Information Processing Letters 113 (2013) 210–216

Contents lists available at SciVerse ScienceDirect

Information Processing Letters

www.elsevier.com/locate/ipl

On the complexity of model checking interval-valued discrete time Markov chains ☆

Taolue Chen *, Tingting Han, Marta Kwiatkowska

Department of Computer Science, University of Oxford, United Kingdom

ARTICLE INFO

Article history:
Received 29 June 2012
Received in revised form 17 December 2012
Accepted 6 January 2013
Available online 9 January 2013
Communicated by A. Muscholl

Keywords:
Formal methods
Discrete time Markov chain
Reachability
PCTL
Complexity
Linear programming

ABSTRACT

We investigate the complexity of model checking (finite) interval-valued discrete time Markov chains, that is, discrete time Markov chains where each transition is associated with an interval in which the actual transition probability must lie. Two semantics are considered, the uncertain Markov chain (UMC) semantics and the interval Markov decision process (IMDP) semantics. We show that, for reachability, these two semantics coincide and the problem is P-complete. This entails that PCTL model checking problem under the IMDP semantics is also P-complete. We also show that model checking PCTL under the UMC semantics is Square-Root-Sum hard, meaning that one can reduce the Square-Root-Sum problem to it.

1. Introduction

Discrete time Markov chains (DTMCs) are a well-established stochastic model which has numerous applications in a wide range of areas such as physics, chemistry, biology, engineering, etc. A DTMC usually consists of a set of states and a *fixed* transition probability between each pair of states. However, in practice it is often not realistic to assume that transition probabilities are exact. In particular, these probabilities are usually estimated by statistical experiments, which can only give bounds instead of fixed values. *Interval-valued discrete time Markov chains* (IDTMCs) have been introduced [15,18] to faithfully capture this type of uncertainty. IDTMCs are DTMCs where each transition probability is assumed to be within a range (interval). Two different semantics for IDTMCs are discussed in [24], i.e., the *uncertain Markov chain* (UMC) semantics and the *interval Markov decision process* (IMDP) semantics. The former [18] is an interpretation of an IDTMC as a family of (possibly uncountably many) DTMCs, each of which is a DTMC whose transition probabilities lie within the specified interval. As for the latter, the uncertainty is resolved similarly to *Markov decision processes* (MDPs), namely, each time a state is visited, a transition distribution which respects the interval constraints is selected, and then a probabilistic step according to the chosen distribution is taken. Thus, IMDPs allow the modelling of nondeterministic choices from a set of (possibly) uncountably many such choices. An IMDP can be seen as a generalisation of MDPs (which normally assume only finitely many choices).

We investigate the problem of model checking *probabilistic computation tree logic* (PCTL, [14]) specifications for IDTMCs. In [24,7] various complexity results have been obtained; in particular, it is shown that the PCTL model checking problem is Np- and coNp-hard, and in Pspace under the UMC semantics, while it is P-hard and in coNp under the IMDP semantics. An obvious question is how to fill these complexity gaps.

1.1. Contributions

In this paper, we either improve the upper bound obtained in [24,7], obtaining matching upper and low

☆ This work is supported by the ERC Advanced Grant VERIWARE.
* Corresponding author.
E-mail address: taolue.chen@gmail.com (T. Chen).

0020-0190/$ – see front matter
http://dx.doi.org/10.1016/j.ipl.2013.01.004

complexity bounds, or provide "evidence" that certain problems are intractable. In particular, we obtain the following results:

- We show that the two semantics for IDTMCs coincide for *reachability problems*; further we show that the problem is P-complete.
- We show that the PCTL model checking problem is P-complete under the IMDP semantics, while it is Square-Root-Sum hard under the UMC semantics.

The first result tightens the complexity bounds achieved in [24,7]. It can also be used for [11,17], where only heuristic approaches are proposed which cannot guarantee polynomial running time. We discuss briefly the techniques introduced here. Following [24,7], the reachability problem for IDTMCs can be reduced to solving the same problem for MDPs with exponentially many nondeterministic choices, which in turn can be reduced to solving a linear programming (LP) problem with exponentially many constraints. We then apply the ellipsoid method to solve the obtained LP. The crux of this approach is to find a polynomial-time separation oracle which will be made clear in Section 3. Based on this, the ellipsoid method yields a polynomial-time algorithm. Together with an easy reduction from the reachability problem for MDPs showing P-hardness, we obtain P-completeness. A direct application of this result implies that model checking PCTL under the IMDP semantics is also P-complete.

As for PCTL model checking under the UMC semantics, we reduce the Square-Root-Sum problem to it, by slightly adapting a construction in [5]. The Square-Root-Sum problem is known to be in Pspace, but its containment even in Np has been a long-standing open problem since 1976. (This problem arises often and has been studied extensively, especially in computational geometry where the square root sum represents the sum of Euclidean distances between given pairs of points with integer/rational coordinates; as an example, determining whether the length of a TSP tour of a set of points on the plane is bounded by a given threshold can be easily encoded as the Square-Root-Sum problem.) In [1] it is shown that this problem can be decided in the 4-th level of the Counting Hierarchy (an analogue of the polynomial-time hierarchy for counting classes), hence it is unlikely to be Pspace-hard, but it remains open whether the problem can be decided in P or even in Np. Our result suggests that the Pspace upper bound in [24,7] cannot be substantially improved without a breakthrough concerning this long-standing open problem. Interesting examples of this type of argument in formal verification can also be found in, among others, [10].

1.2. Related work

IDTMCs are probably the simplest stochastic model addressing uncertainties. More general models indeed exist, for example, constraint Markov chains [6], bounded-parameter MDPs [12], and uncertain MDPs [19,26]. Most of these models are from AI where *discounted total reward* objectives are prevailing. In this case, assuming the discounting factor is given as a constant (which is a usual and reasonable assumption), polynomial-time algorithms to compute the optimal value do exist, given that uncertainty is modelled appropriately (e.g. as an interval considered here; cf. [12,26]). In contrast, specification logics like PCTL, whose core component is the *reachability* property, are the main object of the research in verification. We also mention [9], where the complexity of thorough refinement and deciding the existence of a common implementation for an unbounded number of IDTMCs, etc., is studied. These problems stem from compositional modelling methodologies, which are not the main focus of our work. Furthermore, [11,17] use IDTMCs as abstractions for probabilistic systems (DTMCs, CTMCs). Model checking PCTL under the IMDP semantics is essentially adopted there, but no polynomial-time algorithms are given. As mentioned, our results can be applied directly in these settings.

1.3. Structure

This paper is organised as follows. Section 2 provides the necessary background. Section 3 presents the result on reachability problems, and Section 4 presents the results for PCTL. The paper is concluded in Section 5.

2. Preliminaries

Given any finite set S, we write $\Delta(S)$ for the set of *probabilistic distributions* over S, i.e., functions $\mu : S \to [0,1]$ with $\sum_{s\in S} \mu(s) = 1$. Throughout the whole paper, we assume a finite set of *atomic propositions* AP.

2.1. Interval DTMCs

Definition 1 *(DTMC).* A *discrete time Markov chain* (DTMC) is a tuple $\mathcal{D} = (S, s_0, \mathbf{P}, L)$, where

- S is a finite set of states with $s_0 \in S$ as the initial state;
- $\mathbf{P} : S \times S \to [0,1]$ is a transition probability matrix, s.t. $\forall s \in S$, $\sum_{s'\in S} \mathbf{P}(s,s') = 1$; and
- $L : S \to 2^{\mathrm{AP}}$ is a labelling function which maps states to sets of atomic propositions.

A *path* in $\mathcal{D}$ is a sequence $\pi = s_0 s_1 \cdots$ such that $s_i \in S$ and $\mathbf{P}(s_i, s_{i+1}) > 0$ for any $i \geqslant 0$. We write $\mathit{Paths}(s)$ for the set of paths starting at the state s. We can also define a *probability distribution* Pr over $\mathit{Paths}(s_0)$ in a standard way (cf. [2, Chapter 10]).

Definition 2 *(MDP).* A *Markov decision process* (MDP) is a tuple $\mathcal{M} = (S, s_0, \delta, L)$, where

- S, s_0, and L are defined as before; and
- $\delta : S \to 2^{\Delta(S)}$ is a transition function, s.t. $\forall s \in S$, $\delta(s)$ is *finite*.

Without loss of generality, we assume that $\delta(s) \neq \emptyset$ for any $s \in S$. In $\mathcal{M}$, at each state s a probability distribution μ (over S) is chosen *nondeterministically* from the set $\delta(s)$. A successor state s' is then chosen according to μ with probability $\mu(s')$.

A path π in $\mathcal{M}$ is a sequence of the form $s_0 \xrightarrow{\mu_1} s_1 \xrightarrow{\mu_2} \cdots$ where $s_i \in S$, $\mu_{i+1} \in \delta(s_i)$ and $\mu_{i+1}(s_{i+1}) > 0$ for each $i \geqslant 0$. A finite path is a prefix of an infinite path ending in a state. Let $Paths^\star$ be the set of finite paths. A *scheduler* $\sigma : Paths^\star \to \Delta(S)$ maps a finite path (*history*) to a distribution over S. In particular, a *simple* scheduler σ chooses a distribution only based on the current state and $\sigma(s) \in \delta(s)$ for each state s. A (memoryless) *randomised* scheduler σ prescribes, for each state s, a distribution ν over $\delta(s)$, which induces a distribution over S as $\sum_{\mu\in\delta(s)} \nu(\mu)\cdot\mu(s')$ for each $s' \in S$. Note that we may obtain a DTMC by resolving all the nondeterminism in an MDP using a scheduler σ in a standard way (see, e.g., [2,22]). In the sequel, we write $\mathcal{M}_\sigma$ for such a DTMC given an MDP $\mathcal{M}$ and a scheduler σ.

Definition 3 *(IDTMC).* An *interval-valued discrete time Markov chain* (IDTMC) is a tuple $\mathcal{I} = (S, s_0, \mathbf{P}^l, \mathbf{P}^u, L)$, where

- S, s_0 and L are defined as before;
- $\mathbf{P}^l, \mathbf{P}^u : S \times S \to [0,1]$ are two transition probability matrices, where $\mathbf{P}^l(s,s')$ (resp. $\mathbf{P}^u(s,s')$) gives the *lower* (resp. *upper*) bound of the transition probability from state s to s'.

For complexity consideration, we assume that each entry of $\mathbf{P}^l$ and $\mathbf{P}^u$ is a rational number. We define the size of $\mathcal{I}$, denoted by $\sharp(\mathcal{I})$, as the size of the representation of $\mathcal{I}$. Here we represent rational numbers (probabilities) as quotients of integers written in binary. Namely, the size of a rational number is the sum of the bit lengths of its numerator and denominator and the size of a matrix is the sum of the sizes of its entries. We say that a vector is *rational* if all of its coordinates are rational numbers.

2.2. Semantics

We consider two semantic interpretations of IDTMCs [24], i.e., *uncertain Markov chains* (UMC) and *interval Markov decision processes* (IMDP).

2.2.1. UMC semantics

An IDTMC $\mathcal{I} = (S, s_0, \mathbf{P}^l, \mathbf{P}^u, L)$ represents an infinite set of DTMCs, denoted by $[\mathcal{I}]$, where for each DTMC $(S, s_0, \mathbf{P}, L) \in [\mathcal{I}]$ the following holds:

$$\mathbf{P}^l(s,s') \leqslant \mathbf{P}(s,s') \leqslant \mathbf{P}^u(s,s')$$

for all pairs of states $s, s' \in S$. Intuitively, under this semantics we assume that the external environment nondeterministically selects a DTMC from the set $[\mathcal{I}]$ at the beginning and then all the transitions take place according to the chosen DTMC. Note that, here, the external environment makes only one nondeterministic choice.

2.2.2. IMDP semantics

An IDTMC $\mathcal{I} = (S, s_0, \mathbf{P}^l, \mathbf{P}^u, L)$ defines $\lceil\mathcal{I}\rceil = (S, s_0, \delta, L)$ such that

$$\delta(s) = \left\{\mu\in\Delta(S) \;\middle|\; \forall s'\in S, \mathbf{P}^l(s,s') \leqslant \mu(s') \leqslant \mathbf{P}^u(s,s')\right\}.$$

In this case, the external environment makes nondeterministic choices *at each state* by picking a distribution for that state. Technically speaking, $\lceil\mathcal{I}\rceil$ is not an MDP as $\mu(s)$ might contain infinitely many elements. This justifies the name IMDP. However, all the properties of MDPs carry over here, and in the sequel we do not distinguish between MDPs and IMDPs unless stated explicitly.

2.3. PCTL

PCTL [14] is a probabilistic extension of CTL in which state formulas are interpreted over states of a DTMC and path formulas are interpreted over paths in a DTMC. The syntax of PCTL is as follows:

$$\Phi ::= \mathrm{tt} \mid a \mid \neg\Phi \mid \Phi \wedge \Phi \mid \mathcal{P}_{\bowtie p}(\phi)$$

where $p \in [0,1]$ is a probability, $\bowtie \in \{<, \leqslant, >, \geqslant, =, \neq\}$ and ϕ is a path formula defined according to the following grammar:

$$\phi ::= \mathsf{X}\Phi \mid \Phi \,\mathsf{U}\, \Phi.$$

The path formula $\mathsf{X}\Phi$ asserts that the state on the next position on the path satisfies Φ; while $\Phi \,\mathsf{U}\, \Psi$ asserts that Ψ is satisfied and that all preceding states satisfy Φ. We denote by $\Diamond\Phi = \mathrm{tt}\,\mathsf{U}\,\Phi$ the *reachability* property, where states satisfying Φ will eventually be reached.

2.3.1. PCTL semantics for DTMC

Let $\mathcal{D} = (S, s_0, \mathbf{P}, L)$ be a DTMC. The semantics of PCTL is defined by a satisfaction relation, denoted $\models$, which is characterised as the least relation over the states in S (paths in $\mathcal{D}$, respectively) and the state formulas (path formulas) satisfying:

$$s \models \mathcal{P}_{\bowtie p}(\phi) \quad \text{iff} \quad \Pr\left\{\pi \in Paths(s) \;\middle|\; \pi \models \phi\right\} \bowtie p,$$

while the semantics for the other operators are as in CTL. The semantics of PCTL path formulas is defined as:

$$\pi \models \Phi \,\mathsf{U}\, \Psi \quad \text{iff} \quad \exists i \geqslant 0.\left(\pi[i] \models \Psi \wedge \forall 0 \leqslant j < i.\pi[j] \models \Phi\right),$$

$$\pi \models \mathsf{X}\Phi \quad \text{iff} \quad \pi[1] \models \Phi,$$

where $\pi[i]$ denotes the $(i+1)$-th state of π. We then lift the semantics from DTMC to IDTMC.

2.3.2. PCTL semantics for UMC

Given an IDTMC $\mathcal{I}$ and a PCTL state formula Φ, we write $\mathcal{I} \models \Phi$ iff there exists some $\mathcal{D} \in [\mathcal{I}]$ such that $\mathcal{D} \models \Phi$.

2.3.3. PCTL semantics for IMDP

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

$$s \models \mathcal{P}_{\bowtie p}(\phi) \quad \text{iff} \quad \Pr^\sigma\left\{\pi \in Paths(s) \;\middle|\; \pi \models \phi\right\} \bowtie p,$$

for some scheduler σ. Note that here $\Pr^\sigma$ denotes the probability distribution over paths of the Markov chain $\lceil\mathcal{I}\rceil_\sigma$.

Remark 1. Here we follow the traditional approach as in [4,24,7] instead of [5] where MDPs are regarded as a one-player game with PCTL formulas as the winning objective. This "branching-time" view of PCTL would imply undecidability immediately.

3. Reachability

In this section, we focus on the *reachability problem*. Assume an IDTMC $\mathcal{I} = (S, s_0, \mathbf{P}^l, \mathbf{P}^u, L)$ and a set of *goal* states $T \subseteq S$. By *reachability probability to* T we refer to the probability of paths which hit T. Under the UMC semantics, we are interested in computing $\sup_{\mathcal{D}\in[\mathcal{I}]} \Pr(\mathcal{D}, \Diamond T)$ (resp. $\inf_{\mathcal{D}\in[\mathcal{I}]} \Pr(\mathcal{D}, \Diamond T)$), i.e., the maximum (resp. minimum) reachability probability among all DTMCs of $[\mathcal{I}]$. Under the IMDP semantics, we are interested in computing $\sup_\sigma \Pr(\lceil\mathcal{I}\rceil_\sigma, \Diamond T)$ (resp. $\inf_\sigma \Pr(\lceil\mathcal{I}\rceil_\sigma, \Diamond T)$), i.e., the maximum (resp. minimum) reachability probability among all schedulers for $\lceil\mathcal{I}\rceil$.

We first show that, for reachability, the two semantics coincide in the following sense (note that the sup case can be shown similarly):

Proposition 1. *Given any IDTMC* $\mathcal{I} = (S, s_0, \mathbf{P}^l, \mathbf{P}^u, L)$ *and* $T \subseteq S$*, we have that*

$$\inf_\sigma \Pr\big(\lceil\mathcal{I}\rceil_\sigma, \Diamond T\big) = \inf_{\mathcal{D}\in[\mathcal{I}]} \Pr(\mathcal{D}, \Diamond T).$$

Proof. For any $\mathcal{D} \in [\mathcal{I}]$, we construct a (simple) scheduler σ such that $\sigma(s)$ is the transition probability distribution of $\mathcal{D}$ at s. Clearly, we have that $\lceil\mathcal{I}\rceil_\sigma$ is identical to $\mathcal{D}$. We hence obtain that

$$\inf_\sigma \Pr\big(\lceil\mathcal{I}\rceil_\sigma, \Diamond T\big) \leqslant \inf_{\mathcal{D}\in[\mathcal{I}]} \Pr(\mathcal{D}, \Diamond T).$$

On the other hand, suppose that σ is the scheduler which achieves the minimum reachability probability in $\lceil\mathcal{I}\rceil$. It is well-known that one can assume that σ is a simple scheduler [22,4]. Hence, we can construct a DTMC $\mathcal{D} = \lceil\mathcal{I}\rceil_\sigma$ and clearly $\mathcal{D} \in [\mathcal{I}]$. It follows that

$$\inf_\sigma \Pr\big(\lceil\mathcal{I}\rceil_\sigma, \Diamond T\big) \geqslant \inf_{\mathcal{D}\in[\mathcal{I}]} \Pr(\mathcal{D}, \Diamond T).$$

The conclusion then follows. □

Remark 2. The UMC and IMDP semantics are different when full PCTL is considered, cf. [24].

Owing to Proposition 1, we only need to focus on the IMDP semantics, which turns out to be technically easier to tackle. Note that it also follows that, under the UMC semantics, the supremum (resp. infimum) can actually be achieved. (This is evidently true for IMDP semantics, following standard results for MDPs.)

3.1. Complexity

In the rest of this section, let us fix an IMDP (S, s_0, δ, L) as the semantics of an IDTMC $\mathcal{I} = (S, s_0, \mathbf{P}^l, \mathbf{P}^u, L)$, and we shall focus on the minimum probability of reaching $T \subseteq S$, as the maximum one can be dealt with dually. A standard approach is to reduce to linear programming (LP) which proceeds as follows.

For any state $s \in S$, let ∇_s be a polytope in $\mathbb{R}^{|S|}_{\geqslant 0}$ defined by the following constraints:

$$\sum_{s'\in S} \mu_{s'} = 1 \quad \text{and} \quad \mathbf{P}^l(s,s') \leqslant \mu_{s'} \leqslant \mathbf{P}^u(s,s'). \tag{1}$$

We write $\mathcal{V}(\nabla_s)$ for the (finite) set of *vertices* of ∇_s.

Lemma 1. *Each vertex of* ∇_s *is rational of size polynomial in* $\sharp(\mathcal{I})$.

Proof. It suffices to observe that:

- Each vertex of ∇_s is an intersection of hyperplanes given by the (in)equalities in (1); and
- The coefficients of each (in)equality are rationals bounded by $\sharp(\mathcal{I})$ in size.

The conclusion follows from [23, Theorem 10.1]. □

We consider the following LP problem:

$$\begin{aligned} \text{maximise} \quad & \sum_{s\in S\setminus T} x_s \\ \text{subject to} \quad & x_s \leqslant \sum_{s'\in S\setminus T} \mu_{s'} x_{s'} + \sum_{s'\in T} \mu_{s'}, \\ & \qquad \text{if } \mu \in \mathcal{V}(\nabla_s) \text{ and } s \notin T \\ & x_s = 1, \quad \text{if } s \in T. \end{aligned} \tag{2}$$

We remark that, in principle one should write ∇_s instead of $\mathcal{V}(\nabla_s)$ in (2) and standard results yield that x_{s_0} is the minimum reachability probability [22,4]. However, ∇_s is *infinite* and therefore the set of constraints would be infinite as well. For each fixed $(x_s)_{s\in S}$, standard result from LP yields that the linear function $\sum_{s'\in S\setminus T} \mu_{s'} x_{s'} + \sum_{s'\in T} \mu_{s'}$ achieves its minimum on ∇_s at some vertex (see, e.g., [23]). Hence we can replace ∇_s by $\mathcal{V}(\nabla_s)$, obtaining an LP with the same feasible region, in its current form.

Remark 3. In [24], the notion of *basic feasible solutions* is introduced, which is essentially $\mathcal{V}(\nabla_s)$.

One obstacle is that the LP given in (2) still has *exponentially* many constraints, as, for each s, $\mathcal{V}(\nabla_s)$ in general is of exponential size. Hence, it is not clear how a polynomial-time upper bound can be achieved immediately, as opposed to the normal MDP case. (The issue is also mentioned in [24,7].) To overcome this difficulty, we now leverage the celebrated *ellipsoid method*. A remarkable observation is that it is not necessary to have an explicit list of all constraints (in terms of (in)equalities) ready at hand; instead it suffices to be able to test if a given vector (in particular, the centre of the ellipsoid) is a solution of the constraint system, and, if not, to find one violated constraint. Historically this was observed independently by Karp and Papadimitriou [16], Padberg and Rao [20], and Grötschel, Lováze, and Schrijver [13]; we refer the reader to [23, Chapter 14] for details.

Hence, to conclude that the LP (2) admits a polynomial-time algorithm in $\sharp(\mathcal{I})$, one only needs to show that there exists a polynomial-time *separation oracle*, i.e., a polynomial-time algorithm to solve the following *separation problem*, which in our setting reads as follows.

Given $\vec{y} \in \mathbb{R}^{|S|}_{\geqslant 0}$ rational of size polynomial in $\sharp(\mathcal{I})$, decide whether $\vec{y}$ satisfies the constraints of (2) or not, and, if negative, identify in polynomial time a *separating hyperplane*, i.e., $\vec{\alpha} \in \mathbb{R}^{|S|}$ such that $\vec{\alpha} \cdot \vec{y} < \vec{\alpha} \cdot \vec{x}$ for any $\vec{x}$ which satisfies the constraints of (2). Here, $\vec{\alpha}$ should be read as a row vector while $\vec{x}$ and $\vec{y}$ are column vectors, and $\cdot$ is the standard scalar product.

Proposition 2. *The separation oracle exits.*

Proof. Let $\vec{y} \in \mathbb{R}^{|S|}_{\geqslant 0}$ be rational of size polynomial in $\sharp(\mathcal{I})$. For each state s, we consider the following LP with variables $\{\mu_{s'}\}_{s' \in S}$

$$\begin{aligned} \text{minimise} \quad & \sum_{s' \in S \setminus T} \mu_{s'} \cdot y_{s'} + \sum_{s' \in T} \mu_{s'} \\ \text{subject to} \quad & \sum_{s' \in S} \mu_{s'} = 1 \\ & \mathbf{P}^l(s, s') \leqslant \mu_{s'} \leqslant \mathbf{P}^u(s, s'). \end{aligned} \tag{3}$$

We write z_s for the optimal value corresponding to s. We have that $\vec{y}$ satisfies the constraints in (2) iff $y_s \leqslant z_s$ for each $s \in S$. To see this, suppose that $\vec{y}$ satisfies the constraints in (2), then for each state s, if $s \in T$, we clearly have $z_s = 1$ and thus $y_s \leqslant z_s$; if $s \notin T$, then by standard result of LP, there must exist some $\mu \in \mathcal{V}(\nabla_s)$ such that $\sum_{s' \in S \setminus T} \mu_{s'} y_{s'} + \sum_{s' \in T} \mu_{s'}$, achieves its minimum. It follows from the definition of z_s that $y_s = z_s$. On the other hand, suppose that $y_s \leqslant z_s$ for each state s, clearly for each $\mu \in \mathcal{V}(\nabla_s)$, $y_s \leqslant \sum_{s' \in S \setminus T} \mu_{s'} y_{s'} + \sum_{s' \in T} \mu_{s'}$, which implies that $\vec{y}$ satisfies the constraints in (2).

Now suppose that $\vec{y}$ does *not* satisfy the constraints in (2) for which we need to define the separating hyperplane. In this case, there must exist (at least) a state $\hat{s} \in S$ such that $y_{\hat{s}} > z_{\hat{s}}$. Let μ be a vertex that realises $\vec{z}$ at state $\hat{s}$. By Lemma 1, μ is rational of size polynomial in $\sharp(\mathcal{I})$. We define $\vec{\alpha}$ as: $\alpha(t) = \mu(\hat{s}) - 1$ if $t = \hat{s}$; and $\mu(t)$ otherwise. As $\vec{y}$ does *not* satisfy the LP (2), we have that

$$y(\hat{s}) > \sum_{s' \in S \setminus T} \mu(s') \cdot y(s') + \sum_{s' \in T} \mu_{s'}. \tag{4}$$

For any $\vec{x}$ satisfying the LP (2), we have that

$$x(\hat{s}) \leqslant \sum_{s' \in S \setminus T} \mu(s') \cdot x(s') + \sum_{s' \in T} \mu_{s'}. \tag{5}$$

It follows that

$$\begin{aligned} \sum_{s \in S} \alpha(s) y(s) &= \big(\mu(\hat{s}) - 1\big) y(\hat{s}) + \sum_{s \neq \hat{s}} \mu(s) y(s) \\ &= \sum_{s \in S} \mu(s) y(s) - y(\hat{s}) \\ &< 0 \quad \text{by (4)} \\ &\leqslant \sum_{s \in S} \mu(s) x(s) - x(\hat{s}) \quad \text{by (5)} \\ &= \big(\mu(\hat{s}) - 1\big) x(\hat{s}) + \sum_{s \neq \hat{s}} \mu(s) x(s) \\ &= \sum_{s \in S} \alpha(s) x(s). \end{aligned}$$

This completes the proof. □

Proposition 2 immediately entails the following.

Proposition 3. *Given an IDTMC $\mathcal{I}$, the reachability problem can be solved in polynomial time.*

We note that this algorithm is not only of theoretical interest; it also has practical impact since, based on the construction of separating oracles, one can use, for instance, a very efficient randomised algorithm [3], or the algorithm based on interior methods [25] (or its numerous improvements/variants).

It is well-known that computing the maximum (or minimum) reachability probabilities for MDPs is P-hard [21]. As solving reachability for IDTMCs essentially involves MDPs with infinitely (or at least exponentially) many actions, it is not hard to imagine that this is also P-hard. We show the following.

Proposition 4. *Given an IDTMC $\mathcal{I}$, the reachability problem is P-hard.*

Proof. Given any MDP $\mathcal{M} = (S, s_0, \delta, L)$, we construct an IDTMC $\mathcal{I}$ as follows: the state space of $\mathcal{I}$ consists of $\{s, (s, \mu) \mid s \in S \text{ and } \mu \in \delta(s)\}$. For any $\mu \in \delta(s)$ in $\mathcal{M}$, we have that $\mathbf{P}^l(s, (s, \mu)) = 0$ and $\mathbf{P}^u(s, (s, \mu)) = 1$, and for each s' we have $\mathbf{P}^l((s, \mu), s') = \mathbf{P}^u((s, \mu), s') = \mu(s, s')$. For any other pairs of states, the upper bound (hence the lower bound) of the transition probabilities are 0. Moreover, $\mathcal{M}$ and $\mathcal{I}$ have the same goal states $T \subseteq S$. We show that the minimum (resp. maximum) reachability probability to T in $\mathcal{M}$ equals the minimum (resp. maximum) reachability probability to T in $\mathcal{I}$.

Suppose that σ is the scheduler which achieves the minimum reachability probability in $\mathcal{M}$. Evidently we can assume that σ is simple, i.e., for each state s, σ assigns a unique μ. We then construct σ' for $\lceil \mathcal{I} \rceil$ such that $\sigma'(s) = (s, \sigma(s))$ and $\sigma'((s, \mu)) = \mu$. Clearly, $\mathcal{M}_\sigma$ and $\lceil \mathcal{I} \rceil_{\sigma'}$ yield the same reachability probability to T. Hence $\inf_\sigma \Pr(\mathcal{M}_\sigma, \Diamond T) \geqslant \inf_\sigma \Pr(\lceil \mathcal{I} \rceil_\sigma, \Diamond T)$.

On the other hand, suppose that σ is the scheduler which achieves the minimum reachability probability in $\lceil \mathcal{I} \rceil$. W.l.o.g. we can assume that σ is history-independent. Clearly, for a state of the form (s, μ), $\sigma((s, \mu)) = \mu$ as, according to the definition of $\mathcal{I}$, μ is the only option there. For state s, $\sigma(s)$ gives rise to a randomised scheduler σ' for $\mathcal{M}$, namely, for $\mu \in \delta(s)$, $\sigma'(s)(\mu) = \sigma(s)((s, \mu))$. Clearly, $\mathcal{M}_{\sigma'}$ and $\lceil \mathcal{I} \rceil_\sigma$ yield the same reachability probability T. However, it is well-known that the minimum reachability probability of $\mathcal{M}$ is achieved by simple schedulers [22], hence

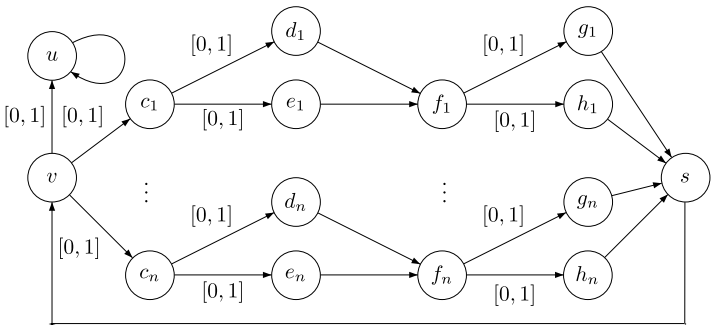

Fig. 1. IDTMC for the proof, adapted from [5].

$\inf_\sigma \Pr(\lceil\mathcal{I}\rceil_\sigma, \Diamond T) \geqslant \inf_\sigma \Pr(\mathcal{M}_\sigma, \Diamond T)$. It follows that $\inf_\sigma \Pr(\mathcal{M}_\sigma, \Diamond T) = \inf_\sigma \Pr(\lceil\mathcal{I}\rceil_\sigma, \Diamond T)$, which completes the proof. (The maximum case can be proved in exactly the same way.) □

Combining Proposition 3 and Proposition 4, we obtain the following result.

Theorem 1. *The reachability problem for IDTMCs is P-complete.*

4. PCTL

In this section, we consider PCTL model checking. As mentioned before, UMC and IMDP semantics differ in this case, and we treat them separately below.

4.1. UMC semantics

An instance of the Square-Root-Sum problem is a tuple $(x_1, \ldots, x_n, y)$ of integers, and the question is to decide whether $\sum_{i=1}^n \sqrt{x_i} \leqslant y$. By a minor adaption of the construction in [5] for MDPs, we show the following.

Theorem 2. *Model checking IDTMCs against PCTL under the UMC semantics is* Square-Root-Sum *hard.*

Proof. Let $(x_1, \ldots, x_n, y)$ be an instance of the Square-Root-Sum problem. Define an IDTMC $\mathcal{I}$ as follows (a schematic description is given in Fig. 1):

The state space consists of $v, u, s, c_i, d_i, e_i, f_i, g_i, h_i$ for $1 \leqslant i \leqslant n$. For the transitions $v \to u$, $v \to c_i$, $c_i \to d_i$, $c_i \to e_i$, $f_i \to g_i$, and $f_i \to h_i$, the associated intervals are $[0, 1]$, while for transitions $e_i \to f_i$, $d_i \to f_i$, $h_i \to s$, $g_i \to s$, $s \to v$, and $u \to u$, the associated intervals are $[1, 1]$. Moreover, we assume that the label of each state (see Fig. 1) is valid *only* in that state.

Let $q = y + \sum_{i=1}^n x_i$. We now construct a formula

$$\Xi = \mathcal{P}_{\geqslant 1-\frac{b}{q}}(\mathsf{X}\, u) \wedge \bigwedge_{1\leqslant i\leqslant n} \Phi_i \wedge \Psi_i \wedge \Theta_i$$

where for each $1 \leqslant i \leqslant n$, let

$$\Phi_i = \mathcal{P}_{=\frac{x_i}{q^2}}\big((v \vee c_i)\, \mathsf{U}\, e_i\big),$$

$$\Psi_i = \mathcal{P}_{>0}\big(\mathsf{X}\big(c(i) \vee \mathcal{P}_{=\frac{x_i}{q^2}}(e_i \vee f_i)\, \mathsf{U}\, h_i\big)\big)$$

and

$$\Theta_i = \mathcal{P}_{>0}\big(\mathsf{X}\big(\mathcal{P}_{>0}\big(\mathsf{X}\mathcal{P}_{>0}\big(\mathsf{X}\mathcal{P}_{=\frac{x_i}{q^2}}(f_i \vee h_i \vee s \vee v)\, \mathsf{U}\, c_i\big)\big)\big)\big).$$

For any DTMC $\mathcal{D} \in [\mathcal{I}]$, we claim that $v \models \Xi$ iff

$$\mathbf{P}(v, c_i) = \mathbf{P}(c_i, e_i) = \mathbf{P}(f_i, h_i) = \frac{\sqrt{x_i}}{q} \tag{6}$$

for each $1 \leqslant i \leqslant n$. To see this, note that $v \models \Phi_i$ iff $\mathbf{P}(v, c_i) \cdot \mathbf{P}(c_i, e_i) = \frac{x_i}{q^2}$, $v \models \Psi_i$ iff $\mathbf{P}(c_i, e_i) \cdot \mathbf{P}(f_i, h_i) = \frac{x_i}{q^2}$, and $v \models \Theta_i$ iff $\mathbf{P}(f_i, h_i) \cdot \mathbf{P}(v, c_i) = \frac{x_i}{q^2}$. Thus (6) follows immediately.

It is easy to see that in $\mathcal{D}$ it must be the case that $\mathbf{P}(v, u) = 1 - \sum_{i=1}^n \frac{\sqrt{x_i}}{q}$. The conjunct $\mathcal{P}_{\geqslant 1-\frac{y}{q}}(\mathsf{X}\, u)$ in Ξ stipulates that $1 - \sum_{i=1}^n \frac{\sqrt{x_i}}{q} \geqslant 1 - \frac{y}{q}$, i.e., $\sum_{i=1}^n \sqrt{x_i} \leqslant y$. This completes the proof. □

4.2. IMDP semantics

Following the same argument as [24, Theorem 3], one can reduce the model checking problem for PCTL under the IMDP semantics to the same problem for MDPs. It is routine to apply the labelling algorithm (in a bottom-up fashion) as in [4] to solve this problem, and the core is to compute the optimal (i.e., maximum or minimum) probability of path formulas under all schedulers. We note that a slightly more complicated case is to check the operator $\mathcal{P}_{=p}(\phi)$. Chen et al. [8] show that $\mathcal{P}_{=p}(\phi)$ holds iff $p \in [\inf_\sigma \Pr^\sigma(\phi), \sup_\sigma \Pr^\sigma(\phi)]$, hence it can be reduced to computing the optimal probability of path formula ϕ as well. The path formula of the form $\mathsf{X}\Phi$ can be tackled trivially and the path formula $\Phi\, \mathsf{U}\, \Psi$ can be reduced to the reachability problem by making states not satisfying Φ absorbing. We refer the readers to [4] or textbook [2, Chapter 10] for details. It follows from Theorem 1 that model checking IDTMCs against PCTL under the IMDP semantics is also P-complete. This improves an upper bound of [7].

Theorem 3. *Model checking IDTMCs against PCTL under the IMDP semantics is P-complete.*

Remark 4. In [7] a slightly more complicated logic called ω-PCTL is considered. One can use the same construction given in [7, Section 4], together with Theorem 1, to show that model checking problem for this logic is also P-complete. We choose not to present this result simply for

Table 1
Summary of the results.

	Reachability	(ω-)PCTL
UMC	P-complete	Square-Root-Sum-hard
IMDP	P-complete	P-complete

simplicity, as the techniques are the same as those demonstrated above.

5. Conclusion

We have studied the complexity of model checking IDTMCs under the UMC semantics and the IMDP semantics. Table 1 summarises the main results.

We believe that techniques introduced in this paper can be generalised (at least partially) to tackle more general models like constraint Markov chains [6] or robust MDPs [19], as far as reachability problems are concerned. This is the subject of future work.

Acknowledgements

We are grateful to Lu Feng for bringing the problem to our attention, and to Vojtěch Forejt, Daniel Klink, and Aistis Simaitis for inspiring discussions.

References

[1] Eric Allender, Peter Bürgisser, Johan Kjeldgaard-Pedersen, Peter Bro Miltersen, On the complexity of numerical analysis, SIAM J. Comput. 38 (5) (2009) 1987–2006.
[2] Christel Baier, Joost-Pieter Katoen, Principles of Model Checking, MIT Press, 2008.
[3] Dimitris Bertsimas, Santosh Vempala, Solving convex programs by random walks, J. ACM 51 (4) (2004) 540–556.
[4] Andrea Bianco, Luca de Alfaro, Model checking of probabilistic and nondeterministic systems, in: P.S. Thiagarajan (Ed.), FSTTCS, in: Lecture Notes in Computer Science, vol. 1026, Springer, 1995, pp. 499–513.
[5] Tomás Brázdil, Václav Brozek, Vojtech Forejt, Antonín Kucera, Stochastic games with branching-time winning objectives, in: LICS, IEEE Computer Society, 2006, pp. 349–358.
[6] Benoît Caillaud, Benoît Delahaye, Kim G. Larsen, Axel Legay, Mikkel L. Pedersen, Andrzej Wasowski, Constraint Markov chains, Theoret. Comput. Sci. 412 (34) (2011) 4373–4404.
[7] Krishnendu Chatterjee, Koushik Sen, Thomas A. Henzinger, Model-checking omega-regular properties of interval Markov chains, in: Roberto M. Amadio (Ed.), FoSSaCS, in: Lecture Notes in Computer Science, vol. 4962, Springer, 2008, pp. 302–317.
[8] Taolue Chen, Vojtech Forejt, Marta Kwiatkowska, Aistis Simaitis, Ashutosh Trivedi, Michael Ummels, Playing stochastic games precisely, in: CONCUR, 2012.
[9] Benoît Delahaye, Kim G. Larsen, Axel Legay, Mikkel L. Pedersen, Andrzej Wasowski, Consistency and refinement for interval Markov chains, J. Log. Algebr. Program. 81 (3) (2012) 209–226.
[10] Kousha Etessami, Mihalis Yannakakis, Recursive Markov chains, stochastic grammars, and monotone systems of nonlinear equations, J. ACM 56 (1) (2009).
[11] Harald Fecher, Martin Leucker, Verena Wolf, Don't Know in probabilistic systems, in: Antti Valmari (Ed.), SPIN, in: Lecture Notes in Computer Science, vol. 3925, Springer, 2006, pp. 71–88.
[12] Robert Givan, Sonia M. Leach, Thomas Dean, Bounded-parameter Markov decision processes, Artif. Intell. 122 (1–2) (2000) 71–109.
[13] Martin Grötschel, László Lovász, Alexander Schrijver, The ellipsoid method and its consequences in combinatorial optimization, Combinatorica 1 (2) (1981) 169–197.
[14] Hans Hansson, Bengt Jonsson, A logic for reasoning about time and reliability, Form. Asp. Comput. 6 (5) (1994) 512–535.
[15] Bengt Jonsson, Kim Guldstrand Larsen, Specification and refinement of probabilistic processes, in: LICS, 1991, pp. 266–277.
[16] Richard M. Karp, Christos H. Papadimitriou, On linear characterizations of combinatorial optimization problems, SIAM J. Comput. 11 (4) (1982) 620–632.
[17] Joost-Pieter Katoen, Daniel Klink, Martin Leucker, Verena Wolf, Three-valued abstraction for probabilistic systems, J. Log. Algebr. Program. 81 (4) (2012) 356–389.
[18] Igor Kozine, Lev V. Utkin, Interval-valued finite Markov chains, Reliab. Comput. 8 (2) (2002) 97–113.
[19] Arnab Nilim, Laurent El Ghaoui, Robust control of Markov decision processes with uncertain transition matrices, Oper. Res. 53 (5) (2005) 780–798.
[20] Manfred W. Padberg, Rama Mohana Rao, The Russian method and integer programming, Working paper series 203, Salomon Brothers Center for the Study of Financial Institutions, 1980.
[21] Christos Papadimitriou, John Tsitsiklis, The complexity of Markov decision processes, Math. Oper. Res. 12 (3) (1987) 441–450.
[22] Martin L. Puterman, Markov Decision Processes: Discrete Stochastic Dynamic Programming, Wiley, New York, 1994.
[23] Alexander Schrijver, Theory of Linear and Integer Programming, Wiley–Interscience Series in Discrete Mathematics and Optimization, Wiley, 1999.
[24] Koushik Sen, Mahesh Viswanathan, Gul Agha, Model-checking Markov chains in the presence of uncertainties, in: Holger Hermanns, Jens Palsberg (Eds.), TACAS, in: Lecture Notes in Computer Science, vol. 3920, Springer, 2006, pp. 394–410.
[25] Pravin M. Vaidya, A new algorithm for minimizing convex functions over convex sets, Math. Program. 73 (1996) 291–341.
[26] Huan Xu, Shie Mannor, Distributionally robust Markov decision processes, in: John D. Lafferty, Christopher K.I. Williams, John Shawe-Taylor, Richard S. Zemel, Aron Culotta (Eds.), NIPS, Curran Associates, Inc., 2010, pp. 2505–2513.