$ and any formula $\phi' \in \text{subfS}(\phi)$, either $s \models_{\mathcal{M}} \phi'$ or $s \not\models_{\mathcal{M}} \phi'$. For each $s \in S$ and each $\phi' \in \text{subfS}(\phi)$, let the variable $t_s^{\phi'}$ be such that $t_s^{\phi'} = 1$ iff $s \models_{\mathcal{M}} \phi'$; and, $t_s^{\phi'} = 0$ iff $s \not\models_{\mathcal{M}} \phi'$. Following the definition of the various logical operators in PCTL, we can set up a set of constraints among these variables such that for any $\mathcal{M} \in [\mathcal{I}]$, the values taken by these variables is consistent with their intended semantic interpretation. We introduce the following additional variables to aid in setting up these constraints. For every state $s \in S$ and $\phi' \in \text{subfS}(\phi)$, let the auxiliary variables $f_s^{\phi'}$, and $u_s^{\phi'}$ be such that $t_s^{\phi'} = 1 \iff f_s^{\phi'} = 0 \iff u_s^{\phi'} = 1$ and $t_s^{\phi'} = 0 \iff f_s^{\phi'} = 1 \iff u_s^{\phi'} = -1$. Clearly, $t_s^{\phi'}$, $f_s^{\phi'}$, and $u_s^{\phi'}$ are related by the following set of constraints:

$$t_s^{\phi'} f_s^{\phi'} = 0 \quad t_s^{\phi'} + f_s^{\phi'} = 1 \quad 2t_s^{\phi'} = u_s^{\phi'} + 1$$

For every formula $\phi' \in \text{subfS}(\phi)$ of the form $\mathcal{P}_{\triangleright p}(\psi)$ and for every state $s \in S$, let p_s^{ψ} be the variable such that p_s^{ψ} denotes $\text{Prob}\{\pi \in \text{Path}(s) \mid \pi \models_{\mathcal{M}} \psi\}$ in \mathcal{M} .

For each state $s \in S$ and for each $\phi' \in \text{subfS}(\phi)$ exactly one of the following constraints hold depending on the form of ϕ' :

$$\begin{aligned} t_s^{\phi'} = 1 & \quad \text{if } \phi' = a \in L(s) & t_s^{\phi'} = 0 & \quad \text{if } \phi' = a \notin L(s) \\ t_s^{\phi'} = 1 - t_s^{\phi_1} & \quad \text{if } \phi' = \neg\phi_1 & t_s^{\phi_1} t_s^{\phi_2} = t_s^{\phi'} & \quad \text{if } \phi' = \phi_1 \wedge \phi_2 \\ u_s^{\phi'} p_s^{\psi} \geq u_s^{\phi'} p + \delta f_s^{\phi'} & \quad \text{if } \phi' = \mathcal{P}_{\geq p}(\psi) & u_s^{\phi'} p_s^{\psi} \geq u_s^{\phi'} p + \delta t_s^{\phi'} & \quad \text{if } \phi' = \mathcal{P}_{> p}(\psi) \\ u_s^{\phi'} p_s^{\psi} + \delta f_s^{\phi'} \leq u_s^{\phi'} p & \quad \text{if } \phi' = \mathcal{P}_{\leq p}(\psi) & u_s^{\phi'} p_s^{\psi} + \delta t_s^{\phi'} \leq u_s^{\phi'} p & \quad \text{if } \phi' = \mathcal{P}_{< p}(\psi) \end{aligned}$$

where δ is slack variable that is required to be strictly greater than 0.

Note that the above constraints do not reflect the fact that for each $\phi' \in \text{subfS}(\phi)$ of the form $\mathcal{P}_{\triangleright p}(\psi)$, p_s^{ψ} denotes $\text{Prob}\{\pi \in \text{Path}(s) \mid \pi \models_{\mathcal{M}} \psi\}$. To set up such constraints, we introduce the set $\text{subfP}(\phi)$ (of path sub-formulas) as follows:

$$\begin{aligned} \text{subfP}(a) &= \emptyset & \text{subfP}(\neg\phi) &= \text{subfP}(\phi) \\ \text{subfP}(\phi_1 \wedge \phi_2) &= \text{subfP}(\phi_1) \cup \text{subfP}(\phi_2) & \text{subfP}(\mathcal{P}_{\triangleright p}(\psi)) &= \{\psi\} \cup \text{subfP}(\psi) \\ \text{subfP}(\phi_1 \mathcal{U} \phi_2) &= \text{subfP}(\phi_1) \cup \text{subfP}(\phi_2) & \text{subfP}(\mathbf{X}\phi) &= \text{subfP}(\phi) \end{aligned}$$

Thus for all sub-formula of ϕ of the form $\mathcal{P}_{\triangleright p}(\psi)$, $\text{subfP}(\phi)$ contains ψ .

For any $\psi \in \text{subfP}(\phi)$ of the form $\mathbf{X}\phi_1$ and for each $s \in S$ the following constraint holds:

$$p_s^{\psi} = \sum_{s' \in S} p_{ss'} t_s^{\phi_1}$$

For each $\psi \in \text{subfS}(\phi)$ of the form $\phi_1 \mathcal{U} \phi_2$ and $s \in S$ the following constraints hold.

$$p_s^\psi = t_s^{\phi_2} + t_s^{\phi_1 \wedge \neg \phi_2} w_s^\psi \quad w_s^\psi = \sum_{s' \in S} p_{ss'} p_s^\psi$$

As in simple DTMC, if we consider the above constraints only, then we may not have unique solution for certain p_s^ψ . Therefore, we fix a γ such that $0 < \gamma < 1$. Then, as in simple DTMC model-checking, for each $\psi \in \text{subfP}(\phi)$ of the form $\phi_1 \mathcal{U} \phi_2$ and $s \in S$, we introduce the variables $p_s^{\prime\psi}$ and $w_s^{\prime\psi}$, such that the following constraints hold.

$$p_s^{\prime\psi} = t_s^{\phi_2} + t_s^{\phi_1 \wedge \neg \phi_2} w_s^{\prime\psi} \quad w_s^{\prime\psi} = \gamma \sum_{s' \in S} p_{ss'} p_s^{\prime\psi}$$

We want $p_s^\psi = 0$ if $p_s^{\prime\psi} = 0$. To ensure this, for each $\psi \in \text{subfP}(\phi)$ of the form $\phi_1 \mathcal{U} \phi_2$ and $s \in S$, we introduce the auxiliary variable δ_s^ψ and ensure that the following constraint hold.

$$\delta_s^\psi > 0 \quad p_s^\psi = \delta_s^\psi p_s^{\prime\psi}$$

Let $V(\mathcal{I}, \phi) = \{\delta\} \cup \bigcup_{s, s' \in S} \{p_{ss'}\} \cup \bigcup_{s \in S, \phi' \in \text{subfS}(\phi)} \{t_s^{\phi'}, f_s^{\phi'}, u_s^{\phi'}\} \cup \bigcup_{s \in S, \psi \in \text{subfP}(\phi)} \{p_s^\psi, w_s^\psi, p_s^{\prime\psi}, w_s^{\prime\psi}, \delta_s^\psi\}$ denote the set of variables over which the above constraints are described and let $C(\mathcal{I}, \phi)$ denote the above set of constraints.

Lemma 5. *For every solution $I: V(\mathcal{I}, \phi) \rightarrow \mathbb{R}$ of $C(\mathcal{I}, \phi)$, there exists a DTMC $\mathcal{M} = (S, s_I, \mathbf{P}, L) \in [\mathcal{I}]$ such that the following holds:*

1. $I(p_{ss'}) = \mathbf{P}(s, s')$ for any $s, s' \in S$
2. $t_s^{\phi'}, f_s^{\phi'} \in \{0, 1\}$ and $u_s^{\phi'} \in \{-1, 1\}$ for any $s \in S$ and $\phi' \in \text{subfS}(\phi)$
3. $t_s^{\phi'} = 1 \wedge f_s^{\phi'} = 0 \wedge u_s^{\phi'} = 1$ iff $s \models_{\mathcal{M}} \phi'$ for any $s \in S$ and $\phi' \in \text{subfS}(\phi)$
4. $t_s^{\phi'} = 0 \wedge f_s^{\phi'} = 1 \wedge u_s^{\phi'} = -1$ iff $s \models_{\mathcal{M}} \phi'$ for any $s \in S$ and $\phi' \in \text{subfS}(\phi)$
5. $p_s^\psi = \text{Prob}\{\pi \in \text{Path}(s) \mid \pi \models_{\mathcal{M}} \psi\}$ for any $\psi \in \text{subfP}(\phi)$

The proof follows from the observations made while setting up the constraints. An immediate consequence of the Lemma 5 is the following theorem.

Theorem 1. *If there exists a solution I of $C(\mathcal{I}, \phi)$ such that $I(t_{s_I}^\phi) = 1$, then there exists an $\mathcal{M} \in [\mathcal{I}]$ such that $\mathcal{M} \models \phi$.*

In order to check if $\mathcal{I} \models \phi$, the model checking algorithm sets up the constraints $C(\mathcal{I}, \neg\phi)$ and checks its feasibility. Clearly, checking the feasibility of $C(\mathcal{I}, \neg\phi)$ is equivalent to checking if a sentence with existential quantifiers is valid for the reals; the size of the sentence is polynomial in the size of the UMC. However, the constraints $C(\mathcal{I}, \neg\phi)$ are bilinear constraints, and we need to satisfy the conjunction of all these constraints (not an arbitrary boolean function). The feasibility of such constraints can be more efficiently checked viewing them as *bilinear matrix inequalities* (BMIs) for which algorithms [10, 9] and tools [16] have been developed. (More details can be seen in [23].) We also observe that to prove that the model checking problem can be solved in PSPACE, we could have constructed a simpler set of constraints by first guessing the values of the variables $t_s^{\phi'}, u_s^{\phi'}$, and $f_s^{\phi'}$ for the subformulas ϕ' , and then solving the

constraints resulting from those guesses; since $\text{NPSPACE} = \text{PSPACE}$, we can obtain a deterministic algorithm from this. However, we believe that in practice solving this single BMI presented here will be more efficient than solving the exponentially many simpler BMIs that this alternative approach would yield.

5.1 Complexity of Model-Checking UMC

We showed that the model-checking problem for UMC can be reduced to checking the validity of a formula in the existential theory of the reals. Therefore, the model-checking problem of UMC is in PSPACE.

The model checking problem for UMCs is however intractable: we can reduce both the satisfiability and validity of propositional boolean formulas to the model checking problem (details in [23]).

Theorem 2. *The model checking problem for UMC with respect to PCTL is NP-hard and co-NP-hard.*

6 Model-Checking IMDP

We consider the problem of model checking IMDPs in this section. We will solve the problem by showing that we can reduce IMDP model checking to model checking (classical) a Markov Decision Process (MDP) [4, 20]. Before presenting this reduction we recall some basic properties of the feasible solutions of a linear program and the definition of an MDP.

6.1 Linear Programming

Consider an IMDP $\mathcal{I} = (S, s_I, \check{\mathbf{P}}, \hat{\mathbf{P}}, L)$. For a given $s \in S$, let $IE(s)$ be the following set of inequalities over the variables $\{p_{ss'} \mid s' \in S\}$:

$$\sum_{s' \in S} p_{ss'} = 1 \quad \check{\mathbf{P}}(s, s') \leq p_{ss'} \leq \hat{\mathbf{P}}(s, s') \text{ for all } s' \in S$$

Definition 4. *A map $\theta^s: S \rightarrow [0, 1]$ is called a basic feasible solution (BFS) to the above set of inequalities $IE(s)$ iff $\{p_{ss'} = \theta^s(s') \mid s' \in S\}$ is a solution of $IE(s)$ and there exists a set $S' \subseteq S$ such that $|S'| \geq |S| - 1$ and for all $s' \in S'$ either $\theta^s(s') = \check{\mathbf{P}}(s, s')$ or $\theta^s(s') = \hat{\mathbf{P}}(s, s')$.*

Let Θ^s be the set of all BFS of $IE(s)$. The set of BFS of linear program have the special property that every other feasible solution can be expressed as a linear combination of basic feasible solutions. This is the content of the next proposition.

Proposition 1. *Let $\{p_{ss'} = \bar{p}_{ss'} \mid s' \in S\}$ be some solution of $IE(s)$. There there are $0 \leq \alpha_{\theta^s} \leq 1$ for all $\theta^s \in \Theta^s$, such that*

$$\bar{p}_{ss'} = \sum_{\theta^s \in \Theta^s} \alpha_{\theta^s} \theta^s(s') \text{ for all } s' \in S \quad \text{and} \quad \sum_{s \in S} \alpha_{\theta^s} = 1$$

Lemma 6. *The number of basic feasible solutions of $IE(s)$ in the worst case can be $O(|S|2^{|S|-1})$.*

6.2 Markov Decision Processes (MDP)

A Markov decision process (MDP) is a Markov chain that has non-deterministic transitions, in addition to the probabilistic ones. In this section we formally introduce this model along with some well-known observations about them.

Definition 5. *If S is the set of states of a system, a next-state probability distribution is a function $\mu : S \rightarrow [0, 1]$ such that $\sum_{s \in S} \mu(s) = 1$. For $s \in S$, $p(s)$ represents the probability of making a direct transition to s from the current state.*

Definition 6. *A Markov decision Process (MDP) is a 4-tuple $\mathcal{D} = (S, s_I, \tau, L)$, where*

1. S is a finite set of states,
2. $s_I \in S$ is the initial state,
3. $L : S \rightarrow 2^{\text{AP}}$ is a labeling function that maps states to sets of atomic propositions from a set AP,
4. τ is a function which associates to each $s \in S$ a finite set $\tau(s) = \{\mu_1^s, \dots, \mu_{k_s}^s\}$ of next-state probability distributions for transitions from s .

A path π in an MDP $\mathcal{D} = (S, s_I, \tau, L)$ is a non-empty sequence of the form $s_0 \xrightarrow{\mu_1} s_1 \xrightarrow{\mu_2} \dots$, where $s_i \in S$, $\mu_{i+1} \in \tau(s_i)$, and $\mu_{i+1}(s_{i+1}) > 0$ for all $i \geq 0$. A path can be either finite or infinite. We use π_{fin} to denote a finite path. Let $\text{last}(\pi_{\text{fin}})$ be the last state in the finite path π_{fin} . As in DTMC, we denote the i^{th} state in a path π by $\pi[i] = s_i$. We let $\text{Path}(s)$ and $\text{Path}_{\text{fin}}(s)$ be the set of all infinite and finite paths, respectively, starting at state s . To associate a probability measure with the paths, we resolve the non-deterministic choices by a randomized adversary, which is defined as follows:

Definition 7. *A randomized adversary A of an MDP \mathcal{D} is a function mapping every finite path π_{fin} of \mathcal{D} and an element of the set $\tau(\text{last}(\pi_{\text{fin}}))$ to $[0, 1]$, such that for a given finite path π_{fin} of \mathcal{D} , $\sum_{\mu \in \tau(\text{last}(\pi_{\text{fin}}))} A(\pi_{\text{fin}}, \mu) = 1$. Let $\mathcal{A}_{\mathcal{D}}$ denote the set of all possible randomized adversaries of the MDP \mathcal{D} . Let $\text{Path}^A(s)$ denote the subset of $\text{Path}(s)$ which corresponds to an adversary A .*

The behavior of an MDP under a given randomized adversary is purely probabilistic. If an MDP has evolved to the state s after starting from the state s_I and following the finite path π_{fin} , then it chooses the next-state distribution $\mu^s \in \tau(s)$ with probability $A(\pi_{\text{fin}}, \mu^s)$. Then it chooses the next state s' with probability $\mu^s(s')$. Thus the probability that a direct transition to s' takes place is $\sum_{\mu^s \in \tau(s)} A(\pi_{\text{fin}}, \mu^s) \mu^s(s')$. Thus as for IMDPs, one can define DTMC \mathcal{D}^A that captures the probabilistic behavior of MDP \mathcal{D} under adversary A and also associate a probability measure on execution paths. Given a MDP \mathcal{D} and a PCTL formula φ , we can define when $\mathcal{D} \models \varphi$ in a way analogous to the IMDPs (see Figure 2).

6.3 The Reduction

We are now ready to describe the model checking algorithm for IMDPs. Consider an IMDP $\mathcal{I} = (S, s_I, \vec{\mathbf{P}}, \hat{\mathbf{P}}, L)$. Recall from Section 6.1, we can describe the transition probability distributions from state s that satisfy the range constraints as the feasible

solutions of the linear program $IE(s)$. Furthermore, we denote by Θ^s is the set of all BFS of $IE(s)$. Define the following MDP $\mathcal{D} = (S', s'_I, \tau, L')$ where $S' = S$, $s'_I = s_I$, $L' = L$, and for all $s \in S$, $\tau(s) = \Theta^s$. Observe that \mathcal{D} is exponentially sized in \mathcal{I} , since $\tau(s)$ is exponential (see Lemma 6).

The main observation behind the reduction is that the MDP \mathcal{D} “captures” all the possible behaviors of the IMDP \mathcal{I} . This is the formal content of the next proposition.

Proposition 2. *For any adversary A for \mathcal{I} , we can define a randomized adversary A' such that $Prob_s^{\mathcal{I}^A} = Prob_s^{\mathcal{D}^{A'}}$ for every s , where $Prob_s^{X^A}$ is measure on paths from s defined by machine X under A . Similarly for every adversary A for \mathcal{D} , we can find an adversary A' for \mathcal{I} that defines the same probability measure on paths.*

Proof. Consider an adversary A for \mathcal{I} . For a path π_{fin} let $A(\pi_{\text{fin}}) = \mu \in \text{Steps}(\text{last}(\pi_{\text{fin}}))$. We know from Proposition 1, that there are α_{θ^s} for $\theta^s \in \Theta^s$ such that

$$\mu(s') = \sum_{\theta^s \in \Theta^s} \alpha_{\theta^s} \theta^s(s') \text{ for all } s' \in S \quad \text{and} \quad \sum_{s \in S} \alpha_{\theta^s} = 1$$

We now define $A'(\pi_{\text{fin}}, \theta^s) = \alpha_{\theta^s}$. It is straightforward to see that $Prob_s^{\mathcal{I}^A} = Prob_s^{\mathcal{D}^{A'}}$. The converse direction also can be proved similarly. \square

An important consequence of the above observation is the following main theorem.

Theorem 3. *For any PCTL formula φ , $\mathcal{I} \models \varphi$ iff $\mathcal{D} \models \varphi$.*

Thus, in order to model check IMDP \mathcal{I} , we can model check the MDP \mathcal{D} for which algorithms are known [4, 20]. The algorithms for MDP run in time (and space) which is polynomial in the size of the MDP. Thus, if we directly model check \mathcal{D} we get an EXP-TIME model checking algorithm for \mathcal{I} . However, we can improve this to get a PSPACE algorithm. The reason for this is that it is known that as far as model checking MDPs is concerned, we can restrict our attention to *deterministic, memoryless* adversaries, i.e., adversaries that always pick the same single non-deterministic choice whenever a state is visited.

Proposition 3 ([4, 20]). *Let \mathcal{A}_{det} be the set of deterministic, memoryless adversaries for MDP \mathcal{D} , i.e., for all $A \in \mathcal{A}_{\text{det}}$, $A(s, \mu) = 1$ for exactly one $\mu \in \tau(s)$. Consider a PCTL formula $\varphi = \mathcal{P}_{\bowtie p}(\psi)$ such that the truth or falsity of every subformula of ψ in every state of \mathcal{D} is already determined. Then $\mathcal{D} \models \varphi$ iff $\mathcal{D}^A \models \varphi$ for all $A \in \mathcal{A}_{\text{det}}$.*

For every subformula of the form $\mathcal{P}_{\bowtie p}(\psi)$, our model checking algorithm, will model check each of the DTMCs \mathcal{D}^A , where A is a deterministic, memoryless adversary. This will give us the desired PSPACE algorithm.

Theorem 4. *The model-checking algorithm for IMDP is in PSPACE.*

Proof. From Lemma 6, we know that the total number of BFSs is $O(|S|2^{|S|-1})$. Hence the total number of DTMCs \mathcal{D}^A for $A \in \mathcal{A}_{\text{det}}$ is $O(|S|^{|S|}2^{|S|^2-|S|})$. By reusing space for every subformula $\mathcal{P}_{\bowtie p}(\psi)$, all of these model checking problems can be solved in PSPACE. \square

6.4 Iterative Algorithm

The above PSPACE algorithm is computationally expensive for large IMDPs. Therefore, we propose an alternative iterative algorithm motivated by a similar algorithm in [2].

The iterative model checking algorithm for PCTL over IMDPs works exactly as for DTMCs with the exception of handling of $\mathcal{P}_{\bowtie p}(\psi)$. For these, we need to check if $p_s^A(\psi) = \text{Prob}_s^A(\{\pi \in \text{Path}^A(s) \mid \pi \models \psi\})$ satisfies the bound $\bowtie p$ for all adversaries $A \in \mathcal{A}_{\mathcal{I}}$. Let $p_s^{\max}(\psi)$ and $p_s^{\min}(\psi)$ be the *minimum* or *maximum* probability, respectively, for all adversaries $A \in \mathcal{A}_{\mathcal{I}}$, i.e.,

$$p_s^{\max}(\psi) \stackrel{\text{def}}{=} \sup_{A \in \mathcal{A}_{\mathcal{I}}} [p_s^A(\psi)], \quad p_s^{\min}(\psi) \stackrel{\text{def}}{=} \inf_{A \in \mathcal{A}_{\mathcal{I}}} [p_s^A(\psi)].$$

Then if $\bowtie \in \{<, \leq\}$,

$$\text{Sat}(\mathcal{P}_{\bowtie p}(\psi)) = \{s \in S \mid p_s^{\max}(\psi) \bowtie p\}$$

and if $\bowtie \in \{>, \geq\}$,

$$\text{Sat}(\mathcal{P}_{\bowtie p}(\psi)) = \{s \in S \mid p_s^{\min}(\psi) \bowtie p\}$$

We next describe how to compute the values $p_s^{\max}(\psi)$ and $p_s^{\min}(\psi)$ for $\psi = \mathbf{X}\phi$ and $\psi = \phi_1 \mathbf{U} \phi_2$. Recall that Θ^s is the set of all BFS of $IE(s)$. It can be shown following [2] that $p_s^{\max} = \lim_{n \rightarrow \infty} p_s^{\max(n)}$ where:

$$p_s^{\max(n)} = \begin{cases} 1 & \text{if } s \in S^{\text{yes}} \\ 0 & \text{if } s \in S^{\text{no}} \\ 0 & \text{if } s \in S^? \text{ and } n = 0 \\ \max_{\{\bar{p}_{ss'} \mid s' \in S\} \in \Theta^s} \left\{ \sum_{s' \in S} \bar{p}_{ss'} \cdot p_{s'}^{\max(n-1)} \right\} & \text{if } s \in S^? \text{ and } n > 0 \end{cases}$$

and $p_s^{\min} = \lim_{n \rightarrow \infty} p_s^{\min(n)}$ where:

$$p_s^{\min(n)} = \begin{cases} 1 & \text{if } s \in S^{\text{yes}} \\ 0 & \text{if } s \in S^{\text{no}} \\ 0 & \text{if } s \in S^? \text{ and } n = 0 \\ \min_{\{\bar{p}_{ss'} \mid s' \in S\} \in \Theta^s} \left\{ \sum_{s' \in S} \bar{p}_{ss'} \cdot p_{s'}^{\min(n-1)} \right\} & \text{if } s \in S^? \text{ and } n > 0 \end{cases}$$

Note that although the size of Θ^s can be $O(|S|2^{|S|-1})$ (by Lemma 6), the computation of the expressions

$$\max_{\{\bar{p}_{ss'} \mid s' \in S\} \in \Theta^s} \left\{ \sum_{s' \in S} \bar{p}_{ss'} \cdot p_{s'}^{\max(n-1)} \right\} \quad \text{or} \quad \min_{\{\bar{p}_{ss'} \mid s' \in S\} \in \Theta^s} \left\{ \sum_{s' \in S} \bar{p}_{ss'} \cdot p_{s'}^{\min(n-1)} \right\} \quad (5)$$

can be done in $O(|S|)$ time as follows:

We consider the ordering $s_1, s_2, \dots, s_{|S|}$ of the states of S such that $p_{s_1}^{\max(n-1)}, p_{s_2}^{\max(n-1)}, \dots, p_{s_{|S|}}^{\max(n-1)}$ is in descending order. Then the following result holds.

Lemma 7.

- a) *There exists an $1 \leq i \leq |S|$ such that $\{\hat{\mathbf{P}}(s, s_1), \dots, \hat{\mathbf{P}}(s, s_{i-1}), q, \check{\mathbf{P}}(s, s_{i+1}), \dots, \check{\mathbf{P}}(s, s_{|S|})\}$ is a BFS of $IE(s)$, where $q = 1 - \sum_{1 \leq j \leq (i-1)} \hat{\mathbf{P}}(s, s_j) - \sum_{(i+1) \leq j \leq |S|} \check{\mathbf{P}}(s, s_j)$.*
- b) *and for that i*

$$\max_{\{\bar{p}_{ss'} \mid s' \in S\} \in \Theta^S} \left\{ \sum_{s' \in S} \bar{p}_{ss'} \cdot p_{s'}^{\max(n-1)} \right\} = p_{s_i}^{\max(n-1)} \cdot q$$

$$+ \sum_{1 \leq j \leq (i-1)} p_{s_j}^{\max(n-1)} \cdot \hat{\mathbf{P}}(s, s_j) + \sum_{(i+1) \leq j \leq |S|} p_{s_j}^{\max(n-1)} \cdot \check{\mathbf{P}}(s, s_j)$$

Proof.

- a) Let i_0 be defined as follows:

$$i_0 = \min \left\{ i \mid \sum_{j=1}^i \hat{\mathbf{P}}(s, s_j) + \sum_{j=i+1}^{|S|} \check{\mathbf{P}}(s, s_j) \geq 1 \right\}$$

Observe that such an i_0 must exist if the IMDP is well-defined. Consider the solution $\{\hat{\mathbf{P}}(s, s_1), \dots, \hat{\mathbf{P}}(s, s_{i_0-1}), q, \check{\mathbf{P}}(s, s_{i_0+1}), \dots, \check{\mathbf{P}}(s, s_{|S|})\}$ where $q = 1 - \sum_{1 \leq j \leq (i_0-1)} \hat{\mathbf{P}}(s, s_j) - \sum_{(i_0+1) \leq j \leq |S|} \check{\mathbf{P}}(s, s_j)$. This solution is a BFS of $IE(s)$.

- b) Let $\{\bar{p}_{ss_1}, \dots, \bar{p}_{ss_{|S|}}\}$ be any solution (it may be BFS or not) of $IE(s)$. Then by simple algebraic simplification it can be shown that

$$\sum_{1 \leq j \leq (i-1)} p_{s_j}^{\max(n-1)} \cdot \hat{\mathbf{P}}(s, s_j) + p_{s_i}^{\max(n-1)} \cdot q + \sum_{(i+1) \leq j \leq |S|} p_{s_j}^{\max(n-1)} \cdot \check{\mathbf{P}}(s, s_j) \geq \sum_{s' \in S} \bar{p}_{ss'} \cdot p_{s'}^{\max(n-1)}$$

given the fact that $p_{s_1}^{\max(n-1)} \geq p_{s_2}^{\max(n-1)} \geq \dots \geq p_{s_{|S|}}^{\max(n-1)}$, and $\check{\mathbf{P}}(s, s') \leq \bar{p}_{ss'} \leq \hat{\mathbf{P}}(s, s')$ for all $s' \in S$. □

Similarly, if we consider the ordering $s_1, s_2, \dots, s_{|S|}$ of the states of S such that $p_{s_1}^{\min(n-1)}, p_{s_2}^{\min(n-1)}, \dots, p_{s_{|S|}}^{\min(n-1)}$ is in ascending order, then the above Lemma holds with max replaced by min.

The expressions (5) can be computed in $O(|S|)$ time by finding an i as in Lemma 7.

6.5 Lower Bound for IMDP Model-Checking

We can show that the model checking problem for IMDPs is P-hard. The result follows from observing that the problem of determining the truth value of propositional logic formula under an assignment (which is known to be P-complete) can be reduced to the PCTL model checking problem of DTMCs; since DTMCs are special IMDPs, the result follows. The details can be found in [23].

7 Conclusion

We have investigated the PCTL model checking problem for two semantic interpretations of IDTMCs, namely UMC and IMDP. We proved the upper bounds and the lower bounds on the complexity of the model checking problem for these models. Our bounds however are not tight. Finding tight lower and upper bounds for these model-checking problems is an interesting open problem.

Acknowledgment

We would like to thank anonymous referees and Timo Latvala for providing valuable comments. This work is supported in part by the ONR Grant N00014-02-1-0715, the NSF Grants NSF CNS 05-09321, NSF CCF 04-29639, NSF CCF 04-48178, and the Motorola Grant Motorola RPF #23.

References

1. A. Aziz, V. Singhal, R. K. Brayton, and A. L. Sangiovanni-Vincentelli. It usually works: The temporal logic of stochastic systems. In *Proc. of Computer Aided Verification*, volume 939, pages 155–165, 1995.
2. C. Baier. On algorithmic verification methods for probabilistic systems. Habilitation Thesis. Fakultät für Mathematik und Informatik, Universität Mannheim, 1998.
3. C. Baier and M. Z. Kwiatkowska. Model checking for a probabilistic branching time logic with fairness. *Distributed Computing*, 11(3):125–155, 1998.
4. A. Bianco and L. de Alfaro. Model checking of probabilistic and nondeterministic systems. In *Proceedings of 15th Conference on the Foundations of Software Technology and Theoretical Computer Science (FSTTCS'95)*, volume 1026 of *LNCS*.
5. L. Blum, M. Shub, and S. Smale. On a theory of computation and complexity over real numbers: NP-completeness, recursive functions and universal machines. *Bulletin of the American Mathematical Society*, 21:1–46, 1989.
6. J. Canny. Some algebraic and geometric computations in PSPACE. In *20th ACM Symposium on Theory of Computing (STOC'88)*, pages 460–467, 1988.
7. C. Courcoubetis and M. Yannakakis. Markov decision processes and regular events. In *Proceedings of the seventeenth international colloquium on Automata, languages and programming*, pages 336–349, 1990.
8. C. Courcoubetis and M. Yannakakis. The complexity of probabilistic verification. *Journal of ACM*, 42(4):857–907, 1995.
9. M. Fukuda and M. Kojima. Branch-and-cut algorithms for the bilinear matrix inequality eigenvalue problem. *Comput. Optim. Appl.*, 19(1):79–105, 2001.
10. K. C. Goh, M. G. Safonov, and G. P. Papavassilopoulos. Global optimization for the biaffine matrix inequality problem. *Journal of Global Optimization*, 7:365–380, 1995.
11. H. Hansson and B. Jonsson. A logic for reasoning about time and reliability. *Formal Aspects of Computing*, 6(5):512–535, 1994.
12. B. Jonsson and K. G. Larsen. Specification and refinement of probabilistic processes. In *Proceedings of the IEEE Symposium on Logic in Computer Science*, pages 266–277, 1991.
13. J. Kemeny, J. Snell, and A. Knapp. *Denumerable Markov chains*. Springer, 1976.
14. I. O. Kozine and L. V. Utkin. Interval-valued finite markov chains. *Reliable Computing*, 8(2):97–113, 2002.

15. V. P. Kuznetsov. Interval statistical models. *Radio and Communication*, 1991.
16. PENbmi. <http://www.penopt.com/>.
17. M. Puterman. *Markov decision processes: discrete stochastic dynamic programming*. Wiley, New York, 1994.
18. J. Renegar. A faster pspace algorithm for deciding the existential theory of the reals. In *29th Annual IEEE Symposium on Foundations of Computer Science*, pages 291–295, 1988.
19. J. Rutten, M. Kwiatkowska, G. Norman, and D. Parker. *Mathematical Techniques for Analyzing Concurrent and Probabilistic Systems*, volume 23 of *CRM Monograph Series*. American Mathematical Society, 2004.
20. R. Segala. *Modeling and Verification of Randomized Distributed Real-Time Systems*. PhD thesis, MIT, 1995.
21. R. Segala and N. A. Lynch. Probabilistic simulations for probabilistic processes. In *International Conference on Concurrency Theory*, pages 481–496, 1994.
22. K. Sen, M. Viswanathan, and G. Agha. Statistical model checking of black-box probabilistic systems. In *16th conference on Computer Aided Verification (CAV'04)*, volume 3114 of *LNCS*, pages 202–215, 2004.
23. K. Sen, M. Viswanathan, and G. Agha. Model-checking markov chains in the presence of uncertainties. Technical Report UIUCDCS-R-2006-2677, UIUC, 2006.
24. O. Toker and H. Özbay. On the NP-hardness of solving bilinear matrix in equalities and simultaneous stabilization with static output feedback. In *Proc. of American Control Conference*, 1995.
25. M. Y. Vardi. Automatic verification of probabilistic concurrent finite-state programs. In *26th Annual Symposium on Foundations of Computer Science*, pages 327–338. IEEE, 1985.
26. P. Walley. Measures of uncertainty in expert systems. *Artificial Intelligence*, 83:1–58, 1996.