

Model-Checking ω -Regular Properties of Interval Markov Chains^{*}

Krishnendu Chatterjee¹, Koushik Sen¹, and Thomas A. Henzinger^{1,2}

¹ University of California, Berkeley, USA

² EPFL, Switzerland

{c.krish,k.sen,tah}@eecs.berkeley.edu

Abstract. We study 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. We call this semantic interpretation 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 introduce a logic ω -PCTL that can express liveness, strong fairness, and ω -regular properties (such properties cannot be expressed in PCTL). We show that the ω -PCTL model checking problem for Uncertain Markov Chain semantics is decidable in PSPACE (same as the best known upper bound for PCTL) and for Interval Markov Decision Process semantics is decidable in coNP (improving the previous known PSPACE bound for PCTL). We also show that the qualitative fragment of the logic can be solved in coNP for the UMC interpretation, and can be solved in polynomial time for a sub-class of UMCs. We also prove lower bounds for these model checking problems. We show that the model checking problem of IDTMCs with LTL formulas can be solved for both UMC and IMDP semantics by reduction to the model checking problem of IDTMC with ω -PCTL formulas.

1 Introduction

Discrete Time Markov Chains (DTMCs) are often used to model and analyze the reliability and performance of computer systems [6, 9, 15, 12]. A DTMC consists of a finite number of states and a fixed probability of transition from one state

^{*} This research was supported in part by the AFOSR MURI grant F49620-00-1-0327, the NSF grant CNS-0720906, the NSF ITR grant CCR-0225610 and CCR-0234690, the Swiss National Science Foundation (NCCR MICS and Indo-Swiss Research Programme), and the ARTIST2 European Network of Excellence.

to another state. The fixed probability assumption in a DTMC may often not be realistic in practice [11, 14, 22, 13]. For example, in case of an open system that interacts with an environment, transition probabilities may not be known precisely due to incomplete knowledge about the environment. Imprecision in the transition probabilities may arise if the probabilities in the system model are estimated through statistical experiments, which only provide bounds on the transition probabilities.

The model of *Interval-valued Discrete-time Markov Chains (IDTMC)* has been introduced [11, 13] to faithfully capture these system uncertainties. IDTMCs are DTMC models where the exact transition probability is not known, and instead the transition probability is assumed to lie within a range. Three valued abstractions of DTMCs also naturally result in IDTMCs [8]. Two semantic interpretations have been suggested for such models. *Uncertain Markov Chains (UMC)* [11, 18] 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)* [18], the uncertainty is resolved through non-determinism. In other words, each time a state is visited, a transition distribution that respects the interval constraints is adversarially picked, and a probabilistic step is taken 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.

The problem of model checking PCTL specifications for IDTMC was studied in [18]. PSPACE model checking algorithms were given for both UMCs and IMDPs. The model checking problem for UMCs was shown to be both NP-hard and coNP-hard. For IMDPs, a PTIME-hardness was shown; in fact, this is a consequence of the PTIME-hardness of (classical) DTMC model checking [6].

The logic PCTL [9], which extends computation tree logic (CTL) with probabilities, does not allow arbitrarily nested path formulas. Therefore, PCTL cannot express properties that depend on the set of states that appears infinitely often, e.g., liveness properties cannot be expressed in PCTL. *In order to address this limitation of PCTL, we introduce ω -PCTL.* In ω -PCTL, we allow Büchi conditions, that require a set of states to be visited infinitely often, its dual coBüchi conditions, and their boolean combinations. Since we allow Büchi conditions, liveness (or weak-fairness) conditions can be expressed in ω -PCTL. Moreover, since we allow boolean combinations of Büchi and coBüchi conditions, strong fairness conditions can also be expressed in ω -PCTL. The logic ω -PCTL can express all ω -regular conditions, and thus forms a robust specification language to specify properties that commonly arise in verification of probabilistic systems.

In addition to the UMC interpretation, we also consider the sub-class of UMC interpretation that restricts the DTMCs obtained from an IDTMC as follows: if the upper bound of a transition probability is positive, then the actual transition probability is also positive. In many situations the upper bound on a transition probability is positive if the transition is observed (i.e., the actual transition

probability is positive, though no positive lower bound may be known). We call this sub-class as PUMCs (Positive UMCs).

In this paper, we study the problem of model checking ω -PCTL specifications for DTMCs and IDTMCs. We first show that the ω -PCTL model checking problem for DTMCs can be solved in polynomial time. We then show that the ω -PCTL model checking problem for PUMC and UMC interpretations is decidable in PSPACE and for IMDP interpretations is decidable in coNP. These results extend and improve the best known PSPACE bound for PCTL model checking to a much richer logic that can express ω -regular properties. We also show that the qualitative fragment of the logic (called ω -QPCTL) can be solved in polynomial time for the PUMCs and in coNP for the UMCs. The results of PCTL model checking algorithm do not extend straightforwardly to ω -PCTL model checking. We first present the model checking algorithm for PUMC semantics using results on Markov chains, and then reduction to a formula in the existential theory of reals. The result for UMC semantics is then obtained by partitioning the UMCs in equivalence classes of PUMCs. The IMDP model checking algorithm requires a precise characterization of optimal strategies in MDPs with Müller objectives. We also prove lower bounds for these model checking problems: we show that the PCTL model checking problem is both NP-hard and coNP-hard for PUMCs, and the NP and coNP-hardness for PCTL model checking for UMCs follows from [18]. We also present model checking algorithms for IDTMCs with LTL path formulas for PUMC, UMC, and IMDP interpretations, and the result is obtained by reduction to ω -PCTL formulas. Table 1 summarizes the complexity of model checking of the various classes of Markov chains under uncertainty with respect the various fragments of ω -PCTL.

| Models | PCTL | | ω -QPCTL | | ω -PCTL | |
|--------|-------------|-------------|-----------------|-------------|----------------|-------------|
| | Lower Bound | Upper Bound | Lower Bound | Upper Bound | Lower Bound | Upper Bound |
| DTMC | PTIME | PTIME | | PTIME | PTIME | PTIME |
| PUMC | NP and coNP | PSPACE | | PTIME | NP and coNP | PSPACE |
| UMC | NP and coNP | PSPACE | | coNP | NP and coNP | PSPACE |
| IMDP | PTIME | coNP | | coNP | PTIME | coNP |

Table 1. Complexity of DTMC and IDTMC model checking

2 Formal Models

In this section, we recall the definitions of IDTMC, UMC, and IMDP from [18] and introduce the definition of PUMC.

Definition 1. A discrete-time Markov chain (DTMC) is a 3-tuple $\mathcal{M} = (S, \mathbf{P}, L)$, where (1) S is a finite set of states; (2) $\mathbf{P}: S \times S \rightarrow [0, 1]$ is a transition probability matrix, such that $\sum_{s' \in S} \mathbf{P}(s, s') = 1$; and (3) $L: S \rightarrow 2^{\text{AP}}$ 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 $Path(s)$ 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 $\pi \in Path(s_0)$ such that $\pi[i] = s_i$ (for $0 \leq i \leq k$). Let \mathcal{B} be the smallest σ -algebra on $Path(s_0)$ which contains all the cylinders $C(s_0 s_1 \dots s_k)$. The measure μ on cylinder sets can be defined as follows: $\mu(C(s_0 s_1 \dots s_k)) = 1$ if $k = 0$; otherwise $\mu(C(s_0 s_1 \dots s_k)) = \mathbf{P}(s_0, s_1) \dots \mathbf{P}(s_{k-1}, s_k)$. 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 4-tuple $\mathcal{I} = (S, \hat{\mathbf{P}}, \hat{\mathbf{P}}, L)$, where (1) S is a finite set of states; (2) $\hat{\mathbf{P}}: S \times S \rightarrow [0, 1]$ is a transition probability matrix, where each $\hat{\mathbf{P}}(s, s')$ gives the lower bound of the transition probability from the state s to the state s' ; (3) $\hat{\mathbf{P}}: S \times S \rightarrow [0, 1]$ 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' ; and (4) $L: S \rightarrow 2^{\text{AP}}$ is a labeling function that maps states to sets of atomic propositions from a set AP.

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

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

Positive Uncertain Markov Chains (PUMCs). We consider Positive Uncertain Markov Chains (PUMCs) semantics, for which we will obtain more efficient model checking algorithms for the qualitative fragment of the logic that we will consider, and the results will also be useful in the analysis of UMCs. Given an IDTMC \mathcal{I} , we denote by $[\mathcal{I}]_P \subseteq [\mathcal{I}]$ the infinite set of DTMCs (S, \mathbf{P}, L) such that the following conditions hold: (1) $\hat{\mathbf{P}}(s, s') \leq \mathbf{P}(s, s') \leq \hat{\mathbf{P}}(s, s')$ for all pairs of states s and s' in S ; (2) if $\hat{\mathbf{P}}(s, s') > 0$, then $\mathbf{P}(s, s') > 0$, for all $s, s' \in S$. In the semantics for Positive Uncertain Markov Chains (PUMCs), we assume that the external environment non-deterministically picks a DTMC from $[\mathcal{I}]_P$.

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 $Steps(s)$ be the set of probability density functions over S defined as follows: $Steps(s) = \{\mu: S \rightarrow \mathbb{R}^{\geq 0} \mid \sum_{s' \in S} \mu(s') = 1 \text{ and } \check{\mathbf{P}}(s, s') \leq \mu(s') \leq \hat{\mathbf{P}}(s, s') \text{ for all } s' \in S\}$. In an IMDP, at every state $s \in S$, a probability density function μ is chosen non-deterministically from the set $Steps(s)$. A successor state s' is then chosen according to the probability distribution μ over S .

A *path* π in an IMDP $\mathcal{I} = (S, \check{\mathbf{P}}, \hat{\mathbf{P}}, L)$ is a non-empty sequence of the form $s_0 \xrightarrow{\mu_1} s_1 \xrightarrow{\mu_2} \dots$, where $s_i \in S$, $\mu_{i+1} \in Steps(s_i)$, 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 $last(\pi_{\text{fin}})$ 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 $Path(s)$ and $Path_{\text{fin}}(s)$ 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 $Steps(last(\pi_{\text{fin}}))$. Let $\mathcal{A}_{\mathcal{I}}$ denote the set of all possible adversaries of the IMDP \mathcal{I} . Let $Path^A(s)$ denote the subset of $Path(s)$ which corresponds to A .

The behavior of an IMDP $\mathcal{I} = (S, \check{\mathbf{P}}, \hat{\mathbf{P}}, L)$ under a given adversary A is purely probabilistic. The behavior of a IMDP \mathcal{I} from a state s can be described by an infinite-state DTMC $\mathcal{M}^A = (S^A, \mathbf{P}^A, L^A)$ where (a) $S^A = Path_{\text{fin}}(s)$; (b) $\mathbf{P}^A(\pi_{\text{fin}}, \pi'_{\text{fin}}) = A(\pi_{\text{fin}})(s')$ if π'_{fin} is of the form $\pi_{\text{fin}} \xrightarrow{A(\pi_{\text{fin}})} s'$; and 0 otherwise. There is a one-to-one correspondence between the paths of \mathcal{M}^A and $Path^A(s)$ of \mathcal{I} . Therefore, we can define a probability measure $Prob_s^A$ over the set of paths $Path^A(s)$ using the probability measure of the DTMC \mathcal{M}^A .

3 ω -Probabilistic Computation Tree Logic (ω -PCTL)

In this paper, we consider an extension of PCTL that can express ω -regular properties. We call the logic ω -PCTL. The formal syntax and semantics of this logic is as follows.

ω -PCTL Syntax. We define the syntax of ω -PCTL and its qualitative fragment as follows:

$$\begin{aligned} \phi &::= true \mid a \mid \neg\phi \mid \phi \wedge \phi \mid \mathcal{P}_{\bowtie p}(\psi) \\ \psi &::= \phi \mathcal{U} \phi \mid \mathbf{X}\phi \mid \psi^\omega \\ \psi^\omega &::= \text{Buchi}(\phi) \mid \text{coBuchi}(\phi) \mid \psi^\omega \wedge \psi^\omega \mid \psi^\omega \vee \psi^\omega \end{aligned}$$

where $a \in \text{AP}$ is an atomic proposition, and $\bowtie \in \{<, \leq, >, \geq\}$, $p \in [0, 1]$. Here ϕ represents a *state* formula, ψ represents a *path* formula, and ψ^ω represents path formulas that depend on the set of states that appear infinitely often in a path (we call them infinitary path formulas). The qualitative fragment of the logic,

denoted as ω -QPCTL, consists of formulas ϕ_Q such that in all sub-formulas $\mathcal{P}_{\bowtie p}(\psi)$ of ϕ_Q we have $p \in \{0, 1\}$, i.e., the comparison of the probability of satisfying a path formula is only made with 1 and 0 only. The logic PCTL is obtained from ω -PCTL where only path formulas of the form $\phi \mathcal{U} \phi$ and $\mathbf{X}\phi$ are considered, i.e., formulas obtained as ψ^ω are not allowed. The canonical Rabin and Streett conditions (strong fairness conditions) can be expressed as conjunction and disjunction of Büchi and coBüchi conditions. Hence ω -PCTL can express Rabin and Streett conditions. Since Rabin and Streett conditions are canonical forms to express ω -regular properties [19], ω -PCTL can express ω -regular properties.

ω -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:

| | |
|--|--|
| $s \models_{\mathcal{M}} \text{true}$ | |
| $s \models_{\mathcal{M}} a$ | iff $a \in L(s)$ |
| $s \models_{\mathcal{M}} \neg\phi$ | iff $s \not\models_{\mathcal{M}} \phi$ |
| $s \models_{\mathcal{M}} \phi_1 \wedge \phi_2$ | iff $s \models_{\mathcal{M}} \phi_1$ and $s \models_{\mathcal{M}} \phi_2$ |
| $s \models_{\mathcal{M}} \mathcal{P}_{\bowtie p}(\psi)$ | iff $\text{Prob}\{\pi \in \text{Path}(s) \mid \pi \models_{\mathcal{M}} \psi\} \bowtie p$ |
| $\pi \models_{\mathcal{M}} \mathbf{X}\phi$ | iff $\pi[1] \models_{\mathcal{M}} \phi$ |
| $\pi \models_{\mathcal{M}} \phi_1 \mathcal{U} \phi_2$ | iff $\exists i \geq 0 (\pi[i] \models_{\mathcal{M}} \phi_2 \text{ and } \forall j < i. \pi[j] \models_{\mathcal{M}} \phi_1)$ |
| $\pi \models_{\mathcal{M}} \text{Buchi}(\phi)$ | iff $\forall i \geq 0. \exists j \geq i. (\pi[j] \models_{\mathcal{M}} \phi)$ |
| $\pi \models_{\mathcal{M}} \text{coBuchi}(\phi)$ | iff $\exists i \geq 0. \forall j \geq i. (\pi[j] \models_{\mathcal{M}} \phi)$ |
| $\pi \models_{\mathcal{M}} \psi_1^\omega \wedge \psi_2^\omega$ | iff $\pi \models_{\mathcal{M}} \psi_1^\omega$ and $\pi \models_{\mathcal{M}} \psi_2^\omega$ |
| $\pi \models_{\mathcal{M}} \psi_1^\omega \vee \psi_2^\omega$ | iff $\pi \models_{\mathcal{M}} \psi_1^\omega$ or $\pi \models_{\mathcal{M}} \psi_2^\omega$. |

It can be shown that for any path formula ψ and any state s , the set $\{\pi \in \text{Path}(s) \mid \pi \models_{\mathcal{M}} \psi\}$ is measurable [21]. For a path formula ψ we denote by $\text{Prob}_s(\psi)$ the probability of satisfying ψ from s , i.e., $\text{Prob}_s(\psi) = \text{Prob}[\{\pi \in \text{Path}(s) \mid \pi \models_{\mathcal{M}} \psi\}]$. A formula $\mathcal{P}_{\bowtie p}(\psi)$ is satisfied by a state s if $\text{Prob}_s[\psi] \bowtie p$. 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 . The formula $\text{Buchi}(\phi)$ is true over a path π if the path infinitely often visits states that satisfy ϕ . The formula $\text{coBuchi}(\phi)$ is true over a path π if after a finite prefix the path visits only states that satisfy ϕ . Given a DTMC \mathcal{M} and an ω -PCTL state formula ϕ , we denote by $[\phi]_{\mathcal{M}} = \{s \mid s \models_{\mathcal{M}} \phi\}$ the set of the states that satisfy ϕ . Given a DTMC \mathcal{M} and an ω -PCTL path formula ψ we denote by $W_{\mathcal{M}}(\psi) = \{s \mid \text{Prob}_s(\psi) = 1\}$ the set of states that satisfy ψ with probability 1.

ω -PCTL Semantics for UMC. Given an IDTMC \mathcal{I} and an ω -PCTL state formula ϕ , we denote by $[\phi]_{\mathcal{I}} = \bigcap_{\mathcal{M} \in [\mathcal{I}]} [\phi]_{\mathcal{M}}$. Note that $s \notin [\phi]_{\mathcal{I}}$ does not imply that $s \in [\neg\phi]_{\mathcal{I}}$. This is because there may exist $\mathcal{M}, \mathcal{M}' \in [\mathcal{I}]$ such that $s \models_{\mathcal{M}} \phi$ and $s \models_{\mathcal{M}'} \neg\phi$. The semantics of ω -PCTL for PUMCs are obtained similarly: given an IDTMC \mathcal{I} and an ω -PCTL state formula ϕ , we denote by $[\phi]_{\mathcal{I}_P} = \bigcap_{\mathcal{M} \in [\mathcal{I}]_P} [\phi]_{\mathcal{M}}$.

ω -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 $s \models_{\mathcal{I}} \phi$ (or $\pi \models_{\mathcal{I}} \phi$), and the semantics is very similar to the one of DTMC other than path formulas with probabilistic operator which is defined below:

$$s \models_{\mathcal{I}} \mathcal{P}_{\bowtie p}(\psi) \text{ iff } \text{Prob}_s^A(\{\pi \in \text{Path}^A(s) \mid \pi \models_{\mathcal{I}} \psi\}) \bowtie p \text{ for all } A \in \mathcal{A}$$

The model checking of IDTMC with respect to the two semantics can give different results. An example illustrating this fact for the PCTL logic can be found in [18].

4 DTMC Model Checking

In this section we outline the basic model checking algorithm for (classical) DTMCs for ω -PCTL. We start with a few notations.

Graph of a DTMC. Given a DTMC $\mathcal{M} = (S, \mathbf{P}, L)$ we define a graph $G_{\mathcal{M}} = (S_{\mathcal{M}}, E_{\mathcal{M}}, L_{\mathcal{M}})$ for \mathcal{M} where $S_{\mathcal{M}} = S$, $L_{\mathcal{M}} = L$, and the set of edges $E_{\mathcal{M}} = \{(s, s') \mid \mathbf{P}(s, s') > 0\}$ consists of state pairs (s, s') such that the transition probability from s to s' is positive. Given two DTMCs \mathcal{M}_1 and \mathcal{M}_2 , they are graph equivalent, denoted by $\mathcal{M}_1 \equiv \mathcal{M}_2$, iff $S_{\mathcal{M}_1} = S_{\mathcal{M}_2}$, $E_{\mathcal{M}_1} = E_{\mathcal{M}_2}$, and $L_{\mathcal{M}_1} = L_{\mathcal{M}_2}$, i.e., the set of states, the set of edges, and the labeling function in \mathcal{M}_1 and \mathcal{M}_2 coincide. Observe that though the set of edges in \mathcal{M}_1 and \mathcal{M}_2 coincide, the exact transition probabilities in \mathcal{M}_1 and \mathcal{M}_2 can be different. For a state formula ϕ (resp. a set $U \subseteq S$ of states) we denote by $\diamond\phi$ (resp. $\diamond U$) eventually ϕ (resp. eventually U), i.e., the PCTL formula $\text{true } \mathcal{U} \phi$ (resp. $\text{true } \mathcal{U} U$).

Lemma 1. *Given a DTMC \mathcal{M} and an infinitary path formula ψ^ω , we have $\text{Prob}_s(\psi^\omega) = \text{Prob}_s(\diamond(W_{\mathcal{M}}(\psi^\omega)))$.*

Graph equivalence and ω -QPCTL. The truth of a qualitative PCTL formula ϕ (i.e., a QPCTL formula) does not depend on the precise transition probabilities of a DTMC, but depends only on the underlying graph structure of the DTMC. Lemma 2 extends the result to ω -QPCTL formulas. Formally, we have the following lemma.

Lemma 2. *For all DTMCs \mathcal{M}_1 and \mathcal{M}_2 , if $\mathcal{M}_1 \equiv \mathcal{M}_2$, then for all ω -QPCTL state formulas ϕ we have $[\phi]_{\mathcal{M}_1} = [\phi]_{\mathcal{M}_2}$.*

Model checking ω -PCTL for DTMCs. The model checking algorithm for ω -PCTL for DTMCs is as follows. Given a DTMC \mathcal{M} the set of closed recurrent sets of states in \mathcal{M} can be computed in linear time by computing the maximal strongly connected components of $G_{\mathcal{M}}$ [5]. From the proof of Lemma 1 it follows that once the set of closed recurrent set of states in \mathcal{M} is computed, the computation of an ω -PCTL formula can be reduced to a PCTL formula. The model checking algorithm for QPCTL formulas on DTMCs is very similar to CTL model checking on graphs, and the CTL like model checking algorithm is applied on the graph of the DTMC. The model checking of PCTL for DTMCs can be solved in polynomial time [6] by solving a set of linear constraints. Thus we have the following result.

Theorem 1. *Given a DTMC \mathcal{M} and an ω -PCTL state formula ϕ , the following assertions hold: (1) the set $[\phi]_{\mathcal{M}}$ can be computed in time polynomial in $|\mathcal{M}|$ times ℓ ; (2) if ϕ is an ω -QPCTL formula, then the set $[\phi]_{\mathcal{M}}$ can be computed in $O(|\mathcal{M}| \cdot \ell)$ time; where $|\mathcal{M}|$ denotes the size of \mathcal{M} and ℓ denotes the length of ϕ .*

Reduction to existential theory of reals. We now present a reduction of the model checking problem for DTMCs with ω -PCTL formulas to the existential theory of reals, which is decidable in PSPACE [2]. The reduction will be later useful for model checking algorithms for IDTMCs under the PUMC and UMC semantics. Since the model checking of DTMCs for ω -PCTL formulas can be done in polynomial time and the NP-complete SAT problem can be reduced to the existential theory of reals, it follows that the model checking problem of DTMCs with ω -PCTL formulas can be reduced to the existential theory of reals. Formally, for all DTMCs \mathcal{M} , for all ω -PCTL formulas ϕ , for all states s of \mathcal{M} , there is a formula $\Gamma(\mathcal{M}, \phi, s)$ in the existential theory of reals such that (a) $\Gamma(\mathcal{M}, \phi, s)$ is true iff $s \models_{\mathcal{M}} \phi$, (b) $\Gamma(\mathcal{M}, \phi, s)$ is polynomial in size in \mathcal{M} and ϕ ; (c) $\Gamma(\mathcal{M}, \phi, s)$ can be constructed in polynomial time in size of \mathcal{M} and ϕ .

Here we make an important observation. For two DTMCs \mathcal{M}_1 and \mathcal{M}_2 , if $\mathcal{M}_1 \equiv \mathcal{M}_2$, then $\Gamma(\mathcal{M}_1, \phi, s)$ and $\Gamma(\mathcal{M}_2, \phi, s)$ have the same structure in which the transition probabilities only differ. However, the converse is not true. This important observation makes the model checking algorithms for PUMC and UMC different—the UMC model checking algorithm gets more complex.

5 PUMC Model Checking

We first present a polynomial time model checking algorithm for ω -QPCTL for PUMC interpretation of IDTMCs. We then present a PSPACE model checking algorithm for ω -PCTL for PUMC interpretation of IDTMCs, and show that the problem is both NP-hard and coNP-hard. The algorithms exploit the fact that for an IDTMC \mathcal{I} and for all $\mathcal{M}_1, \mathcal{M}_2 \in [\mathcal{I}]_P$, we have $\mathcal{M}_1 \equiv \mathcal{M}_2$.

Model checking ω -QPCTL. Given an IDTMC \mathcal{I} , all the DTMCs in the PUMC interpretation of \mathcal{I} are graph equivalent. Formally, for all $\mathcal{M}_1, \mathcal{M}_2 \in [\mathcal{I}]_P$ we have $\mathcal{M}_1 \equiv \mathcal{M}_2$. The above observation and Lemma 2 lead directly to the following model checking algorithm: given an IDTMC \mathcal{I} , pick a DTMC $\mathcal{M}_1 \in [\mathcal{I}]_P$, then for all ω -QPCTL state formulas ϕ , we have $[\phi]_{\mathcal{M}_1} = [\phi]_{\mathcal{I}_P}$. This is because for all $\mathcal{M}_2 \in [\mathcal{I}]_P$ we have $[\phi]_{\mathcal{M}_1} = [\phi]_{\mathcal{M}_2}$. Thus we obtain a polynomial time model checking algorithm for PUMC semantics for ω -QPCTL, by just picking a DTMC \mathcal{M}_1 from $[\mathcal{I}]_P$ and model checking \mathcal{M}_1 .

Theorem 2. *Given an IDTMC \mathcal{I} and an ω -QPCTL state formula ϕ_Q , the set $[\phi_Q]_{\mathcal{I}_P}$ can be computed in $O(|\mathcal{I}| \cdot \ell)$ time, where $|\mathcal{I}|$ denotes the size of \mathcal{I} and ℓ denotes the length of the formula ϕ_Q .*

Model checking ω -PCTL. We will now present a PSPACE model checking algorithm for ω -PCTL. The result is obtained by reduction to the existential theory of reals, and using the PSPACE decision procedure for the existential theory of reals [2]. Recall that for a DTMC \mathcal{M} , an ω -PCTL formula ϕ and a state

s of \mathcal{M} , there is a formula $\Gamma(\mathcal{M}, \phi, s)$ in the existential theory of reals such that $s \models_{\mathcal{M}} \phi$ if and only if $\Gamma(\mathcal{M}, \phi, s)$ is true; moreover, $\Gamma(\mathcal{M}, \phi, s)$ is polynomial in the size of \mathcal{M} and length of ϕ , and $\Gamma(\mathcal{M}, \phi, s)$ can be constructed in polynomial time. Given an IDTMC $\mathcal{I} = (S, \check{\mathbf{P}}, \hat{\mathbf{P}}, L)$, consider values $0 \leq p_{s,s'} \leq 1$ for all $s, s' \in S$ such that (a) $\check{\mathbf{P}}(s, s') \leq p_{s,s'} \leq \hat{\mathbf{P}}(s, s')$, for all $s, s' \in S$; and (b) $\sum_{s' \in S} p_{s,s'} = 1$, for all $s \in S$. Let us denote \mathbf{p} for all the values $p_{s,s'}$. We denote by $\mathcal{I}(\mathbf{p}) = (S, \mathbf{P}, L)$ the DTMC obtained by assigning $p_{s,s'}$ for the transition probability $\mathbf{P}(s, s')$. Given an IDTMC \mathcal{I} , an ω -PCTL formula ϕ and a state s of \mathcal{I} , we first observe that $s \in [\phi]_{\mathcal{I}_P}$ if and only if for all $\mathcal{M} \in [\mathcal{I}]_P$ we have $s \in [\phi]_{\mathcal{M}}$, i.e., in other words, $s \notin [\phi]_{\mathcal{I}_P}$ if and only if there is a DTMC $\mathcal{M} \in [\mathcal{I}]_P$ such that $s \models \neg\phi$. Thus for an IDTMC $\mathcal{I} = (S, \check{\mathbf{P}}, \hat{\mathbf{P}}, L)$, an ω -PCTL formula and a state s we obtain a formula $\Phi(\mathcal{I}, \phi, s)$ in the existential theory of reals such that $s \notin [\phi]_{\mathcal{I}_P}$ if and only if $\Phi(\mathcal{I}, \phi, s)$ is true. The formula $\Phi(\mathcal{I}, \phi, s)$ is as follows:

$$\Phi(\mathcal{I}, \phi, s) = \exists \mathbf{p}. \bigwedge_{s,s' \in S} (\check{\mathbf{P}}(s, s') \leq p_{s,s'} \leq \hat{\mathbf{P}}(s, s')) \wedge \bigwedge_{s \in S} (\sum_{s' \in S} p_{s,s'} = 1) \\ \bigwedge_{s,s' \in S} (\hat{\mathbf{P}}(s, s') > 0 \Rightarrow p_{s,s'} > 0) \wedge \bigwedge \Gamma(\mathcal{I}(\mathbf{p}), \neg\phi, s)$$

The first two sets of constraints specify the transition probability restriction on \mathbf{p} such that \mathbf{p} represents a valid probability transition for $\mathcal{M} \in [\mathcal{I}]$. The third set of constraints specify that if $\hat{\mathbf{P}}(s, s') > 0$, then $p_{s,s'} > 0$, and thus ensures that \mathbf{p} represents a valid probability transition for $\mathcal{M} \in [\mathcal{I}]_P$. The last constraint specifies that the DTMC $\mathcal{I}(\mathbf{p})$ satisfies $\neg\phi$ at s . Note that the formula $\Gamma(\mathcal{I}(\mathbf{p}), \neg\phi, s)$ has the same form for all $\mathcal{M} \in [\mathcal{I}]_P$, because for all $\mathcal{M}_1, \mathcal{M}_2 \in [\mathcal{I}]_P$, $\mathcal{M}_1 \equiv \mathcal{M}_2$. This is not the case if $\bigwedge_{s,s' \in S} (\hat{\mathbf{P}}(s, s') > 0 \Rightarrow p_{s,s'} > 0)$ does not hold as in UMC. Therefore, this model checking algorithm is not applicable for UMCs. Since the existential theory of reals can be decided in PSPACE [2], we have the following theorem.

Theorem 3. *Given an IDTMC \mathcal{I} and an ω -PCTL state formula ϕ , the set $[\phi]_{\mathcal{I}_P}$ can be computed in space polynomial in size of \mathcal{I} times the length of ϕ .*

Hardness of PCTL model checking. We next demonstrate the intractability of the model checking problem for PUMC by reducing the satisfiability and validity of propositional boolean formulas to the model checking problem. Consider a propositional boolean formula φ over the propositions $\{p_1, \dots, p_m\}$. We consider the UMC $\mathcal{I} = (S, \check{\mathbf{P}}, \hat{\mathbf{P}}, L)$ where

- $S = \{s_I, s_1, \dots, s_m, s_{\perp}\}$
- $L(s_I) = L(s_{\perp}) = \{\}$, $L(s_i) = \{p_i\}$ for each $1 \leq i \leq m$
- $\check{\mathbf{P}}(s_I, s_i) = 1/m^3$ and $\hat{\mathbf{P}}(s_I, s_i) = 1/m$ for all $1 \leq i \leq m$
- $\check{\mathbf{P}}(s_I, s_{\perp}) = 1/m^3$ and $\hat{\mathbf{P}}(s_I, s_{\perp}) = 1$
- $\check{\mathbf{P}}(s_i, s_i) = \hat{\mathbf{P}}(s_i, s_i) = 1$ for all $1 \leq i \leq m$
- $\check{\mathbf{P}}(s_i, s_j) = \hat{\mathbf{P}}(s_i, s_j) = 0$ for all $1 \leq i \leq m$ and $1 \leq j \leq m$ and $i \neq j$
- $\check{\mathbf{P}}(s_{\perp}, s_{\perp}) = \hat{\mathbf{P}}(s_{\perp}, s_{\perp}) = 1$

We consider the PCTL formula ϕ' obtained from ϕ by syntactically replacing every occurrence of p_i in ϕ by $\mathcal{P}_{> \frac{1}{2m}}(\mathbf{X}p_i)$ for $1 < i < m$.

Lemma 3. *The following assertions hold: (a) φ is satisfiable iff $s_I \in [\neg\phi]_{\mathcal{I}_P}$; and (b) φ is valid iff $s_I \in [\phi]_{\mathcal{I}_P}$.*

Proof. Suppose φ is satisfiable and let a be the satisfying assignment. Consider the DTMC \mathcal{M}^a , where $\mathbf{P}(s_I, s_i) = \frac{1}{2m}$ if $a(p_i) = \text{false}$ and $\mathbf{P}(s_I, s_i) = \frac{1}{m+1}$ if $a(p_i) = \text{true}$; $\mathbf{P}(s_I, s_\perp)$ is thus determined by this assignment. It is easy to see that $\mathcal{M}^a \in [\mathcal{I}]$ and $\mathcal{M}^a \models \phi$. Similarly, if $\mathcal{M} \in [\mathcal{I}]$ such that $\mathcal{M} \models \phi$, then we can construct a satisfying assignment for φ : $a(p_i) = \text{false}$ if $\mathbf{P}(s_I, s_i) \leq \frac{1}{2m}$ and $a(p_i) = \text{true}$ if $\mathbf{P}(s_I, s_i) > \frac{1}{2m}$. These observations also imply that φ is valid iff $s_I \in [\phi]_{\mathcal{I}_P}$. \square

Since the satisfiability of general propositional boolean formulas is NP-hard and the validity of general propositional boolean formulas is coNP-hard [10], the lower bounds follow immediately from Lemma 3.

Theorem 4. *Given an IDTMC \mathcal{I} , a PCTL formula ϕ , and a state s of \mathcal{I} the decision problem of whether $s \in [\phi]_{\mathcal{I}_P}$ is NP-hard and coNP-hard.*

6 UMC Model Checking

In this section we present a PSPACE model checking algorithm for UMC semantics. The PSPACE algorithm is obtained by a reduction to PUMC model checking. The basic reduction is obtained by partitioning the set of DTMCs $[\mathcal{I}]$ of IDTMC \mathcal{I} into several PUMCs.

Partitioning $[\mathcal{I}]$ of an IDTMC \mathcal{I} . Given an IDTMC $\mathcal{I} = (S, \check{\mathbf{P}}, \hat{\mathbf{P}}, L)$, let $\mathcal{B} = \{(s, s') \mid s, s' \in S, \check{\mathbf{P}}(s, s') = 0 \text{ and } \hat{\mathbf{P}}(s, s') > 0\}$ be the set of transitions that have a positive upper bound and the lower bound is 0. We consider the following set of IDTMCs \mathcal{I}^B for $B \subseteq \mathcal{B}$: we have $\mathcal{I}^B = (S, \check{\mathbf{P}}, \hat{\mathbf{P}}^B, L)$ such that $\hat{\mathbf{P}}^B(s, s') = 0$ if $(s, s') \in B$; and $\hat{\mathbf{P}}(s, s')$ otherwise. In other words, in \mathcal{I}^B the upper of the transition probabilities for the set B is set to 0, and otherwise it behaves like \mathcal{I} . The key partitioning property is as follows: $[\mathcal{I}] = \bigcup_{B \subseteq \mathcal{B}} [\mathcal{I}^B]_P$, i.e., the union of the DTMCs obtained from the PUMCs semantics of \mathcal{I}^B is the set of DTMCs obtained from the UMC semantics of $[\mathcal{I}]$. Thus we obtain that for all ω -PCTL formulas ϕ we have $[\phi]_{\mathcal{I}} = \bigcap_{B \subseteq \mathcal{B}} [\phi]_{\mathcal{I}^B}$.

Model checking ω -QPCTL. The model checking problem for IDTMCs for ω -QPCTL formulas under UMC semantics can be solved in coNP. Given an IDTMC \mathcal{I} , an ω -QPCTL formula ϕ , and a state s , to show that $s \notin [\phi]_{\mathcal{I}}$, it suffices to guess $B \subseteq \mathcal{B}$ and prove that $s \notin [\phi]_{\mathcal{I}^B}$. Hence the guess (or the witness) is B and Theorem 2 provides the polynomial time verification procedure. Hence we obtain the following theorem.

Theorem 5. *Given an IDTMC \mathcal{I} , an ω -QPCTL state formula ϕ_Q , and a state s of \mathcal{I} whether $s \in [\phi_Q]_{\mathcal{I}}$ can be decided in coNP.*

Model checking ω -PCTL. Similar to the model checking algorithm for the ω -QPCTL, we can obtain a NPSpace model checking algorithm, by guessing $B \subseteq \mathcal{B}$ and then using the PSPACE model checking algorithm for ω -PCTL for PUMC semantics.

Theorem 6. *Given an IDTMC \mathcal{I} and an ω -PCTL state formula ϕ , the set $[\phi]_{\mathcal{I}}$ can be computed in space polynomial in size of \mathcal{I} times the length of ϕ .*

Hardness of PCTL model checking. The hardness result follows from the result for PUMC. In the hardness proof for PUMC, the IDTMCs \mathcal{I} considered satisfied that $[\mathcal{I}] = [\mathcal{I}]_P$; and hence the UMC and PUMC semantics coincide for \mathcal{I} . This gives us the following result.

Theorem 7. *Given an IDTMC \mathcal{I} , a PCTL formula ϕ , and a state s of \mathcal{I} the decision problem of whether $s \in [\phi]_{\mathcal{I}}$ is NP-hard and coNP-hard.*

7 IMDP Model Checking

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). Before presenting this reduction we recall some basic properties of the feasible solutions of a linear program and the definition of an MDP.

Linear programming. Consider an IMDP $\mathcal{I} = (S, \check{\mathbf{P}}, \hat{\mathbf{P}}, L)$. For a given $s \in S$, let $IE(s)$ be the following set of inequalities over the variables $\{p_{ss'} \mid s' \in S\}$: $\sum_{s' \in S} p_{ss'} = 1$, where $\check{\mathbf{P}}(s, s') \leq p_{ss'} \leq \hat{\mathbf{P}}(s, s')$ 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 a linear program has 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 $IE(s)$. Then 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 4. *The number of basic feasible solutions of $IE(s)$ in the worst case can be $O(|S|2^{|S|-1})$.*

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 $\sum_{s \in S} \mu(s) = 1$.*

Definition 6. A Markov decision process (MDP) is a 3-tuple $\mathcal{D} = (S, \tau, L)$, where (1) S is a finite set of states; (2) $L: S \rightarrow 2^{\text{AP}}$ is a labeling function that maps states to sets of atomic propositions from a set AP; and (3) 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 $\mathcal{D} = (S, \tau, L)$ 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 history dependent 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 history dependent adversaries of the MDP \mathcal{D} . An adversary is memoryless if it is independent of the history and only depends on the current state. 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 an MDP \mathcal{D} , an ω -PCTL formula φ , and a state s we can define when $s \models_{\mathcal{D}} \varphi$ in a way analogous to the IMDPs.

The reduction. We are now ready to describe the model checking algorithm for IMDPs. Consider an IMDP $\mathcal{I} = (S, \mathbf{P}, \mathbf{P}, L)$. Recall from the description of linear programming that 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 the set of all BFS of $IE(s)$. Define the following MDP $\mathcal{D} = (S', \tau, L')$ where $S' = S$, $L' = L$, 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 4). 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. Theorem 8 follows from the following proposition.

Proposition 2. For any adversary A for \mathcal{I} , we can define a randomized adversary A' such that $\text{Prob}_s^{\mathcal{I}^A} = \text{Prob}_s^{\mathcal{D}^{A'}}$ for every s , where $\text{Prob}_s^{X^A}$ is measure on

paths from s defined by X under A . Similarly for every adversary A for \mathcal{D} , there is an adversary A' for \mathcal{I} that defines the same probability measure on paths.

Theorem 8. *Given an IMDP \mathcal{I} , for all ω -PCTL formulas φ and for all states s , we have $s \models_{\mathcal{I}} \varphi$ iff $s \models_{\mathcal{D}} \varphi$.*

Thus, in order to model check IMDP \mathcal{I} , we can model check the MDP \mathcal{D} . The model checking algorithm for MDPs requires the solution of MDPs with infinitary path formulas, and solution of MDPs with PCTL formulas. Algorithms that run in polynomial time (and space) for MDPs with Büchi and coBüchi conditions are known from [4, 7], and it is straightforward to extend the algorithms to infinitary path formulas that are obtained as conjunction and disjunction of Büchi and coBüchi conditions. Algorithms that run in polynomial time (and space) for MDPs with PCTL formulas are available in [1, 17]. Thus, if we directly model check \mathcal{D} we get an EXPTIME model checking algorithm for \mathcal{I} . However, we can improve this to get a coNP procedure. The reason for this is that it is known that as far as model checking MDPs is concerned, we can restrict our attention to certain special class of *memoryless* adversaries, i.e., adversaries that always pick a fixed probability distribution over a set of non-deterministic choices whenever a state is visited. It follows from the results of [3] that in MDPs with Müller conditions (that subsumes the infinitary path formulas of ω -PCTL) an uniform randomized memoryless optimal strategy exists such that the size of the support of the memoryless optimal strategy is bounded by the size of the state space. Formally, we have the following lemma.

Lemma 5. *For an MDP $\mathcal{D} = (S, \tau, L)$ and an infinitary path formula ψ^ω , there exists an randomized memoryless adversary A such that (1) (Support of size at most $|S|$). for all $s \in S$ we have $|Supp(A(s))| \leq |S|$; (2) (Uniform). for all $s \in S$ and $\mu \in Supp(A(s))$ we have $A(s)(\mu) = \frac{1}{|Supp(A(s))|}$; and (3) (Optimal). for all $s \in S$ we have $Prob_s^{\mathcal{D}^A}(\psi^\omega) = \sup_{A' \in \mathcal{A}} Prob_s^{\mathcal{D}^{A'}}(\psi^\omega)$.*

The existence of deterministic memoryless strategies for formulas in PCTL (where the sub-formulas are already evaluated) for MDPs follows from the results of [1, 17]. Thus we obtain the following theorem.

Proposition 3 ([1, 17, 3]). *Let $\mathcal{D} = (S, \tau, L)$ be an MDP. Let \mathcal{A}_{unf} be the set of uniform randomized memoryless adversaries with support of size at most $|S|$ for MDP \mathcal{D} , i.e., for all $A \in \mathcal{A}_{unf}$, $A(s)(\mu) = \frac{1}{|Supp(A(s))|}$ for $\mu \in Supp(A(s))$ and $|Supp(A(s))| \leq |S|$. Consider an ω -PCTL formula $\varphi = \mathcal{P}_{\triangleright\triangleleft p}(\psi)$ 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 $A \in \mathcal{A}_{unf}$.*

For every subformula of the form $\varphi = \mathcal{P}_{\triangleright\triangleleft p}(\psi)$, if the formula φ is not true at a state s , in the IMDP semantics, then we can guess $A \in \mathcal{A}_{unf}$ and then verify that in \mathcal{D}^A the formula φ is not true at s . The witness A is the polynomial witness and the polynomial time algorithm for Markov chains presents the polynomial time verification procedure. In case of general formulas, the above procedure needs to be applied in a bottom up fashion.

Theorem 9. *Given an IDTMC \mathcal{I} and an ω -PCTL state formula ϕ , and a state s , whether the state $s \models \phi$ under the IMDP semantics can be decided in coNP .*

Lower bound. It follows from the results of [6] that the model checking problem for DTMCs with PCTL formulas is PTIME-hard. Since DTMCs are a special case of IMDPs, the PTIME-time lower bound follows for model checking IMDPs with PCTL and ω -PCTL formulas.

8 Model Checking of Linear Time Formulas

Finally, we consider the model checking problem of IDTMCs with LTL formulas. In other words, we consider LTL path formulas ψ , and formulas of the form $\mathcal{P}_{\triangleright\triangleleft p}(\psi)$. For the model checking problem we apply the following procedure: we first convert ψ to an equivalent non-deterministic Büchi automata [20], and then determinize it to obtain an equivalent deterministic Rabin automata $Q(\psi)$ [16]. The deterministic Rabin automata $Q(\psi)$ has 2^{2^l} states, where l is the length of the formula ψ , and has 2^l Rabin pairs. Given a IDTMC \mathcal{I} and a formula $\varphi = \mathcal{P}_{\triangleright\triangleleft p}(\psi)$, the model checking problem for the UMC and IMDP semantics are solved as follows. In both case we construct the Rabin automata $Q(\psi)$.

1. For the IMDP semantics, we construct the product IDTMC of \mathcal{I} and $Q(\psi)$, denoted as $\mathcal{I} \times Q(\psi)$, and solve it under the IMDP semantics with respect to a Rabin objective (applying the results of Section 7).
2. For the PUMC semantics, we construct the product IDTMC of \mathcal{I} and $Q(\psi)$, denoted as $\mathcal{I} \times Q(\psi)$. For the formula φ , we write a formula in the existential theory of reals: the formula is similar to the formula of Section 5 with the additional constraints that for two states in $\mathcal{I} \times Q(\psi)$, if the state component of \mathcal{I} is the same, then the chosen distribution at the states must also be same, i.e., for two states (s, q_1) and (s, q_2) we require that the probability distribution chosen from the interval must be the same. The result for UMC semantics is similar. We thus obtain the following result.

Theorem 10. *Given an IDTMC \mathcal{I} , an LTL path formula ψ , and a state formula $\phi = \mathcal{P}_{\triangleright\triangleleft p}(\psi)$, the following assertions hold: (1) the sets $[\phi]_{\mathcal{I}_P}$ and $[\phi]_{\mathcal{I}}$ can be computed in PSPACE in the size of \mathcal{I} and 2EXPTIME in the length of the formula ψ ; and (2) given a state s , whether the state $s \models \phi$ under the IMDP semantics can be decided in coNP in the size of \mathcal{I} and 2EXPTIME in the length of the formula ψ .*

9 Conclusion

We have investigated the model checking problem of ω -PCTL and its qualitative fragment for three semantic interpretations of IDTMCs, namely UMC, PUMC and IMDP. We proved upper bounds and lower bounds on the complexity of the model checking problem for these models. Some of our bounds however are not tight. Finding tight lower and upper bounds for these model checking problems is an interesting open problem. We also present model checking algorithm for LTL formulas.

References

1. A. Bianco and L. de Alfaro. Model checking of probabilistic and nondeterministic systems. In *FSTTCS'95*, vol. 1026 *LNCS*, Springer, 1995.
2. J. Canny. Some algebraic and geometric computations in PSPACE. In *STOC'88*, pp 460–467, ACM, 1988.
3. K. Chatterjee, L. de Alfaro, and T. Henzinger. Trading memory for randomness. In *QEST'04*, IEEE, 2004.
4. K. Chatterjee, M. Jurdziński, and T. Henzinger. Quantitative stochastic parity games. In *SODA'04*, ACM-SIAM, 2004.
5. T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms, Second Edition*. MIT Press and McGraw-Hill, 1990.
6. C. Courcoubetis and M. Yannakakis. The complexity of probabilistic verification. *Journal of ACM*, 42(4):857–907, 1995.
7. L. de Alfaro. *Formal Verification of Probabilistic Systems*. PhD thesis, Stanford University, 1997.
8. H. Fecher, M. Leucker, and V. Wolf. Don't know in probabilistic systems. In *SPIN'06*, vol. 3925 *LNCS*, Springer, 2006.
9. H. Hansson and B. Jonsson. A logic for reasoning about time and reliability. *Formal Aspects of Computing*, 6(5):512–535, 1994.
10. J. E. Hopcroft and J. D. Ullman. *Introduction to Automata Theory, Languages and Computation*. Addison-Wesley, Reading, MA, 1979.
11. B. Jonsson and K. G. Larsen. Specification and refinement of probabilistic processes. In *LICS'91*, IEEE, 1991.
12. J. Kemeny, J. Snell, and A. Knapp. *Denumerable Markov chains*. Springer, 1976.
13. I. O. Kozine and L. V. Utkin. Interval-valued finite Markov chains. *Reliable Computing*, 8(2):97–113, 2002.
14. V. P. Kuznetsov. Interval statistical models. *Radio and Communication*, 1991.
15. J. Rutten, M. Kwiatkowska, G. Norman, and D. Parker. *Mathematical Techniques for Analyzing Concurrent and Probabilistic Systems*, American Mathematical Society, 2004.
16. S. Safra. *Complexity of automata on infinite objects*. PhD thesis, Weizmann Institute of Science, 1989.
17. R. Segala. *Modeling and Verification of Randomized Distributed Real-Time Systems*. PhD thesis, MIT, 1995.
18. K. Sen, M. Viswanathan, and G. Agha. Model-checking Markov chains in the presence of uncertainties. In *TACAS'06*, vol. 3920 *LNCS*, Springer, 2006.
19. W. Thomas. Languages, automata, and logic. In *Handbook of Formal Languages*, volume 3, Beyond Words, chapter 7, Springer, 1997.
20. M. Vardi and P. Wolper. An automata-theoretic approach to automatic program verification. In *LICS'86*, IEEE, 1986.
21. M. Y. Vardi. Automatic verification of probabilistic concurrent finite-state programs. In *FOCS'85*, IEEE, 1985.
22. P. Walley. Measures of uncertainty in expert systems. *Artificial Intelligence*, 83:1–58, 1996.