Robust Model Checking with Imprecise Markov Reward Models

Alberto Termine ALBERTO.TERMINE@UNIMI.IT
University of Milan, Milan, Italy

Alessandro Antonucci ALESSANDRO@IDSIA.CH
Alessandro Facchini* ALESSANDRO.FACCHINI@IDSIA.CH
Istituto Dalle Molle di Studi sull'Intelligenza Artificiale (IDSIA) - Lugano, Switzerland

Giuseppe Primiero GIUSEPPE.PRIMIERO@UNIMI.IT
University of Milan, Milan, Italy

Abstract

In recent years probabilistic model checking has become an important area of research because of the diffusion of computational systems of stochastic nature. Despite its great success, standard probabilistic model checking suffers the limitation of requiring a sharp specification of the probabilities governing the model behaviour. The theory of imprecise probabilities offers a natural approach to overcome such limitation by a sensitivity analysis with respect to the values of these parameters. However, only extensions based on discrete-time imprecise Markov chains have been considered so far for such a robust approach to model checking. We present a further extension based on imprecise Markov reward models. In particular, we derive efficient algorithms to compute lower and upper bounds of the expected cumulative reward and probabilistic bounded rewards based on existing results for imprecise Markov chains. These ideas are tested on a real case study involving the spend-down costs of geriatric medicine departments.

Keywords: Probabilistic Computational Tree Logic, Model-Checking, Imprecise Markov Chains, Imprecise Markov Reward Models.

1. Introduction

Model checking is a fully-automatic logic-based technique to decide whether a model satisfies some given requirements. As such, it has been mostly applied as a tool to decide correctness and compliance in systems design. Relevant examples are software verification (e.g., Bérard et al., 2001), system communication and also computational biology (e.g., Brim et al., 2013; Benes et al., 2019). During the last decades, the spread of stochastic computational systems led to the development of appropriate formal models and logical languages. A notable success has been achieved by *probabilistic computational tree logic* (PCTL, Hansson and Jonsson, 1994), and its extensions (e.g., PCTL*, PRCTL, PCTLK), that are conceived to specify properties of Markov processes. Standard PCTL and PCTL*, in particular, are the reference languages for *discrete-time Markov chains* (Baier and Katoen, 2008) and *Markov decision processes* (Henriques et al., 2012), whereas PRCTL and CSL are used, respectively, for *Markov reward models* (Andova et al., 2003) and *continuous-time Markov chains* (Baier et al., 2000). Other extensions have been also proposed to cope with *hidden Markov models* (Zhang et al., 2005) and *probabilistic interpreted systems* (Chen et al., 2016).

Despite their wide applicability, PCTL and its extensions suffer from a significant limitation: the probabilities governing the model behaviour require a sharp quantification. This is a serious issue in many real-case scenarios. In computational biology, for instance, non-homogeneous processes are quite common (Benes et al., 2019). Similarly, ignorance about probabilities represents an important challenge for epistemic multi-agent systems. To overcome these limitations, different approaches have been proposed. *Parametric Markov chains* (Daws, 2004; Ilie and Worrell, 2020) model such ignorance by treating transition probabilities as parameters. The applicability of these models is however limited because of the computational complexity of the relative model-checking procedures, which is exponential with respect to the number of states of the model, even for the most advanced procedures based on fraction-free Gaussian elimination (Baier et al., 2020).

An alternative and less demanding approach is provided by the theory of *imprecise probabilities* (Walley, 1991) and, in particular, by the most recent works on imprecise Markovian processes, which provide a robust approach to the modelling of non-homogeneous Markov processes, as well as to the modelling of partial ignorance about transition probabilities (e.g., de Cooman et al., 2016). Troffaes and Škulj (2013) have outlined the first attempt to extend the framework of probabilistic model checking to imprecise probability models. On the side of properties specification, they replace the standard PCTL probability operator with a new one weighted by an interval whose bounds correspond to the respective lower and upper bounds of probabilistic inferences in an imprecise Markov chain. On the semantic

* Supported by the Hasler foundation grant n. 20061.

side, they prove that relevant model-checking tasks concerning probabilistic formulae can be reduced to well-known marginalization tasks in imprecise Markov chains. Finally, they provide specific algorithms that exploit marginalization to check whether a model satisfy a given imprecise probabilistic formula.

The contribution of Troffaes and Škulj (2013) is focused on discrete-time Markov chains. A very natural way to extend such seminal work is to consider an imprecise version of Markov reward models. The present work explores this extension. Imprecise Markov reward models are simply intended as imprecise Markov chains provided with a labelling function that assigns a numerical reward to each state in the model. Two new robust inference tasks are considered, these being the computation of the bounds of *expected cumulative reward* and of the *reward-bounded* probabilities. The IPCTL language proposed by Troffaes and Škulj (2013) is therefore extended here including new operators to represent those inferences. We call the language we obtain IPRCTL. We provide the new language with a proper semantics and define satisfiability conditions for the new operators. Hence, we present specific algorithms to compute the relevant inferences specified by those operators. These algorithms are derived by the schema introduced by T'Joens et al. (2019) to compute robust inferences in imprecise Markov chains. Finally, we outline some considerations about the computational complexity of those new algorithms. Notably shifting from precise to imprecise probabilities does not affect the overall computational complexity of the relevant model-checking procedures.

The paper is organized as follows. In Section 2 we review some basic material. The definitions and algorithms for, respectively, precise and imprecise Markov chains are discussed in Sections 3.1 and 3.2. Section 4 contains the syntax and semantics of both PCTL and PRCTL, while the imprecise-probabilistic extensions are discussed in Section 5. Before the conclusions in Section 7, we validate these algorithms in Section 6 with a case study about the spend-down costs of geriatric medicine departments based on a sensitivity analysis of the results in McClean et al. (1998).

2. Background

We first review the necessary background material and notation about the theory of imprecise probability. Let S be a variable taking its values from a finite non-empty set of states $\mathscr{S}$ whose generic elements are denoted by $s \in \mathscr{S}$. A *probability mass function* (PMF, for short) over S, denoted by $P(S)$, is defined as a non-negative normalized real map over $\mathscr{S}$. Given a real-valued function f of S, i.e., $f : \mathscr{S} \to \mathbb{R}$, its *expectation* based on P is defined as $E[f] := \sum_{s\in\mathscr{S}} f(s)P(s)$. Notation $P(S'|s) := \{P(s'|s)\}_{s\in\mathscr{S}}$ and $P(S'|S) := \{P(S'|s)\}_{s\in\mathscr{S}}$ is used instead for conditional probabilities.

A *credal set* (CS) over S, denoted as $K(S)$, is a collection of PMFs over S. We consider here only finitely generated CSs, i.e., CSs whose convex hull has only a finite number of extreme points. Given a real-valued function f of $\mathscr{S}$, the upper expectation of f with respect to $K(S)$ is defined as $\overline{E}[f] := \sup_{P(S)\in K(S)} E_P[f]$. The lower expectation $\underline{E}$ is similarly defined. Suprema (infima) of upper (lower) expectations can be equivalently reduced to maxima (mimima) over the extreme points of the CS convexification. Consequently, we can identify a CS with the extreme points of its convex hull. Conditional expectations can be similarly considered.

3. Markovian Models

Les us first discuss how *discrete-time Markov chains* can model the behaviour of stochastic, time-homogeneous and memory-less, agents with a finite number of possible states.

3.1. (Precise) Markov Chains

Consider an agent defined over a finite non-empty set of possible states $\mathscr{S}$. The agent evolves from a state $s \in \mathscr{S}$ to another state $s' \in \mathscr{S}$. Any possible temporal evolution of the agent across time is described by a countable sequence of states that is called a *path* and denoted by π. Similarly, we use Π to denote the set of all possible paths, whereas we use $\pi(t)$ to denote the generic state of the path π at time $t \in \mathbb{N}$. The agent is stochastic, meaning that there is a certain degree of *uncertainty* about which path describes its actual evolution. This uncertainty can be measured. To do so, we endow Π with a σ-algebra $\sigma(\Pi)$ and augment it to a probability space $(\Pi, \sigma(\Pi), P)$. [1] Over this probability space, we define a family $\{S_t\}_{t\in\mathbb{N}}$ of categorical stochastic variables S_t such that, for all $t \in \mathbb{N}$, $S_t : \pi \mapsto \pi(t)$. For each $t \in \mathbb{N}$, $P(S_t)$ denotes a PMF that assigns to each $s \in \mathscr{S}$ the probability of s to be the state of the agent at time $t \in \mathbb{N}$. Similarly, for each $t \in \mathbb{N}$, $P(S_{t+1} = s'|S_t = s)$ are the conditional probabilities modelling, for each pair of states $s, s' \in \mathscr{S}^2$, the probability of the agent to reach s' at time $t+1$ given that it is in state s at time t. Because the agent is memory-less, it satisfies the Markov property, i.e.,

$$P(S_{t+1}|S_t, \dots, S_0) = P(S_{t+1}|S_t). \tag{1}$$

Furthermore, we assume the behaviour of the agent to be *time-homogeneous*, i.e., $P(S_{t+1}|S_t)$ is the same for all $t \in \mathbb{N}$. Given the Markov property and the time-homogeneity, a compact specification of the agent behaviour is possible in terms of an initial PMF $P(S_0)$ and a transition matrix $T : \mathscr{S}^2 \mapsto [0,1]$, whose elements are the values in $P(S_{t+1}|S_t)$

1. In particular, $\sigma(\Pi)$ is the σ-algebra generated by the cylinder sets, also called *cylinder σ-algebra* (Revuz, 2008). This allows all the functions we introduce to be measurable.

and where the choice of t is arbitrary because of time-homogeneity. Such a model is called here a *precise Markov chain* (PMC) and denoted by M.

Typical inferential tasks in PMCs can be computed by means of a *transition operator* $\hat{T}$ mapping a real-valued function f of $\mathscr{S}$ to its (left) scalar product by T, i.e.,

$$(\hat{T}f)(s) := \sum_{s' \in \mathscr{S}} T(s',s) \cdot f(s'), \tag{2}$$

for each $s \in \mathscr{S}$. The *dual* $\hat{T}^\dagger$ of this linear operator is obtained by a right scalar product as follows:

$$(\hat{T}^\dagger f)(s) := \sum_{s' \in \mathscr{S}} T(s,s') \cdot f(s'), \tag{3}$$

for each $s \in \mathscr{S}$. By total probability it is easy to check that $\hat{T}P(S_t) = P(S_{t+1})$ and hence $\hat{T}^t P(S_0) = P(S_t)$. By the notion of conditional expectation, $\hat{T}^\dagger f(S_t) = E_P[f(S_{t+1})|S_t]$, and hence $((\hat{T}^\dagger)^t f(S_0))(s) = E_P[f(S_t)|S_0 = s]$.

We similarly compute the *hitting probability* vector $h^{\leq t}_{\mathscr{A}}$ for a finite time-horizon $t \in \mathbb{N}$. For each $s \in \mathscr{S}$, and $\mathscr{A} \subseteq \mathscr{S}$, $h^{\leq t}_{\mathscr{A}}(s)$ is defined as the probability of having at least a state $s_\tau \in \mathscr{A}$ for some $\tau \leq t$, provided that $S_0 = s$. Clearly, $h^{\leq 0}_{\mathscr{A}} = \mathbb{I}_{\mathscr{A}}$, i.e., the indicator function of $\mathscr{A}$ gives the hitting vector for $t = 0$ being one for states consistent with $\mathscr{A}$ and zero otherwise. We say that a state $s \in \mathscr{S}$ is *absorbing* if $T(s,s') = 0$ for each $s' \neq s$, i.e., once the model is in an absorbing state, the transition probabilities to other states are all zero. Let $T_{\mathscr{A}}$ denote the transition matrix obtained from T by making absorbing all states $s \in \mathscr{A}$. We can obtain the hitting probability as:

$$h^{\leq t}_{\mathscr{A}}(s) = \sum_{s' \in \mathscr{A}} T^t_{\mathscr{A}}(s,s') = [(\hat{T}^\dagger)^t \mathbb{I}_{\mathscr{A}}](s), \tag{4}$$

for $s \in \mathscr{A}^c := \mathscr{S} \setminus \mathscr{A}$, while, trivially, $h^{\leq t}_{\mathscr{A}}(s) = 1$ for $s \in \mathscr{A}$. The dual of the above computation corresponds to the recursion:

$$h^{\leq t}_{\mathscr{A}} = \mathbb{I}_{\mathscr{A}} + \mathbb{I}_{\mathscr{A}^c} \hat{T}^\dagger h^{\leq t-1}_{\mathscr{A}}, \tag{5}$$

for each $t > 0$, where sums and products by indicator functions are intended as element-wise operations on arrays.

The unbounded hitting probability vector $h_{\mathscr{A}}$ whose elements are the values of the probability of the agent to reach at least a state $s \in \mathscr{A}$ *evetually in the future* computed for each $s \in \mathscr{S}$, can be achieved as $\lim_{t\to\infty} h^{\leq t}_{\mathscr{A}}$. Proves of the existence of this limit are available in literature, see Revuz (2008). Here, simply note that $\lim_{t\to\infty} h^{\leq t}_{\mathscr{A}}$ corresponds to the fixed point of Equation (5). We remind the reader to classical references, e.g., Revuz (2008), for the formal proofs of these results.

Finally, following Katoen et al. (2005), we define a *Markov reward model* (MRM) as a PMC paired with a reward function $rew : \mathscr{S} \mapsto \mathbb{N}$. We call the natural number $rew(s)$, the reward of state $s \in \mathscr{S}$, while rew also denotes the array of all the rewards for the different values of $\mathscr{S}$.

The *cumulative* reward of the t-th state of a path π in the time range $[0,t]$ is intended as:

$$Rew(\pi,t) := \sum_{\tau=0}^{t} rew(\pi(\tau)). \tag{6}$$

3.2. Imprecise Markov Chains

The generalization of PMCs to imprecise probability can be achieved in different ways, possibly leading to different inferential results on specific tasks (Krak et al., 2019). The approach we adopt here, also called *model-theoretic*, is based on a direct sensitivity analysis interpretation approach): an IMC is intended as a family of PMCs all compatible with the assessments about the system uncertain behaviour. Under this interpretation, a generalization of a PMC to an IMC can be achieved by replacing $P(S_0)$ with a CS $K(S_0)$ and the transition matrix T with a *credal transition matrix* $\mathscr{T}$ made of conditional CSs, i.e., $\mathscr{T} := \{K(S'|s)\}_{s\in\mathscr{S}}$ and characterizing the transitions from S to S'. Time-homogeneity consists instead in assuming the specification of the (collections of) CSs $K(S_{t+1}|S_t)$ independent of t. The dual (linear) transition operator in Equation (3) admits the following (non-linear) extension to IMCs:

$$(\overline{\mathscr{T}} f)(s) := \sup_{T(s,S') \in K(S'|s)} \sum_{s' \in \mathscr{S}} T(s,s') \cdot f(s'), \tag{7}$$

to be considered for each $s \in \mathscr{S}$, with $K(S'|s) \in \mathscr{T}$ (Troffaes and Škulj, 2013). An analogous *lower* operator $\underline{\mathscr{T}}$ can be defined by replacing the supremum in Equation (7) with an infimum. Note that the optimization in Equation (7) is a linear programming task whose feasible region is the convex hull of $K(S'|s)$, that in our assumptions can be described by a finite number of linear constraints. It is easy to check that:

$$(\overline{\mathscr{T}}^t f)(s) = \overline{E}(f(S_t)|S_0 = s), \tag{8}$$

for each $s \in \mathscr{S}$. A similar relation holds for the lower operator and the lower expectation. As recently shown by T'Joens et al. (2019), the recursion in Equation (5) to compute the hitting probabilities in a PMC can be easily generalized to IMCs as follows:

$$\overline{h}^{\leq t}_{\mathscr{A}} = \mathbb{I}_{\mathscr{A}} + \mathbb{I}_{\mathscr{A}^c} \overline{\mathscr{T}}\, \overline{h}^{\leq t-1}_{\mathscr{A}}, \tag{9}$$

with starting point $\overline{h}^{\leq 0}_{\mathscr{A}} = \mathbb{I}_{\mathscr{A}}$. The upper hitting probability vector $\overline{h}^{\leq t}_{\mathscr{A}}$ is intended as the upper bound, computed with respect to the joint CS induced by the IMC, of the hitting probability vector defined in Section 3.2. A similar recursion, involving the lower operator and giving the lower hitting probability also holds. As shown by Troffaes and Škulj (2013), those recursions can equivalently provide a generalization of Equation (4):

$$\overline{h}^{\leq t}_{\mathscr{A}}(s) = (\overline{\mathscr{T}}^t_{\mathscr{A}} \mathbb{I}_{\mathscr{A}})(s), \tag{10}$$

for each $s \in \mathscr{A}^c$ being instead one for $s \in \mathscr{A}$, where $\overline{\mathscr{T}}_{\mathscr{A}}$ is such that $\overline{\mathscr{T}}_{\mathscr{A}} f(s) = f(s)$ for $s \in \mathscr{A}$ and $\overline{T}_{\mathscr{A}} f(s) = \overline{T} f(s)$ for $s \in \mathscr{A}^c$. When paired with a reward function *rew*, the IMC is called *imprecise* MRM (IMRM). The problem of computing lower and upper expectations of the cumulative reward of a path as in Equation (6) is one of the algorithmic challenges we want to address in the rest of the paper.

4. Probabilistic Computational Tree Logic

We open the discussion by reviewing the syntax (Section 4.1) and semantics (Section 4.2) of PCTL, the reference language for MC based on PMCs. A demonstrative model-checking task involving PCTL is in Section 4.3. Afterwards, we show how the PCTL syntax has been extended to take into account reward functions in PRCTL (Section 4.4), whose semantics is reduced to MRM queries (Section 4.5).

4.1. PCTL Syntax

The PCTL syntax is recursively defined as follows:

$$\phi := \top \mid p \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid P_{\nabla b}\psi\,, \tag{11}$$

$$\psi := \bigcirc\,\phi \mid \phi_1 \bigcup \phi_2 \mid \phi_1 \overset{\leq t}{\bigcup} \phi_2\,. \tag{12}$$

The language includes the standard notation $\top$ for *true*, atoms (such as p) and standard Boolean formulae, whose meaning is the same as in standard propositional logic. It also includes path formulae denoted by ψ and representing properties of paths with the following informal reading (Baier and Katoen, 2008):

- $\bigcirc\phi$ means that *in the next state of the path ϕ hold*;
- $\phi_1 \bigcup^{\leq t} \phi_2$ means that ϕ_2 *holds at a certain time $\tau \leq t$ and ϕ_1 holds in all the previous states of the path*;
- $\phi_1 \bigcup \phi_2$ means that ϕ_2 *holds eventually along the path and ϕ_1 holds in all the previous states of the path.*

Finally, for the probability formulae, where $b \in [0,1]$ and $\nabla \in \{<,\leq,=,\geq,>\}$, $P_{\nabla b}\psi$ means that *there is a probability ∇b to reach a path satisfying ψ.*

4.2. PCTL Semantics

To present PCTL semantics we first augment standard PMCs (Section 3.1) with a finite non-empty set of atomic propositions AP and a *labelling* function $l : \mathscr{S} \mapsto 2^{AP}$ that assigns to each state $s \in \mathscr{S}$ a set of propositions $l(s) \subseteq AP$. The resulting model is called *labelled PMC*. In the following, by PMC we always denote a labelled PMC.

Boolean Formulae. Given PMC M and state $s \in \mathscr{S}$, the semantics for Boolean formulae is as follows:

$$M,s \models \top \;\forall s \in \mathscr{S}\,; \tag{13}$$

$$M,s \models p \text{ iff } p \in l(s)\,; \tag{14}$$

$$M,s \models \phi_1 \wedge \phi_2 \text{ iff } M,s \models \phi_1 \text{ and } M,s \models \phi_2\,; \tag{15}$$

$$M,s \models \neg\phi \text{ iff } M,s \not\models \phi\,. \tag{16}$$

Path Formulae. Given PMC M and path π, the following conditions hold:

$$M,\pi \models \bigcirc\phi \text{ iff } M,\pi(1) \models \phi\,; \tag{17}$$

$$M,\pi \models \phi_1 \overset{\leq t}{\bigcup} \phi_2 \text{ iff } \exists\tau \leq t : \begin{array}{l} M,\pi(\tau) \models \phi_2\,, \\ M,\pi(\tau') \models \phi_1\,, \\ \forall\tau' : 0 \leq \tau' < \tau\,; \end{array} \tag{18}$$

$$M,\pi \models \phi_1 \bigcup \phi_2 \text{ iff } \exists\tau \geq 0 : \begin{array}{l} M,\pi(\tau) \models \phi_2\,, \\ M,\pi(\tau') \models \phi_1\,, \\ \forall\tau' : 0 \leq \tau' < \tau\,. \end{array} \tag{19}$$

The *check* of the models with respect to Boolean formulae can be achieved by SAT solvers (Davis and Putnam, 1960), while the *parsing-tree* technique is used instead for path formulae (Baier and Katoen, 2008).

Probability Formulae. Given PMC M and state $s \in \mathscr{S}$ the following condition holds:

$$M,s \models P_{\nabla b}\psi \text{ iff } P(s \models \psi)\nabla b\,, \tag{20}$$

provided that $P(s \models \psi)$ is the probability of the PMC to reach a path $\pi \in \Pi$ such that $\pi \models \psi$ given that $S_0 = s$ (Baier and Katoen, 2008). For each one of the possible values of ψ, the computation of $P(s \models \psi)$ can be reduced to a PMC inference as in the following. Let us begin by considering the case $\psi := \bigcirc\phi$. The computation of $P(s \models \psi)$ can be achieved simply defining $\Phi := \{s' \in \mathscr{S} : M,s' \vDash \phi\}$ and computing the trivial inference:

$$P(S_1 \in \Phi | S_0 = s) := (\hat{T}^\dagger \mathbb{I}_\Phi)(s) = \sum_{s' \in \Phi} T(s,s')\,. \tag{21}$$

We now consider the case $\psi := \phi_1 \bigcup^{\leq t} \phi_2$. To compute $P(s \models \phi_1 \bigcup^{\leq t} \phi_2)$, we first define Φ_1 and Φ_2 as the subsets of $\mathscr{S}$ satisfying, respectively, ϕ_1 and ϕ_2. The probability in Equation (20), hence, can be computed as the hitting probability of event $\Phi_2|\Phi_1$, that is the set of all $s \in \mathscr{S}$ such that $s \in \Phi_2$ and all the states $s' \in \mathscr{S}$ visited before reaching s are in Φ_1. We denote such hitting probability as $h^{\leq t}_{\Phi_2|\Phi_1}$. A recursion analogous to that in Equation (5) is obtained by multiplying the complement of the hitting event Φ_2 by the indicator of Φ_1, thus obtaining the indicator of the set difference , i.e.,

$$h^{\leq\tau}_{\Phi_2|\Phi_1} := \mathbb{I}_{\Phi_2} + \mathbb{I}_{\Phi_1 \setminus \Phi_2} \hat{T}^\dagger h^{\leq\tau-1}_{\Phi_2|\Phi_1}\,, \tag{22}$$

for each $\tau = 1,\ldots,t$, with $h^{\leq 0}_{\Phi_2|\Phi_1} := \mathbb{I}_{\Phi_2}$. Case $\psi := \phi_1 \bigcup \phi_2$ is analogous to $\phi_1 \bigcup^{\leq t} \phi_2$ when $t \to \infty$. The value of $P(s \models$

$\phi_1 \bigcup \phi_2$), indeed, corresponds to $\lim_{t\to\infty} h^{\leq t}_{\Phi_2|\Phi_1}$. Remind that, this limit exists and it corresponds to the fixed point of the schema in Equation (22) (Revuz, 2008). For this reason, an easy way to approximate $\lim_{t\to\infty} h^{\leq t}_{\Phi_2|\Phi_1}$ consists in iterating the schema in Equation (22) until convergence. Notice that, limited to PMCs, there exist other procedures that allow to directly compute the value of $P(s \models \phi_1 \bigcup \phi_2)$ solving a system of linear equations through linear programming (Baier and Katoen, 2008, pp. 761-762).

4.3. A PCTL Example

Example 1 (Baier and Katoen (2008)) *Consider a simple communication model operating with a single channel. The channel is error-prone, meaning that messages can be lost. In this particular example, the (four) states of such model M are in one-to-one correspondence with the atomic propositions, i.e., $\mathscr{S} = AP :=$ {start, try, lost, delivered}. Transition probabilities are shown in Figure 1 as label arrows, impossible transitions correspond to missing edges.*

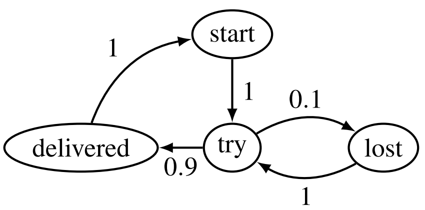

Figure 1: Transitions in a four-state PMC.

M is compliant if and only if "the probability of a message to be lost within seven time steps is smaller than or equal to 0.25"*. As the starting state for M is by construction $s =$ start, this corresponds to the following PCTL formula:*

$$M, \{\text{start}\} \models P_{\leq 0.25} \top \bigcup^{\leq 7} (\{\text{lost}\}), \tag{23}$$

whose semantics is such as in Equation (20). Yet, the task reduces to compute $h^{\leq 7}_{\{\text{lost}\}}(\mathscr{S}|\{\text{start}\})$. Its computation can be achieved either by the recursion in Equation (5) or by the matrix product in Equation (4), both leading to the numerical values in Table 1. Since the corresponding values is $h^{\leq 7}_{\{\text{lost}\}}(\text{start}) = 0.190 \leq 0.25$, the system satisfies the requirement and can be considered compliant.

4.4. PRCTL Syntax

PRCTL (Katoen et al., 2005) is a PCTL extension achieved by augmenting Equations (11) and (12) as follows:

$$\phi := \cdots \mid E_{\nabla r}\phi\,, \tag{24}$$

$$\psi := \cdots \mid \phi_1 \bigcup_{\leq r} \phi_2\,, \tag{25}$$

	s			
t	start	deliv.	try	lost
0	0.000	0.000	0.000	1.000
1	0.000	0.000	0.100	1.000
2	0.100	0.000	0.100	1.000
3	0.100	0.100	0.100	1.000
4	0.100	0.100	0.190	1.000
5	0.190	0.100	0.190	1.000
6	0.190	0.190	0.190	1.000
7	**0.190**	0.190	0.271	1.000

Table 1: Hitting probabilities $h^{\leq t}_{\{lost\}}(s)$ for $t \leq 7$.

where $\nabla := \{<, \leq, =, \geq, >\}$ and $r \in \mathbb{N}$. [2] The informal reading of the state formula $E_{\nabla r}\phi$ is that *the expected cumulative reward earned until reaching a state that satisfies ϕ is ∇r*. The *bounded reward* path formula $\phi_1 \bigcup_{\leq r} \phi_2$ means instead that *there exists $t \geq 0$ such that $\pi(t)$ satisfies ϕ_2 while all the previous states along the path satisfy ϕ_1 and the cumulative reward of path π from 0 to t is less than or equal to r.*

4.5. PRCTL Semantics

To define the PRCTL semantics we add three new satisfiability conditions to those in Section 4.2.

Bounded Reward. Given MRM (M, rew) and path π, the following condition holds:

$$(M, rew), \pi \models \phi_1 \bigcup_{\leq r} \phi_2 \text{ iff } \exists \tau : \begin{array}{l} (M, rew), \pi(\tau) \models \phi_2\,, \\ (M, rew), \pi(\tau') \models \phi_1, \forall \tau' < \tau\,, \\ Rew(\pi, \tau) \leq r\,. \end{array} \tag{26}$$

The procedure for this formula is analogous to that of the PCTL path formulae.

Expected Cumulative Reward. Given MRM (M, rew) and state $s \in \mathscr{S}$, the following condition holds:

$$(M, rew), s \models E_{\nabla r}\phi \text{ iff } E[Rew^{\leq t}_{\Phi}](s) \nabla r\,. \tag{27}$$

Notice that, when $t \to \infty$, i.e., when no state satisfying ϕ is reached until a finite time $t \in \mathbb{N}$, by definition $E[Rew^{\leq t}_{\Phi}](s) = +\infty$. Instead, for each $t \in \mathbb{N}$ and $s \in \mathscr{S}$ the expected cumulative reward is defined as follows:

$$E[Rew^{\leq t}_{\Phi}](s) := \sum_{\substack{\pi \in \text{Paths}(s): \\ \exists \tau \leq t : \pi(\tau) \models \phi}} Rew(\pi, \tau) \cdot P(\pi)\,, \tag{28}$$

That is, an expected value for paths starting in $S_0 = s$ and reaching Φ at some time step $\tau \leq t$. A recursion analogous

2. Following Baier and Katoen (2008), we assume rewards to be natural instead of real numbers. This assumption is very common in the field of model checking because it aids to avoid possible issues concerning the overall computational complexity of the main tasks.

to that in Equation (5) for the hitting event Φ can be therefore derived for the quantity in Equation (28). The only difference is that, if event Φ is achieved for $S_0 = s$, the value one in the indicator function should be multiplied by the corresponding reward for s. If this is not the case, the expected rewards of the previous time step should be propagated by the transition operator, and increased by the original rewards, i.e.,

$$E[Rew_{\Phi}^{\leq\tau}] = \mathbb{I}_{\Phi} \cdot rew + \mathbb{I}_{\Phi^c}(rew + \hat{T}^{\dagger}E[Rew_{\Phi}^{\leq\tau-1}]), \quad (29)$$

and hence:

$$E[Rew_{\Phi}^{\leq\tau}] = rew + \mathbb{I}_{\Phi^c}\hat{T}^{\dagger}E[Rew_{\Phi}^{\leq\tau-1}], \quad (30)$$

where both equations are array relations to be considered for each $\tau = 1,\dots,t$, sums and products should be intended as element-wise operations on arrays, and *rew* is the array of (state) rewards defined in Section 3.1. Note that, when $rew = \mathbb{I}_{\Phi}$, Equation (30) becomes Equation (5), i.e., if rewards are one on the hitting event and zero otherwise, the expected cumulative reward equals the hitting probability.

Bounded-Reward Probability. Given MRM (M, rew) and state $s \in \mathscr{S}$, the following condition holds:

$$(M, rew), s \models P_{\nabla b}\phi_1 \bigcup_{\leq r} \phi_2 \text{ iff } P(s \models \phi_1 \bigcup_{\leq r} \phi_2)\nabla b, \quad (31)$$

where, as in Equation (20), $P(s \models \phi_1 \bigcup_{\leq r} \phi_2)$ is the probability of the PMC to reach a path $\pi \models \phi_1 \bigcup_{\leq r} \phi_2$ given that $S_0 = s$. This can be intended as a hitting probability for event Φ_2 with respect to state $s \in \mathscr{S}$ provided that all the states visited before reaching Φ_2 satisfy ϕ_1, and the expected cumulative reward earned before reaching Φ_2 is less than or equal to r. This remark allows to derive a recursion analogous to that in Equation (22), that also takes into account the reward earned after any iteration. For each $\tau \geq 0$, $\rho \leq r$ and $s \in \mathscr{S}$, we denote by $x_{\Phi_2|\Phi_1}^{\leq\tau,\rho}(s)$ the hitting probability of reaching Φ_2 for some $\tau \geq 0$ provided that all the states visited before reaching Φ_2 are in Φ_1 and the expected cumulative reward earned until reaching Φ_2 is $\leq r$. Trivially, for $\tau = 0$:

$$x_{\Phi_2|\Phi_1}^{\leq 0,\rho}(s) := \begin{cases} 1 \text{ if } s \in \Phi_2 \text{ and } rew(s) \leq \rho, \\ 0 \text{ otherwise.} \end{cases} \quad (32)$$

In array notation, this rewrites as follows:

$$x_{\Phi_2|\Phi_1}^{\leq 0,\rho} := \mathbb{I}_{\mathscr{S}_{\rho}^{rew}} \cdot \mathbb{I}_{\Phi_2}, \quad (33)$$

where $\mathscr{S}_{\rho}^{rew} := \{s \in \mathscr{S} : rew(s) \leq \rho\}$. For $\tau > 0$, the following recursion holds (Baier and Katoen, 2008):

$$x_{\Phi_2|\Phi_1}^{\leq\tau,\rho}(s) := \begin{cases} 1 \text{ if } s \in \Phi_2 \text{ and } rew(s) \leq \rho, \\ 0 \text{ if } rew(s) > \rho, \\ \sum_{s' \in \Phi_1} T(s,s') x_{\Phi_2|\Phi_1}^{\leq\tau-1,\rho-rew(s')}(s') \text{ otherwise,} \end{cases} \quad (34)$$

i.e.,

$$x_{\Phi_2|\Phi_1}^{\leq\tau,\rho} := \mathbb{I}_{\mathscr{S}_{\rho}^{rew}}\left(\mathbb{I}_{\Phi_2} + \mathbb{I}_{\Phi_1\setminus\Phi_2} \cdot \hat{T}^{\dagger}\chi^{\leq\tau-1,\rho}\right), \quad (35)$$

where, for each $s \in \mathscr{S}$:

$$\chi^{\leq\tau-1,\rho}(s) := \begin{cases} 1 & \text{if } \rho \leq rew(s), \\ x_{\Phi_2|\Phi_2}^{\leq\tau-1,\rho-rew(s)}(s) & \text{otherwise.} \end{cases} \quad (36)$$

This defines a simple algorithmic scheme where $x_{\Phi_2|\Phi_2}^{\leq\tau,\rho}(s)$ should be computed as in Equation (34) for each $s \in \mathscr{S}$ and $\rho \leq r$ before moving to the subsequent value of τ. According to Equation (34), the recursion is blocked for each $s \in \Phi_2$, with the final probability being equal to one, and for each $s \in \mathscr{S} : rew(s) > \rho$, with the final probability being equal to zero. The recursion is eventually blocked when, for all the reached states $s \in \mathscr{S}$ it holds that either $s \in \Phi_2$ or $rew(s) > \rho$. This always happens within a finite time horizon $\tau \in \mathbb{N}$ because, at each further recursive step, the reward threshold ρ is reduced, for each $s' \in \mathscr{S}$, of a value $rew(s')$ as specified by Equation (36). The final value of $P(s \models \phi_1 \bigcup_{\leq r} \phi_2)$ is then equal to $x_{\Phi_2|\Phi_1}^{\leq\tau,r}$ where $\tau \in \mathbb{N}$ is the time step at which (35) is eventually blocked.

5. Towards an Imprecise PRCTL

In this section we discuss how both PCTL and PRCTL can be extended to an imprecise-probabilistic setting. An *imprecise* PCTL (IPCTL) syntax can be obtained by replacing PCTL probability operator $P_{\nabla b}$ with its lower (upper) variant $\underline{P}_{\nabla b}$ ($\overline{P}_{\nabla b}$) (Section 5.1). The semantics (Section 5.2) is consequently defined as in Section 4.2 with the checking tasks performed in the corresponding IMCs by means of the inference algorithms described in Section 3.2.

IPCTL has been proposed by Troffaes and Škulj (2013). Here, we take advantage of the alternative approaches to the computation of the lower and upper hitting probabilities recently proposed by T'Joens et al. (2019) and corresponding to Equation (9). Compared to the absorbing-state approach in Equation (10), proposed by Troffaes and Škulj (2013), the more recent approach allows to easily achieve an analogous computation for the lower and upper expected cumulative rewards in the underlying IMRM, thus providing an imprecise-probabilistic version of Equation (30). This is the key to define an *imprecise* PRCTL (IPRCTL) whose syntax (Section 5.4), semantics and model checking (Section 5.5) represent the main contribution of this work.

5.1. IPCTL Syntax

The non-probabilistic syntax of IPCTL coincides with that of PCTL. Accordingly, to define IPCTL syntax, we keep the non-probabilistic specification in Equation (12) and, in Equation (11), replace $P_{\nabla b}\phi$ with:

$$\overline{P}_{\nabla b}\psi \mid \underline{P}_{\nabla b}\psi. \quad (37)$$

The informal reading of $\overline{P}_{\nabla b}\psi$ is *the upper bound of the probability to reach a path that satisfies* ψ *is* ∇b. A similar reading for the lower bound is associated with $\underline{P}_{\nabla b}\psi$. Note that IPCTL coincides with PCTL for $\underline{P} = \overline{P}$.

5.2. IPCTL Semantics

The only difference between IPCTL and PCTL semantics is in the probability formulae. Given labelled IMC $\mathscr{M}$ and state $s \in \mathscr{S}$, Equation (20) becomes:

$$\mathscr{M}, s \models \overline{P}_{\nabla b}\psi \text{ iff } \overline{P}(s \models \psi)\nabla b\,, \tag{38}$$

and analogously for $\underline{P}$. As in Section 4.2 for PCTL, the satisfiability check in Equation (38) leads to different inferential tasks in $\mathscr{M}$ depending on ψ.

For $\psi := \bigcirc\phi$, following Equation (21), the lower and upper bounds of $P(S_1 \in \Phi | S_0 = s)$ should be computed. An imprecise-probabilistic version of this equation is achieved by replacing the linear operator in Equation (3) with its non-linear analogous in Equation (7), i.e.,

$$\overline{P}(S_1 \in \Phi | S_0 = s) = (\overline{\mathscr{T}}\mathbb{I}_\Phi)(s)\,, \tag{39}$$

and analogously for $\underline{P}$ and $\underline{\mathscr{T}}$. The case $\psi := \phi_1 \bigcup^{\leq t} \phi_2$ requires instead the computation of the lower and upper bounds of the conditional hitting probability vector $h^{\leq t}_{\Phi_2|\Phi_1}$. Exactly as Equation (22) was obtained as a conditional version of Equation (5), from Equation (9) we can obtain the recursion:

$$\overline{h}^{\leq \tau}_{\Phi_2|\Phi_1} := \mathbb{I}_{\Phi_2} + \mathbb{I}_{\Phi_1 \setminus \Phi_2}\overline{\mathscr{T}}\,\overline{h}^{\leq \tau-1}_{\Phi_2|\Phi_1}\,, \tag{40}$$

for each $\tau = 1, \ldots, t$, and analogously for the upper bound, with the same initialization for both cases, i.e.,

$$\underline{h}^{\leq 0}_{\Phi_2|\Phi_1} = \overline{h}^{\leq 0}_{\Phi_2|\Phi_1} = \mathbb{I}_{\Phi_2}\,. \tag{41}$$

As for PMCs, case $\psi := \phi_1 \bigcup \phi_2$ can be treated analogously to $\phi_1 \bigcup^{\leq t} \phi_2$ when $t \to \infty$. The value of $\overline{P}(s \models \phi_1 \bigcup \phi_2)$ hence corresponds to $\lim_{t\to\infty} \overline{h}^{\leq t}_{\Phi_2|\Phi_1}$. Proposition 16 in (Krak et al., 2019) proves that such limit exists and corresponds to the fixed point of $h^{\leq \tau}_{A\subseteq\mathscr{S}}$. Notice that, the result can be considered valid also for the recursion in Equation (40) because the additional condition that all the states visited before reaching Φ_2 are in Φ_1 does not alter the validity of the proof. Notice also that the result by Krak et al. (2019) is obtained within a game-theoretic approach to IMCs, which is different from the sensitivity-analysis approach we adopt here. However, as the author clearly points out, the result is to be considered valid for all the foundational approaches to IMCs (Krak et al., 2019).

Regarding computational complexity, the linear programming tasks in Equation (40) take only polynomial time with respect to $|\mathscr{S}|$, while the maximum number of iterations is t. This shows that shifting to imprecise probabilities does not affect the overall computational complexity of the task. The only computational issue with computing $\overline{P}(s \models \phi)$ (as well as $\underline{P}(s \models \phi)$) concerns the nesting depth of ϕ, i.e., the number of iterated nested formulae in ϕ. The overall complexity of computing $\overline{P}(s \models \phi)$, indeed, is exponential with respect to the nesting depth of ϕ. However, the same holds in the precise case and this does not significantly affect applications, where small nesting depths are typically considered.

5.3. A IPCTL Example

Example 2 *Consider an imprecise probabilistic version of the model M discussed in Example 1. The same, precise, transition probabilities are kept as in Figure 1 apart from* $P(S_{t+1}|s_t = \text{try})$. *For a sensitivity analysis parametrized by* $\varepsilon \in [0,1]$, *we replace such (conditional) PMF with the CS induced by the linear constraint:*

$$P(\text{delivered}|\text{try}) \in [(1-\varepsilon)0.9 - \varepsilon, (1-\varepsilon)0.9 + \varepsilon]\,. \tag{42}$$

with the impossible transitions remaining impossible.[3] *This makes the model a IMC to be used to answer IPCTL queries, such as* "the upper probability to lose a message within seven time steps is less than or equal to 0.25". *Having this formula for the upper probability satisfied ensures that any MC consistent with the IMC would satisfy the analogous formula in Example 1, thus providing the desired sensitivity analysis. Figure 2 shows the corresponding (upper) hitting probabilities for increasing values of* ε *computed by means of the recursion in Equation* (9) *and two different time horizons. Even for the higher perturbation level we consider* ($\varepsilon = 0.03$) *within seven time steps the hitting probability remains under the threshold level. In the limit of longer chains, for this model, both bounds converge to one.*

5.4. IPRCTL Syntax

We are now in the condition of extending IPCTL syntax in order to cope with IMRMs. We call the corresponding language IPRCTL. The IPRCTL syntax is obtained by augmenting the IPCTL syntax in Section 5.1 with the PRCTL path formulae in Equation (25), while expression $E_{\nabla r}\phi$ in Equation (24) is replaced by:

$$\underline{E}_{\nabla r}\phi \mid \overline{E}_{\nabla r}\phi\,, \tag{43}$$

whose informal reading is as for PRCTL in Section 4.4.

5.5. IPRCTL Semantics

The IPRCTL semantics is obtained extending the IPCTL semantics with respective satisfiability conditions for expected reward and reward-bounded probabilities.

3. As only two states are possible, the bounds on the probability of $\{\text{lost}\}$ can be also induced by ε-contamination. Deterministic PMFs are unaffected by the contamination.

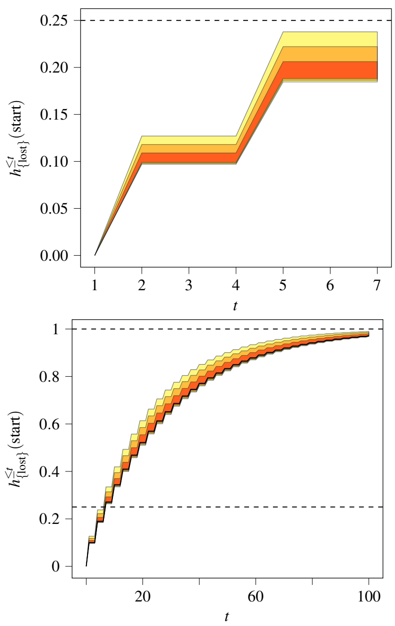

Figure 2: Hitting probability ranges for increasing perturbation levels. Red, orange and yellow plots correspond to $\varepsilon \in \{0.01, 0.02, 0.03\}$.

Expected Cumulative Reward. Given an IMRM $(\mathscr{M}, rew)$ and state $s \in \mathscr{S}$, the analogous of Equation (27) corresponds to condition:

$$(\mathscr{M}, rew), s \models \overline{E}_{\nabla r}\phi \text{ iff } \overline{E}[Rew_{\Phi}^{\leq t}](s)\nabla r, \quad (44)$$

and analogously for $\underline{E}_{\nabla r}$ and $\underline{E}[Rew_{\Phi}^{\leq t}]$. Those lower and upper expected cumulative reward arrays $\underline{E}[Rew_{\Phi}^{\leq t}]$ and $\overline{E}[Rew_{\Phi}^{\leq t}]$ represent the lower and upper bounds of the precise expectations in Equation (28) with respect to the CSs in the specification of $\mathscr{M}$ (see Section 3.2).

Exactly as for the derivation of Equation (40), we rely on the results by T'Joens et al. (2019) to achieve an imprecise-probabilistic version of the recursion in Equation (30). This simply corresponds to:

$$\overline{E}[Rew_{\Phi}^{\leq \tau}] := rew + \mathbb{I}_{\Phi^c}\overline{\mathscr{T}}\,\overline{E}[Rew_{\Phi}^{\leq \tau-1}], \quad (45)$$

for each $\tau = 1, \ldots, t$, and analogously for the lower expectation, with initialization for both cases:

$$\underline{E}[Rew_{\Phi}^{\leq 0}] = \overline{E}[Rew_{\Phi}^{\leq 0}] = rew. \quad (46)$$

The recursion in Equation (45), exactly as that in Equation (9), requires $t \in \mathbb{N}$ iterative applications of the non-linear transition operator in Equation (7). As for Equation (40), each further iterative application of (45) requires to solve $|\mathscr{S}|$ linear programming tasks. Also in this case, hence, shifting to imprecise probabilities does not increases the overall computational complexity with respect to $|\mathscr{S}|$. Notice that, the complexity with respect to the nesting depth of ϕ results to be exponential also in this case, but the same practical considerations stated for Equations (40) hold.

Bounded-Reward Probability. Given IMRM $(\mathscr{M}, rew)$ and state $s \in \mathscr{S}$, the following condition holds:

$$(\mathscr{M}, rew), s \models \overline{P}_{\nabla b}\delta \text{ iff } \overline{P}(s \models \phi_1 \bigcup_{\leq r} \phi_2)\nabla r, \quad (47)$$

where the event on the right-hand side is as in Equation (31) and an analogous condition holds for the lower probability. By a discussion similar to that in Section 4.5, we obtain a recursive relation analogous to Equation (35) for the upper probabilities, denoted here as $\overline{x}_{\Phi_2|\Phi_1}^{\leq \tau, \rho}$, by simply replacing the linear operator with its non-linear, upper, version, i.e.,

$$\overline{x}_{\Phi_2|\Phi_1}^{\leq \tau, \rho} := \mathbb{I}_{\mathscr{S}_\rho^{rew}}\left(\mathbb{I}_{\Phi_2} + \mathbb{I}_{\Phi_1 \setminus \Phi_2} \cdot \overline{\mathscr{T}}\,\overline{\chi}^{\leq \tau-1, \rho}\right), \quad (48)$$

where $\overline{\chi}^{\leq \tau-1, \rho}$ is obtained as in Equation (36) but from the upper probabilities for the same time step. An analogous derivation holds for the lower bound. Both the initializations are as in the precise case in Equation (33). Notice that, Equation (48) can be obtained from Equation (40) by simply including the indicator vector $\mathbb{I}_{\mathscr{S}_\rho^{rew}}$, which blocks the recursion for each $s \in \mathscr{S} : rew(s) > \rho$, and by replacing the upper hitting probability with the function $\overline{\chi}^{\leq \tau-1, \rho}$. Note that $\overline{\chi}^{\leq \tau-1, \rho}$ coincides with $\overline{x}_{\Phi_2|\Phi_1}^{\leq \tau, \rho - rew(s)}$. For each further iteration hence, the reward threshold ρ is reduced, for each $s \in \mathscr{S}$, of a value $rew(s)$. The recursion is eventually blocked when all the reached states $s' \in \mathscr{S}$ are either such that $s' \in \Phi_2$ or such that $rew(s') > \rho$. Also in this case this always happens for a finite time horizon $t \in \mathbb{N}$ because the reward threshold ρ is reduced at each further iteration.

Concerning computational complexity, since Equation (48) also requires an application of the dual imprecise transition operator $\mathscr{T}^{\dagger}$ for each further recursive step, the same considerations stated for Equation (40) hold. The overall computational complexity of (48) is therefore polynomial with respect to $|\mathscr{S}| \cdot t$, where t is the total number of iterations occured until any further iteration is blocked.

6. A Case Study on IPRCTL

As a very first IPRCTL application we perform a sensitivity analysis in the MRM originally proposed by McClean et al. (1998). In that paper expected cumulative rewards are used to estimate the cost of annual recovery of geriatric patients. Let us briefly describe their model and report the results of our IPRCTL-based sensitivity analysis.[4]

In the considered geriatric departments, there are two kinds of recovery: short-term recoveries for acute cares have a daily cost estimated as £100, while long-term recoveries cost £50. From a cumulative perspective, long-term recoveries are more expensive, since those patients typically

4. Code available at github.com/IDSIA-papers/2021-ISIPTA-IPRCTL.

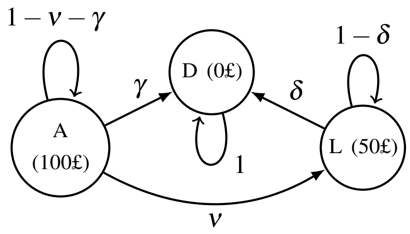

Figure 3: Transitions in a three-state MRM.

remain in the hospitals for longer periods. The scenario can be naturally described as a PMC M whose three states are in one-to-one correspondence with the singletons of the three atomic propositions: A (acute care), L (long stay), and D (discharge or death). The first two states represent short and long-term recoveries, while the latter represents the end of a recovery. D should be regarded as an absorbing state and a parametrized version of the transition matrix for this model is in Figure 3. The parameters have the following interpretation: the *conversion rate* ν corresponds to the probability of passing from a short-term to a long-term recovery, while the *dismissing rates* γ and δ correspond to the probability of being discharged/die, respectively, in a short- and long-term recovery. Rates γ, ν and δ vary depending on the patient and disease. An assessment of these parameters for different departments is in Table 2.

Rate (%)	Dep.1	Dep.2	Dep.3
γ	1.750	3.540	2.810
ν	0.031	0.187	0.149
δ	0.120	0.130	0.180

Table 2: Conversion and dismissing rates.

The reward *rew* associated with each state represents the daily cost per patient. In a scale where one corresponds to one pound, we already assumed $rew(A) = 100$, $rew(L) = 50$, while the reward of D is set to zero. Under these assumption, the corresponding MRM (M, rew) is used to predict the expected annual cumulated cost of each department. This is obtained from the initial numbers $k(A)$ and $k(L)$ of patients in acute care and long stay:

$$cost := \sum_{s \in \{A,L\}} k(s) \cdot E[Rew_D^{\leq n}](s)\,. \tag{49}$$

This cost can be therefore computed separately for each department. A sensitivity analysis with respect to the transition probabilities in Table 2 consists in considering the interval spanned by the extreme values of each one of these parameters and consequently define a credal transition matrix. The corresponding IMRM can be used to compute the lower and upper cumulative costs for different values of s, to be eventually combined as in Equation (49) with the aggregated numbers about the patients of the three departments. Figure 4 show the result for a horizon of 20 years.

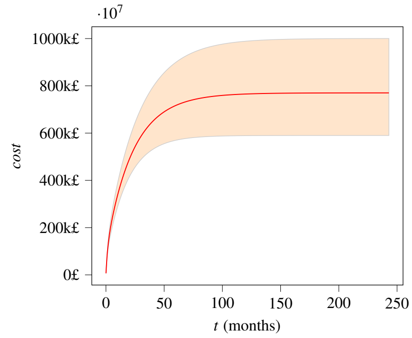

Figure 4: Aggregated cumulative costs and bounds.

Finally, assume that departments are sustainable if and only if the total cumulative cost per patient until the patient is discharged or dies is less than or equal to a given threshold $r := 15'000$ in a time horizon of one year. This corresponds to $E_{\leq 15'000}\{D\}$ in PRCTL and $\overline{E}_{\leq 15'000}\{D\}$ in IPRCTL. Following Equations (27) and (44) we can check this formula by computing the corresponding expected cumulative rewards, whose values for the different starting states are depicted in Table 3.

s	$E[Rew_{\{D\}}^{\leq 365}](s)$			$\underline{E}$	$\overline{E}$
Dep.	1	2	3		
A	5′832	3′372	4′009	2′910	6′421
L	14′850	14′600	13′437	13′437	14′850

Table 3: Yearly cumulative costs for single patients.

Both formulae are satisfied, thus making the sustainability of each department robust even with respect to an imprecise evaluation of the conversion and dismissing rates.

7. Conclusions

An imprecise-probabilistic generalization of PRCTL, called IPRCTL, has been presented together with inference algorithms to compute expected cumulative rewards and bounded-reward probabilities. IPRCTL represent a first step toward the development of an imprecise PCTL based on *imprecise Markov decision processes*. Although such processes have been already considered (e.g., Delgado et al. (2011)) their application to model checking is an open area of investigation. The same holds for *imprecise continuous-time Markov chains*, that have been subject of intense research in the very last years (e.g., Krak et al. (2017)) and whose application to model checking represent an open challenge we want to explore as a necessary future work.

References

Suzana Andova, Holger Hermanns, and Joost-Pieter Katoen. Discrete-time rewards model-checked. In *International Conference on Formal Modeling and Analysis of Timed Systems*, pages 88–104. Springer, 2003.

Christel Baier and Joost-Pieter Katoen. *Principles of model checking*. MIT Press, 2008.

Christel Baier, Boudewijn Haverkort, Holger Hermanns, and Joost-Pieter Katoen. Model checking continuous-time Markov chains by transient analysis. *LNCS*, 1855: 358–372, 07 2000.

Christel Baier, Christian Hensel, Lisa Hutschenreiter, Sebastian Junges, Joost-Pieter Katoen, and Joachim Klein. Parametric Markov chains: PCTL complexity and fraction-free gaussian elimination. *Information and Computation.*, 272:104504, 2020.

Nikola Benes, Lubos Brim, Samuel Pastva, and David Safránek. Model checking approach to the analysis of biological systems. In Pietro Liò and Paolo Zuliani, editors, *Automated Reasoning for Systems Biology and Medicine*, volume 30 of *Computational Biology*, pages 3–35. Springer, 2019.

Béatrice Bérard, Michel Bidoit, Alain Finkel, François Laroussinie, Antoine Petit, Laure Petrucci, Philippe Schnoebelen, and Pierre McKenzie. *Systems and Software Verification, Model-Checking Techniques and Tools*. Springer, 2001.

Lubos Brim, Milan Ceska, and David Safránek. Model checking of biological systems. In Marco Bernardo, Erik P. de Vink, Alessandra Di Pierro, and Herbert Wiklicky, editors, *Formal Methods for Dynamical Systems - 13th International School on Formal Methods for the Design of Computer, Communication, and Software Systems, SFM 2013*, volume 7938 of *Lecture Notes in Computer Science*, pages 63–112. Springer, 2013.

Taolue Chen, Giuseppe Primiero, Franco Raimondi, and Neha Rungta. A computationally grounded, weighted doxastic logic. *Studia Logica*, 104(4):679–703, 2016.

Martin Davis and Hilary Putnam. A computing procedure for quantification theory. *Journal of the ACM (JACM)*, 7 (3):201–215, 1960.

Conrado Daws. Symbolic and parametric model checking of discrete-time Markov chains. In Zhiming Liu and Keijiro Araki, editors, *Theoretical Aspects of Computing - ICTAC 2004, First International Colloquium, Guiyang, China, September 20-24, 2004, Revised Selected Papers*, volume 3407 of *Lecture Notes in Computer Science*, pages 280–294. Springer, 2004.

Gert de Cooman, Jasper De Bock, and Stavros Lopatatzidis. Imprecise stochastic processes in discrete time: global models, imprecise Markov chains, and ergodic theorems. *International Journal of Approximate Reasoning*, 76: 18–46, 2016.

Karina Valdivia Delgado, Leliane Nunes De Barros, Fabio Gagliardi Cozman, and Scott Sanner. Using mathematical programming to solve factored Markov decision processes with imprecise probabilities. *International journal of approximate reasoning*, 52(7):1000–1017, 2011.

Hans Hansson and Bengt Jonsson. A logic for reasoning about time and reliability. *Formal Aspects of Computing*, 6(5):512–535, 1994.

David Henriques, João Martins, Paolo Zuliani, André Platzer, and Edmund Clarke. Statistical model checking for Markov decision processes. In *Proceedings - 2012 9th International Conference on Quantitative Evaluation of Systems, QEST 2012*, pages 84–93, 09 2012.

Catalin-Andrei Ilie and James B. Worrell. Parametric model checking continuous-time Markov chains. In Emilio Muñoz-Velasco, Ana Ozaki, and Martin Theobald, editors, *27th International Symposium on Temporal Representation and Reasoning, TIME 2020, September 23-25, 2020, Bozen-Bolzano, Italy*, volume 178 of *LIPIcs*, pages 7:1–7:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.

J-P Katoen, Maneesh Khattri, and IS Zapreevt. A Markov reward model checker. In *Second International Conference on the Quantitative Evaluation of Systems (QEST'05)*, pages 243–244. IEEE, 2005.

Thomas Krak, Jasper De Bock, and Arno Siebes. Imprecise continuous-time Markov chains. *International Journal of Approximate Reasoning*, 88:452–528, 2017.

Thomas Krak, Natan T'Joens, and Jasper De Bock. Hitting times and probabilities for imprecise Markov chains. In *International Symposium on Imprecise Probabilities: Theories and Applications*, pages 265–275. PMLR, 2019.

Sally I. McClean, B. McAlea, and Peter H. Millard. Using a Markov reward model to estimate spend-down costs for a geriatric department. *J. Oper. Res. Soc.*, 49(10): 1021–1025, 1998.

Daniel Revuz. *Markov chains*. Elsevier, 2008.

Natan T'Joens, Thomas E. Krak, Jasper De Bock, and Gert de Cooman. A recursive algorithm for computing inferences in imprecise Markov chains. In Gabriele Kern-Isberner and Zoran Ognjanovic, editors, *Symbolic* *and Quantitative Approaches to Reasoning with Uncertainty, 15th European Conference, ECSQARU 2019, Proceedings*, volume 11726 of *Lecture Notes in Computer Science*, pages 455–465. Springer, 2019.

Matthias C. M. Troffaes and Damjan Škulj. Model checking for imprecise Markov chains. In Fabio Cozman, Thierry Denoeux, Sebastien Destercke, and Teddy Seidenfeld, editors, *ISIPTA '13 : Proceedings of the Eighth International Symposium on Imprecise Probability: Theories and Applications*, pages 337–344. Society for Imprecise Probability: Theories and Applications (SIPTA), July 2013.

Peter Walley. *Statistical reasoning with imprecise probabilities*. Chapman and Hall, 1991.

Lijun Zhang, Holger Hermanns, and David Jansen. Logic and model checking for hidden Markov models. In *International Conference on Formal Techniques for Networked and Distributed Systems*, volume 3731, pages 98–112. Springer, 10 2005.