Quantitative Analysis of Interval Markov Chains

Giovanni Bacci[1], Benoît Delahaye[2], Kim G. Larsen[1], and Anders Mariegaard[1(✉)]

[1] Department of Computer Science, Aalborg University, Aalborg, Denmark
{giovbacci,kgl,am}@cs.aau.dk
[2] Université de Nantes/LS2N UMR CNRS, 6004 Nantes, France
benoit.delahaye@univ-nantes.fr

Abstract. Interval Markov chains (IMCs), as first introduced by Larsen and Jonsson in 1991 are succinct specifications for probabilistic systems that generalise Markov chains (MCs) by allowing state transition probabilities to lie within an interval. In this work, we address the study of IMCs in a quantitative setting by extending the notion of IMCs by associating with each state a reward that is gained when leaving the state. Specifically, we compare three different semantic interpretations proposed in the literature (once-and-for-all, interval Markov decision process and at-every-step) in the context of model-checking rPCTL, an extension of PCTL where each path-formula is equipped with the specification of a bound on the accumulated reward. We prove that for the full logic, the three semantics are not equivalent, but for the fragment of reward-bounded reachability properties, the interval Markov decision process semantics and the at-every-step semantics are equivalent. Finally, we discuss model-checking algorithms for the three semantics by reduction to the model-checking problem for parametric Markov chains.

1 Introduction

The early work of Bengt Jonsson contains several contributions to the verification of distributed systems [20]. This still very active research direction [7] has been dominated by two schools: the North American school stressing automata and temporal logics, and the European school with focus on process algebra and behavioural equivalences. Both directions have their pros and cons with respect to compositionality and refinement: in the process algebraic approach compositional reasoning was guaranteed by congruence properties of the considered equivalences. However, specifications are typically very explicit being single equivalence classes leaving no room for a stepwise refinement process. In contrast, in the temporal logic approach logical implication between specifications provides the basis for stepwise refinement. However, it is notoriously hard to derive logical properties of composite systems from properties of their components, see [2,30]. Within the process algebraic approach, the introduction of Modal Transition Systems [27] (MTS) may be seen as a step towards support of a true stepwise refinement process. In MTS the transitions of a labelled

E.-R. Olderog et al. (Eds.): Jonsson Festschrift, LNCS 13030, pp. 57–77, 2021.
https://doi.org/10.1007/978-3-030-91384-7_4

transition system are classified as either mandatory (must) of optional (may) leading to a modal refinement precongruence generalizing the strict behavioural equivalences.

At the same time, probabilistic extensions of process algebra were introduced, e.g. [15], including the introduction of probabilistic bisimulation [25,26]. In collaboration with the last author of this paper (during a nice sabbatical at SICS in 1990), Bengt Jonsson quickly followed up with a probabilistic extension of MTS [22], originally termed Probabilistic Specifications, but by now better known as Interval Markov Chains (IMC). It is fair to say that IMC has inspired much subsequent research (including this paper).

On the temporal logical side, the introduction of PCTL in the seminal paper [17,21] by Bengt Jonsson and Hans Hansson is by now considered a prime logic for specifying properties of probabilistic systems. Since its introduction significant effort has been made towards efficient model checking algorithms for PCTL. However, there are still open problems foremost the question of decidability of satisfiability. One research direction that we will pursue in this paper is that of model checking PCTL with respect to IMC.

Our Contribution. We consider interval Markov reward models (IMRMs), a class of models that extend interval Markov chains by assigning a (positive) reward to each state. For regular IMCs, three distinct semantics have been proposed in the literature: the once-and-for-all semantics [5], the interval Markov decision process (IMDP) semantics [5,8,29] and the at-every-step semantics [22]. We provide a natural extension of the three semantics to IMRMs and investigate the differences between the three semantics in the context of model-checking. For this we consider the logic Probabilistic CTL (PCTL) [18] with reward-bounded path-formulae (rPCTL). For a given fragment of the logic, we say that two semantics are *equivalent* if for some IMRM specification and rPCTL formula, whenever there exists a satisfying model of one semantics, there exists a satisfying model of the other semantics.

Our contribution is twofold. The first part of the paper concerns the comparison of the above mentioned semantics:

(i) we prove that the three semantics are not equivalent with respect to the full fragment of rPCTL;
(ii) if one restricts the attention to probabilistic bounded reachability queries (a) we show that the once-and-for-all semantics and the IMDP-semantics are not equivalent, whereas (b) the IMDP-semantics and at-every-step semantics are.

The result in (i) can be seen as a generalisation of a similar result by Bart et al. [5] for IMCs against PCTL properties. In contrast to [5], where three IMCs semantics where shown to be equivalent with respect to reachability queries, we show that such an equivalence does not generalise to IMRMs.

In the second part of the paper we present algorithms for model-checking IMRMs for the three semantics. For the full logic and the once-and-for-all semantics, we present a reduction to the (existential) model-checking problem for para-

metric Markov reward models [1] with interval-constraints on the parameters. As for the IMDP semantics, we devise a reduction to the model-checking problem of IMRMs using the once-and-for-all semantics.

Notably, thanks to the semantic equivalence result relative to reachability queries mentioned earlier, such a reduction solves also the model checking problem against reachability queries when one interprets IMRMs using the at-every-step semantics. However, model checking generic rPCTL properties with respect to the at-every-step semantics still remains an open problem.

Related Work. Since their introduction by Jonsson and Larsen [22], IMCs have been investigated from different perspectives. In particular, [12,13] tackles the computational complexity of several decision problems, such as deciding whether or not an IMC has an implementation (the *consistency* problem) and whether the set of implementations of one IMC is entailed by the set of implementation by another IMC (thorough refinement). For model-checking, [6] considers LTL model-checking w.r.t IMCs with the once-and-for-all semantics, while [8,29] presents algorithms for verifying PCTL properties for both the once-and-for-all semantics as well as the IMDP semantics. The work in [8] also considers general ω-regular properties. From a computational complexity perspective, Chen et al. [9] proved that the two variants of the PCTL model-checking problem w.r.t. the once-and-for-all semantics and the IMDP semantics are both P-complete.

Another body of research is the work on *parametric* IMCs (PIMCs) [5,11, 14,28] where, instead of an interval, one can instead place a parameter. All the problems for IMCs can then be re-cast in two variants for PIMCs, depending on the quantification over the parameters (existential or universal). Closest to our work is [5], in which the equivalence between the three different semantics is investigated for IMCs. In the same paper, verifying a probabilistic reachability property for a given PIMCs is reduced to solving a constraint satisfaction problem.

2 Preliminaries and Notation

We denote by $\mathbb{R}$, $\mathbb{Q}$, and $\mathbb{N}$ respectively the set of real, rational, and natural numbers. Given a binary relation $R \subseteq X \times Y$ and $x \in X$, we define the projection of R on x as $R(x) = \{y \in Y \mid (x, y) \in R\}$, and we denote by R^{-1} the inverse of R, i.e., $R^{-1} = \{(y, x) \mid (x, y) \in R\}$.

For a finite nonempty set X, $\mu \colon X \to [0, 1]$ is a probability distribution on X if $\sum_{x \in X} \mu(x) = 1$. Moreover μ is extended to sets $Y \subseteq X$ as $\mu(Y) = \sum_{y \in Y} \mu(y)$. We write $\mathcal{D}(X)$ for the set of probability distributions on X. For $\mu \in \mathcal{D}(X)$ we define the support of μ as $\texttt{support}(\mu) = \{x \in X \mid \mu(x) > 0\}$.

3 Markov Reward Models

In this section we recall the definitions of Markov reward model (MRM), probabilistic reward bisimulation, and Reward-Bounded Probabilistic CTL (rPCTL).

For the rest of the paper, we fix a countable set of atomic propositions A.

Definition 1 (Markov Reward Model). *A* Markov reward model *is a tuple* $\mathcal{M} = (S, s_0, \pi, \rho, \ell)$ *consisting of a finite set of states* S*, an initial state* $s_0 \in S$*, a transition probability function* $\pi \colon S \to \mathcal{D}(S)$*, a state-reward function* $\rho \colon S \to \mathbb{N}_{>0}$ *assigning to each state a positive reward*[1] *and a labelling function* $\ell \colon S \to 2^A$ *mapping states to atomic propositions.*

Intuitively, if $\mathcal{M}$ is in state s it moves to state s' with probability $\pi(s)(s')$, thereby receiving the reward $\rho(s)$. In this sense $\mathcal{M}$ can be seen as a state-machine that generates paths of states starting from the initial state s_0.

We denote by $G_{\mathcal{M}} = (S, \to)$ the *underlying labelled graph of* $\mathcal{M}$, where $s, s' \in S$ are connected by a labelled directed edge $s \xrightarrow{p,r} s'$ if and only if $p = \pi(s)(s') > 0$ and $r = \rho(s)$. We will assume without loss of generality that all states of $\mathcal{M}$ are reachable from the initial state s_0 in its underlying graph. For $s \in S$ we define the set of *successors* of s as $\texttt{succ}(s) = \texttt{support}(\pi(s))$.

Example 1. Figures 1b–d depicts three MRMs. Consider the MRM $\mathcal{M}_o = (T_o, t_0, \pi_o, \rho_o, \ell^{\mathcal{M}_o})$ in Fig. 1b. States $T_o = \{t_i \mid 0 \leq i \leq 4\}$ are visualised by a circle split in two, with the name of a state t_i at the top and the reward $\rho_o(t_i)$ at the bottom. The initial state t_0 is identified by a double-stroke border. State labels $\ell^{\mathcal{M}_o}(t_i)$ are visualised next to the state t_i unless the set is empty, in which case the set is omitted. From the underlying graph $G_{\mathcal{M}_o}$ we have $\texttt{succ}(t_0) = \{t_1, t_2\}$, $t_0 \xrightarrow{0.3,1} t_1$ and $t_0 \xrightarrow{0.7,1} t_2$.

A *path* is an infinite sequence of states $\sigma = s_0, s_1, \ldots \in S^\omega$; for $j \in \mathbb{N}$, we denote by $\sigma[j]$ the $(j+1)$-th state of σ, i.e., $\sigma[j] = s_j$ and by $\mathcal{W}(\sigma)(j) = \sum_{i=0}^{j-1} \rho(s_i)$ the accumulated reward of σ after j transitions. For a finite path $\sigma = s_0, \ldots, s_j \in S^*$ we define the length of σ as $|\sigma| = j$.

To associate probabilities to measurable events, we adopt the classical cylinder set construction from [4, Chapter 10]. For $w \in S^*$, the cylinder set of w is the set of all paths having prefix w, i.e., $cyl(w) = wS^\omega$. Given an *initial probability distribution* $\iota \in \mathcal{D}(S)$, we define the probability space $(S^\omega, \Sigma_{\mathcal{M}}, \mathbb{P}_\iota^{\mathcal{M}})$, where $\Sigma_{\mathcal{M}}$ is the smallest σ-algebra that contains all the cylinder sets, and $\mathbb{P}_\iota^{\mathcal{M}}$ is the unique probability measure such that, for all $w = s_0 \cdots s_n \in S^*$,

$$\mathbb{P}_\iota^{\mathcal{M}}(cyl(w)) = \iota(s_0) \cdot \textstyle\prod_{0 \leq i < n} \pi(s_i)(s_{i+1}) \,.$$

When ι is the Dirac distribution pointed at s, i.e. $\iota(s) = 1$, we write $\mathbb{P}_s^{\mathcal{M}}$, or just $\mathbb{P}_s$ when $\mathcal{M}$ is clear from the context. Similarly, we may write $\mathbb{P}^{\mathcal{M}}$ as a shorthand for $\mathbb{P}_{s_0}^{\mathcal{M}}$ when s_0 is the initial state of $\mathcal{M}$.

Definition 2 (Bisimulation). *Let* $\mathcal{M} = (S, s_0, \pi, \rho, \ell)$ *be an MRM. An equivalence relation* $\mathcal{R} \subseteq S \times S$ *is a* probabilistic reward bisimulation *for* $\mathcal{M}$ *if whenever* $(s, t) \in \mathcal{R}$*, then (i)* $\rho(s) = \rho(t)$*, (ii)* $\ell(s) = \ell(t)$*, and (iii)* $\pi(s)(C) = \pi(t)(C)$ *for all* $C \in S_{/\mathcal{R}}$*.*

[1] All results presented in this paper can be generalized to MRMs having positive rational state rewards by multiplying the vector of rewards by a suitably large scaling factor.

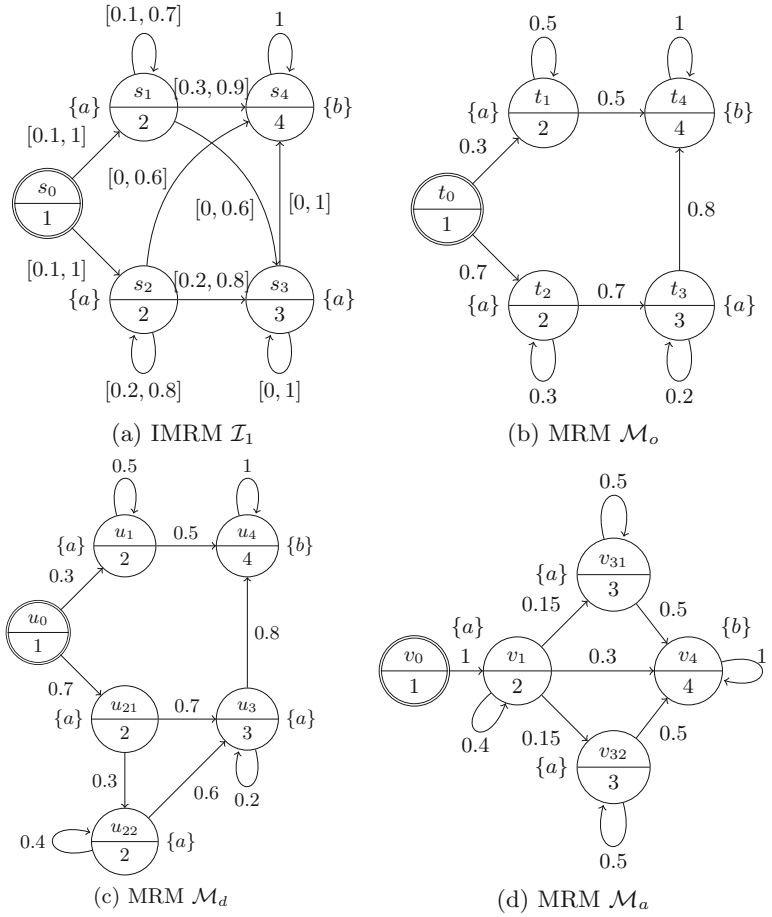

Fig. 1. IMRM $\mathcal{I}_1$ and implementations $\mathcal{M}_o \in [\![\mathcal{I}_1]\!]_o$, $\mathcal{M}_d \in [\![\mathcal{I}_1]\!]_d$ and $\mathcal{M}_a \in [\![\mathcal{I}_1]\!]_a$.

Two states $s, s' \in S$ are probabilistic bisimilar, written $s \sim s'$, if they are related by some probabilistic bisimulation. By abuse of notation we may write $\mathcal{M} \sim \mathcal{M}'$ to indicate that the initial states of the MRMs $\mathcal{M}$ and $\mathcal{M}'$ are bisimilar w.r.t. their disjoint union.

We now present an extension of probabilistic CTL (PCTL) [18], namely reward-bounded PCTL (rPCTL), where the next and the until operators are equipped with the specification of a finite bound on the accumulated reward. As any CTL-based logic, rPCTL allows for state formulae describing properties about states in an MRM and path formulae describing properties about paths in an MRM. State formulae Φ and path formulae Ψ are formed according to the following abstract syntax:

$$\Phi ::= \text{true} \mid a \mid \neg\Phi \mid \Phi \wedge \Phi \mid \mathcal{P}_{\bowtie\lambda}(\Psi)$$
$$\Psi ::= \mathsf{X}_{\trianglelefteq k}\,\Phi \mid \Phi\,\mathsf{U}_{\trianglelefteq k}\,\Phi$$

where $a \in A$, $\bowtie = \{<, \leq, \geq, >\}$, $\trianglelefteq = \{\leq, =, \geq\}$, $\lambda \in \mathbb{Q} \cap [0,1]$, and $k \in \mathbb{N}$. We denote by $\mathtt{rPCTL}$ the set of all rPCTL state-formulae.

Given an MRM $\mathcal{M} = (S, s_0, \pi, \rho, \ell)$, a state $s \in S$, and a path $\sigma \in S^\omega$, we write $\mathcal{M}, s \models \Phi$ (resp. $\mathcal{M}, \sigma \models \Psi$) to indicate that s satisfies the state formula Φ (resp. the path σ satisfies the path formula Ψ). The *satisfiability relation* $\models$ is inductively defined as:

$$\begin{array}{lll}
\mathcal{M}, s \models \text{true} & & \text{always} \\
\mathcal{M}, s \models a & \text{iff} & a \in \ell(s) \\
\mathcal{M}, s \models \neg\Phi & \text{iff} & \mathcal{M} \not\models \Phi \\
\mathcal{M}, s \models \Phi_1 \wedge \Phi_2 & \text{iff} & \mathcal{M}, s \models \Phi_1 \text{ and } \mathcal{M}, s \models \Phi_2 \\
\mathcal{M}, s \models \mathcal{P}_{\bowtie\lambda}(\Psi) & \text{iff} & \mathbb{P}_s(\{\sigma \in S^\omega \mid \mathcal{M}, \sigma \models \Psi\}) \bowtie \lambda \\
\mathcal{M}, \sigma \models \mathsf{X}_{\trianglelefteq k}\Phi & \text{iff} & \rho(\sigma[0]) \trianglelefteq k \text{ and } \mathcal{M}, \sigma[1] \models \Phi \\
\mathcal{M}, \sigma \models \Phi_1 \mathsf{U}_{\trianglelefteq k} \Phi_2 & \text{iff} & \exists j \geq 0.\, \mathcal{W}(\sigma)(j) \trianglelefteq k, \\
 & & \mathcal{M}, \sigma[j] \models \Phi_2 \text{ and} \\
 & & \forall i < j.\, \mathcal{M}, \sigma[i] \models \Phi_1.
\end{array}$$

As usual, we derive the operators false, $\vee$, and $\rightarrow$ as $\text{false} := \neg\text{true}$, $\Phi_1 \vee \Phi_2 := \neg(\neg\Phi_1 \wedge \neg\Phi_2)$, and $\Phi_1 \rightarrow \Phi_2 := \neg\Phi_1 \vee \Phi_2$. Moreover, we define the k-bounded reachability operator as $\Diamond_{\trianglelefteq k}\Phi := \text{true}\,\mathsf{U}_{\trianglelefteq k}\,\Phi$.

The satisfiability relation extends naturally to finite paths: a finite path $\sigma \in S^*$ satisfies a path-formula Ψ if and only if all the infinite paths in the cylinder-set $cyl(\sigma)$ satisfy Ψ. If the MRM is clear from the context, we sometimes write $s \models \Phi$ instead of $\mathcal{M}, s \models \Phi$. We may also write $\mathcal{M} \models \Phi$ as a shorthand for $\mathcal{M}, s_0 \models \Phi$ and $\mathbb{P}^{\mathcal{M}}(\Psi)$ as a shorthand for $\mathbb{P}^{\mathcal{M}}_{s_0}(\{\sigma \in S^\omega \mid \mathcal{M}, \sigma \models \Psi\})$, where s_0 is the initial state of $\mathcal{M}$.

Example 2. Consider the three MRMs $\mathcal{M}_o, \mathcal{M}_d$ and $\mathcal{M}_a$ depicted in Figs. 1b–d and let $\Phi = \mathcal{P}_{\geq 0.15}(\Diamond_{\leq 3}\, b)$. By rPCTL semantics we have $\mathcal{M}_o \models \Phi$, witnessed by the path t_0, t_1, t_4 and similarly, $\mathcal{M}_d \models \Phi$ and $\mathcal{M}_a \models \Phi$. If the probability threshold is increased from 0.15 to 0.3, $\mathcal{M}_o$ and $\mathcal{M}_d$ no longer satisfy the formula, i.e. for formula $\Phi' = \mathcal{P}_{\geq 0.3}(\Diamond_{\leq 3}\, b)$, we have $\mathcal{M}_o \not\models \Phi'$, $\mathcal{M}_d \not\models \Phi'$ but $\mathcal{M}_a \models \Phi'$.

For $s, s' \in S$, we say that s and s' are logically equivalent w.r.t. rPCTL, written $s \cong_{\mathtt{rPCTL}} s'$, if

$$\forall \Phi \in \mathtt{rPCTL}.\, \mathcal{M}, s \models \Phi \iff \mathcal{M}, s' \models \Phi\,.$$

The following theorem states that probabilistic bisimilarity equals logical equivalence w.r.t rPCTL.

Theorem 1. *Let $\mathcal{M} = (S, s_0, \pi, \rho, \ell)$ be an MRM and $s, s' \in S$. Then, $s \sim s' \iff s \cong_{\mathtt{rPCTL}} s'$.*

4 Interval Markov Reward Models

In this section we introduce the notion of interval Markov reward model (IMRM) and present three distinct semantic interpretations of IMRMs, comparing their expressivity with respect to rPCTL.

Before defining IMRMs, it is convenient to introduce some notation. We write $\mathbb{I}$ for the set of all non-empty closed interval subsets of $[0,1]$, and $\mathcal{D}_{\mathbb{I}}(X) = \{f \mid f\colon X \to \mathbb{I}\}$ denotes the set of *interval specifications* on a finite set X. An interval specification $f \in \mathcal{D}_{\mathbb{I}}(X)$ describes a family of probability distributions on X that satisfy the specification i.e., $[\![f]\!] = \{\pi \in \mathcal{D}(X) \mid \forall x \in X.\pi(x) \in f(x)\}$.

Definition 3 (Interval Markov Reward Model). *An* interval Markov reward model *(IMRM) is a tuple* $\mathcal{I} = (S, s_0, \Pi, R, \ell)$ *where*

- S *is a finite nonempty set of states,*
- $s_0 \in S$ *is the initial state,*
- $\Pi\colon S \to \mathcal{D}_{\mathbb{I}}(S)$ *is the interval transition function,*
- $R\colon S \to \mathbb{N}_{>0}$ *is the state-reward function, and*
- $\ell\colon S \to 2^A$ *is the state-labeling function.*

Given an IMRM $\mathcal{I} = (S, s_0, \Pi, R, \ell)$ and state $s \in S$, $\Pi(s) = I_s$ is the interval-specification for state s, defining for each state $s' \in S$ a probability interval $I_s(s')$, within which s moves to s'. By abuse of notation we may refer to MRMs as particular cases of IMRMs having singleton intervals specifications. Hence, an IMRM $\mathcal{I}$ is a succinct specification for a family of MRMs where the transition function satisfies boundary conditions dictated by the interval transition function Π. Hereafter, we will assume that all IMRMs we will be working with have non-empty interval specifications, i.e., $[\![\Pi(s)]\!] \neq \emptyset$ for all $s \in S$. In literature this condition is known as (local) *consistency* [13]. The definition of paths, finite paths and accumulated weight are defined similarly as for MRMs.

Example 3. Consider the IMRM $\mathcal{I}_1 = (S, s_0, \Pi, R, \ell^{\mathcal{I}})$ depicted in Fig. 1a. For any state $s_i \in \{s_0, s_1, s_2, s_3, s_4\}$, the interval specification $\Pi(s_i)$ is depicted by edges connecting s_i to states in $\mathtt{succ}(s_i)$. These edges are labelled by the interval assigned by $\Pi(s_i)(s_j)$. Singleton intervals $[p,p]$ are simply represented by p.

In the literature [5,8,22,29], there have been proposed three different semantic interpretations of IMRMs, namely, the *once-and-for-all semantics*, the *interval Markov decision process semantics* (IMDP), and the *at-every-step semantics*. We now present the three distinct semantics for IMRMs and some basic results showing the relationship among the different semantics. To ease the presentation, we fix an MRM $\mathcal{M} = (T, t_0, \pi, \rho, \ell^{\mathcal{M}})$ and an IMRM $\mathcal{I} = (S, s_0, \Pi, R, \ell^{\mathcal{I}})$ and we will implicitly refer to their components in the remainder of this section.

The *once-and-for-all* semantics [5], also called the Uncertain Markov Chain semantics [29] is the simplest among the three semantics. It requires to choose for each state of the IMRM a probability distribution satisfying the corresponding interval specification.

Definition 4 (Once-and-for-all semantics). *An arbitrary MRM $\mathcal{M}$ satisfies the IMRM $\mathcal{I}$ w.r.t. the once-and-for-all semantics, written $\mathcal{M} \models_o \mathcal{I}$, if and only if $T \subseteq S$, $t_0 = s_0$, and for all $t \in T$, $\rho(t) = R(t)$, $\ell^{\mathcal{M}}(t) = \ell^{\mathcal{I}}(t)$, and $\pi(t) \in [\![\Pi(t)]\!]$.*

Example 4. Consider again the IMRM $\mathcal{I}_1$ in Fig. 1a. Figure 1b depicts an MRM $\mathcal{M}_o$ that satisfies $\mathcal{I}_1$ with the once-and-for-all semantics.

In contrast to the once-and-for-all semantics, in the *interval Markov decision process semantics* (IMDP semantics) [5,8,29], the choice of the transition probability distribution for a state $s \in S$ is performed each time a state is visited.

Definition 5 (IMDP semantics). *An MRM $\mathcal{M}$ satisfies the IMRM $\mathcal{I}$ w.r.t. the IMDP semantics, written $\mathcal{M} \models_d \mathcal{I}$, if and only if there exists a mapping $\tau\colon T \to S$ such that $\tau(t_0) = s_0$, and for all $t \in T$, $\ell^{\mathcal{M}}(t) = \ell^{\mathcal{I}}(\tau(t))$, $\rho(t) = R(\tau(t))$, and there exists $\delta_t \in [\![\Pi(\tau(t))]\!]$ such that for all $t' \in T$, $t' \in \mathtt{succ}(t)$ implies that $\pi(t)(t') = \delta_t(\tau(t'))$.*

As its name suggests, the IMDP semantics is reminiscent of the way one resolves nondeterminism in a Markov decision process (MDP) by means of a deterministic memory-dependent scheduler (*cf.* [4, Ch10]). With respect to similar semantic interpretations given for interval Markov chains [5,8,29], Definition 5 is more similar in spirit to that given in [5] for the fact that the MRM $\mathcal{M}$ needs to be finite.

Example 5. The MRM in Fig. 1c satisfies the IMRM $\mathcal{I}_1$ in Fig. 1a w.r.t. the IMDP semantics. To see this, consider the mapping $\tau(u_0) = s_0, \tau(u_1) = s_1, \tau(u_{21}) = \tau(u_{22}) = s_2, \tau(u_3) = s_3$ and $\tau(u_4) = s_4$. Note that u_{21} and u_{22} are two different implementations of the IMRM state s_2.

Remark 1. Notice that $\mathcal{M} \models_o \mathcal{I}$ implies $\mathcal{M} \models_d \mathcal{I}$ and the mapping $\tau\colon T \to S$ witnessing this fact is the identity function, i.e., $\tau(t) = t$ for all $t \in T$.

The last semantic interpretation for IMRMs is the so-called *at-every-step semantics*. Its definition is a simple extension of the original semantics given for interval Markov chains by Jonsson and Larsen [22]. Its main feature consists in generalizing the mapping $\tau\colon T \to S$ from the IMDP semantics to a relation $\mathcal{R} \subseteq T \times S$. This allows one to "aggregate" compatible states of the IMRM into a single state of the MRM implementation, as well as "redistributing" the successors of a state of the IMRM into multiple states.

Definition 6 (At-every-step semantics). *An MRM $\mathcal{M}$ satisfies the IMRM $\mathcal{I}$ w.r.t. the at-every-step semantics, written $\mathcal{M} \models_a \mathcal{I}$ if and only if there exists a relation $\mathcal{R} \subseteq T \times S$ such that $(t_0, s_0) \in \mathcal{R}$ and for all pairs $(t, s) \in \mathcal{R}$ we have that $\ell^{\mathcal{M}}(t) = \ell^{\mathcal{I}}(s)$, $\rho(t) = R(s)$, and there exists a correspondence function $\delta_{(t,s)}\colon T \to (S \to [0,1])$ such that*

1. *for all* $t' \in \mathtt{succ}(t)$, $\delta_{(t,s)}(t') \in \mathcal{D}(S)$.
2. *for all* $s' \in S$,

$$\left(\sum_{t' \in T} \pi(t)(t') \cdot \delta_{(t,s)}(t')(s')\right) \in \Pi(s)(s').$$

3. *for all* $(t', s') \in T \times S$, *if* $\delta_{(t,s)}(t')(s') > 0$ *then* $(t', s') \in \mathcal{R}$.

Example 6. The MRM $\mathcal{M}_a$ depicted in Fig. 1d is one possible at-every-step implementation of the IMRM $\mathcal{I}_1$ of Fig. 1a. This is witnessed by the relation

$$\mathcal{R} = \{(v_0, s_0), (v_1, s_1), (v_1, s_2), \\ (v_{31}, s_3), (v_{32}, s_3), (v_4, s_4)\}$$

and the following correspondence functions:

$$\delta_{(v_0,s_0)}(v_1)(s_1) = \delta_{(v_0,s_0)}(v_1)(s_2) = \tfrac{1}{2}\,, \\ \delta_{(v_1,s_1)}(v_{31})(s_3) = \delta_{(v_1,s_1)}(v_{32})(s_3) = 1.$$

Note that the state v_1 in $\mathcal{M}_a$ implements both s_1 and s_2, while the state s_3 is "redistributed" into v_{31} and v_{32}. The example illustrates that one is allowed to aggregate and split states under the at-every-step semantics.

As shown in Example 6, the at-every-step semantics allows one MRM state to implement multiple IMRM states by aggregation. Next, we show that for any MRM with such aggregated states, there exists an at-every-step implementation with no aggregated states, which is probabilistic bisimilar to the MRM with aggregated states. The result follows immediately from a similar result for IMCs as presented in [5, Proposition 5]. To formalize the result, we borrow the notion of *degree of satisfaction* from [5].

Definition 7. *Let* $n \in \mathbb{N}$. *The MRM* $\mathcal{M}$ *satisfies the IMRM* $\mathcal{I}$ *w.r.t. the at-every-step semantics with degree of satisfaction* n, *written* $\mathcal{M} \models_a^n \mathcal{I}$, *if there exists a relation* $\mathcal{R} \subseteq T \times S$ *witnessing* $\mathcal{M} \models_a \mathcal{I}$ *such that* $|\mathcal{R}(t)| \leq n$ *for all states* $t \in T$.

Note that if an MRM $\mathcal{M}$ satisfies IMRM $\mathcal{I}$ with degree 1, all correspondence functions $\delta_{(t,s)}$ are Dirac distributions i.e. $\delta_{(t,s)}(t')(s') > 0 \implies \delta_{(t,s)}(t')(s') = 1$.

The following Lemma states that for any at-every-step implementation $\mathcal{M}$ of the IMRM $\mathcal{I}$, there exists an at-every-step implementation $\mathcal{M}'$ of $\mathcal{I}$ with degree 1 that is probabilistic bisimilar to $\mathcal{M}$.

Lemma 1. *Let* $\mathcal{M} \models_a^n \mathcal{I}$ *for some* $n \in \mathbb{N}$. *Then, there exists an MRM* $\mathcal{M}'$ *such that* $\mathcal{M} \sim \mathcal{M}'$ *and* $\mathcal{M}' \models_a^1 \mathcal{I}$.

Remark 2. Note that $\mathcal{M} \models_d \mathcal{I}$ implies $\mathcal{M} \models_a^1 \mathcal{I}$, since the mapping $\tau \colon T \to S$ witnessing $\mathcal{M} \models_d \mathcal{I}$ induces a functional relation $\mathcal{R} = \{(t, \tau(t)) \mid t \in T\}$ which can be easily verified to be a witness for $\mathcal{M} \models_a^1 \mathcal{I}$.

The following result identifies the properties that a relation $\mathcal{R}$ witnessing $\mathcal{M} \models_a \mathcal{I}$ has when the MRM $\mathcal{M}$ satisfies also $\mathcal{M} \models_d \mathcal{I}$.

Proposition 1. *Let $\mathcal{R} \subseteq T \times S$ be a relation witnessing $\mathcal{M} \models_a^1 \mathcal{I}$. Then, $\mathcal{M} \models_d \mathcal{I}$ iff for all $(t, s) \in \mathcal{R}$ there exists no $s' \in \mathtt{succ}(s)$ such that $|\mathcal{R}^{-1}(s') \cap \mathtt{succ}(t)| > 1$.*

We are now ready to establish some basic relationship between the three semantics in terms of their expressivity. For any semantics $x \in \{o, d, a\}$, we denote by $[\![\mathcal{I}]\!]_x = \{\mathcal{M} \mid \mathcal{M} \models_x \mathcal{I}\}$ the family of MRMs that satisfy the IMRM $\mathcal{I}$ with respect to the semantic x.

The following result states that the three semantics presented in this section have different expressivity, with the at-every-step semantics being the most expressive semantics, followed by the IMDP semantics which in turn is more expressive than the once-and-for-all semantics.

Proposition 2. *For any IMRM $\mathcal{I}$, $[\![\mathcal{I}]\!]_o \subseteq [\![\mathcal{I}]\!]_d \subseteq [\![\mathcal{I}]\!]_a$ and for some IMRM $\mathcal{I}'$ these inclusions are strict, i.e., $[\![\mathcal{I}']\!]_o \subset [\![\mathcal{I}']\!]_d \subset [\![\mathcal{I}']\!]_a$.*

Proof. $[\![\mathcal{I}]\!]_o \subseteq [\![\mathcal{I}]\!]_d$ and $[\![\mathcal{I}]\!]_d \subseteq [\![\mathcal{I}]\!]_a$ follow for the arguments sketched respectively in Remarks 1 and 2. For the IMRM $\mathcal{I}_1$ in Fig. 1 in particular it holds that $\mathcal{M}_d \in [\![\mathcal{I}_1]\!]_d \setminus [\![\mathcal{I}_1]\!]_o$ and $\mathcal{M}_a \in [\![\mathcal{I}_1]\!]_a \setminus [\![\mathcal{I}_1]\!]_d$. □

5 Comparing Semantics Against rPCTL

In this section we investigate the IMRM semantics presented in Sect. 4 in the context of rPCTL model-checking. The rPCTL satisfiability relation naturally extends to IMRMs by requiring that an rPCTL formula is satisfied by some MRM implementation.

Definition 8. *We say that an IMRM $\mathcal{I}$ (existentially) satisfies the formula $\Phi \in \mathtt{rPCTL}$ with respect to the semantics $x \in \{o, d, a\}$, written $\mathcal{I} \models_x \Phi$, iff there exists $\mathcal{M} \in [\![\mathcal{I}]\!]_x$ such that $\mathcal{M} \models \Phi$.*

The above definition is implicitly given in terms of the initial state s_0 of $\mathcal{I}$, but can be generalized to arbitrary states $s \in S$, as $\mathcal{I}, s \models \Phi$ by replacing s_0 with s.

In the following, we compare the three different semantics with respect to different classes of rPCTL formulae. To this end introduce a notion of semantic equivalence.

Definition 9 (Semantic equivalence). *For a fragment of rPCTL, $\mathcal{L} \subseteq \mathtt{rPCTL}$ and two IMRM semantics $x, y \in \{o, d, a\}$, we say that the semantics x and y are* equivalent *w.r.t. $\mathcal{L}$ if for any IMRM $\mathcal{I}$ and state formula $\Phi \in \mathcal{L}$, $\mathcal{I} \models_x \Phi \iff \mathcal{I} \models_y \Phi$.*

The next result states that the at-every-step semantics is not equivalent to the IMDP semantics w.r.t. the full logic.

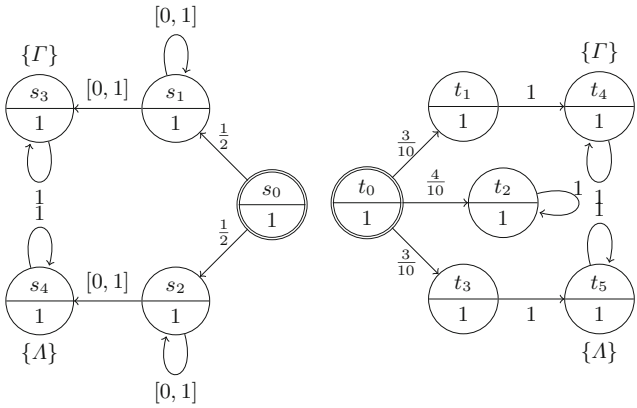

(a) IMRM $\mathcal{I}_2$ (b) MRM $\mathcal{M}_2 \in [\![\mathcal{I}_2]\!]_a$

Fig. 2. IMRM $\mathcal{I}_2$ with at-every-step MRM implementation $\mathcal{M}_2$

Proposition 3. *The at-every-step semantics is not semantically equivalent to the IMDP semantics with respect to* `rPCTL`.

Proof. Consider the IMRM $\mathcal{I}_2$ and MRM $\mathcal{M}_2$ depicted in Fig. 2. One can verify that $\mathcal{M}_2 \models_a \mathcal{I}_2$. Let Φ be the following rPCTL formula

$$\Phi = \mathcal{P}_{>0}(\mathsf{X}_{\leq 1}\Phi_1) \wedge \mathcal{P}_{>0}(\mathsf{X}_{\leq 1}\Phi_2) \wedge \mathcal{P}_{>0}(\mathsf{X}_{\leq 1}\Phi_3),$$

where

$$\Phi_1 = \mathcal{P}_{\geq 1}(\mathsf{X}_{\leq 1}(\neg\Gamma \wedge \neg\Lambda)),\ \Phi_2 = \mathcal{P}_{\geq 1}(\mathsf{X}_{\leq 1}\Gamma) \text{ and } \Phi_3 = \mathcal{P}_{\geq 1}(\mathsf{X}_{\leq 1}\Lambda).$$

Clearly $\mathcal{M}_2 \models \Phi$ as the three outgoing transitions serve to satisfy each of the sub-formulae Φ_i ($i \in \{1, 2, 3\}$).

Consider an MRM $\mathcal{M}' \in [\![\mathcal{I}_2]\!]_d$. By IMDP semantics, $\mathcal{M}'$ must have an initial state with exactly two successors, say t'_1 and t'_2. Therefore there exists $i \in \{1, 2, 3\}$ such that $\mathcal{M}', t_j \not\models \Phi_i$ for any $j = 1, 2$ as no single successor can satisfy each Φ_i simultaneously. Hence, $\mathcal{I}_2 \not\models_d \Phi$. □

The above result is analogous to [5, Section 4.1], where it was proven that for internal Markov chains the at-every-step semantics and the IMDP semantics are not equivalent with respect to PCTL.

Reachability Queries. In the rest of the section we focus our attention on a semantic comparison relative to reachability queries, namely, formulae of the form $\mathcal{P}_{\bowtie\lambda}(\Diamond_{\trianglelefteq k}\Gamma)$, for arbitrary $\Gamma \in AP$, $k \in \mathbb{N}_{>0}$, $\lambda \in [0, 1]$, $\bowtie \in \{<, \leq, \geq, >\}$, and $\trianglelefteq \in \{\leq, \geq\}$. We denote by $\mathcal{L}_{\texttt{reach}}$ the set of reachability queries and we write $\mathcal{L}^{\leq}_{\texttt{reach}}$ (resp. $\mathcal{L}^{\geq}_{\texttt{reach}}$) when we fix $\trianglelefteq = \leq$ (resp. $\trianglelefteq = \geq$).

Reachability properties are one of the fundamental questions for the quantitative analysis of systems. The atomic proposition Γ may represent a set of

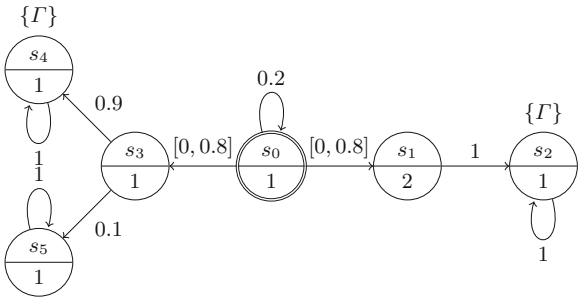

Fig. 3. IMRM $\mathcal{I}_3$

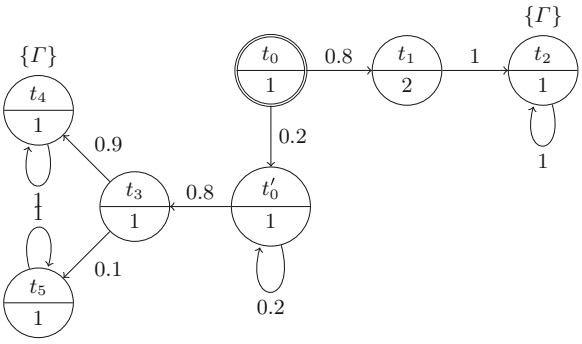

Fig. 4. MRM $\mathcal{M}_3$ such that $\mathcal{M}_3 \models \mathcal{P}_{>0.8}(\Diamond_{\leq 3}\Gamma)$

certain *bad* states which should be unlikely to be visited, or dually, a set of *good* states which should rather be visited with high probability. In the context of interval Markov chains, Bart et al. [5] have shown that the three semantic interpretations are equivalent with respect to reachability queries.

In contrast, the ability to express bounds on the reward accumulated until reaching a some goal state makes the IMPD semantics and the at-every-step semantics, more expressive than the once-and-for-all semantics relative to reachability queries.

Proposition 4. *The once-and-for-all semantics is not equivalent to the IMDP semantics w.r.t.* $\mathcal{L}_{\texttt{reach}}$.

Proof. Consider the IMRM $\mathcal{I}_3$ in Fig. 3 and the formula $\mathcal{P}_{>0.8}(\Diamond_{\leq 3}\Gamma)$. Figure 4 shows $\mathcal{M}_3 \in [\![\mathcal{I}]\!]_d$ witnessed by the mapping $\tau(t_0) = \tau(t'_0) = s_0$ and $\tau(t_i) = s_i$ for $1 \leq i \leq 5$. Clearly $\mathcal{M}_3 \models \mathcal{P}_{>0.8}(\Diamond_{\leq 3}\Gamma)$.

Figure 5 shows the once-and-for-all MRM implementation of $\mathcal{I}_3$, $\mathcal{M}_4$, that maximizes the probability of reaching Γ without exceeding the weight budget of 3. One can see that $\mathcal{M}_4 \not\models \mathcal{P}_{>0.8}(\Diamond_{\leq 3}\Gamma)$. □

It remains to compare the at-every-step semantics and IMDP semantics w.r.t. reachability queries. To this end we first present two technical lemmas.

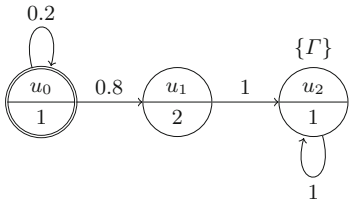

Fig. 5. MRM $\mathcal{M}_4$ such that $\mathcal{M}_4 \not\models \mathcal{P}_{>0.8}(\Diamond_{\leq 3}\Gamma)$

Lemma 2. *Let $\mathcal{I}$ be an IMRM, $\mathcal{M} \in [\![\mathcal{I}]\!]_a$, $\Gamma \in AP$, and $k \in \mathbb{N}$. Then, there exist $\mathcal{M}_{\leq}, \mathcal{M}_{\geq} \in [\![\mathcal{I}]\!]_d$ such that $\mathbb{P}^{\mathcal{M}_{\leq}}(\Diamond_{\leq k}\Gamma) \leq \mathbb{P}^{\mathcal{M}}(\Diamond_{\leq k}\Gamma) \leq \mathbb{P}^{\mathcal{M}_{\geq}}(\Diamond_{\leq k}\Gamma)$.*

Proof (sketch). By Lemma 1 and Theorem 1 we can assume w.l.o.g. that $\mathcal{M} \models_a^1 \mathcal{I}$. To construct $\mathcal{M}_{\leq}$ and $\mathcal{M}_{\geq}$ we proceed in two steps. We present the construction of $\mathcal{M}_{\leq}$ and then explain how to adapt it for $\mathcal{M}_{\geq}$.

(Step 1) We build an MRM $\mathcal{M}'$ from $\mathcal{M}$ by unfolding its structure. The unfolding of each path terminates when its accumulated weight exceeds k or when a state satisfying Γ is reached. Then, the last state of each unfolded path, say t, is replaced with an arbitrary once-and-for-all model of $\mathcal{I}$ with initial state t.
Note that $\mathbb{P}^{\mathcal{M}}(\Diamond_{\leq k}\Gamma) = \mathbb{P}^{\mathcal{M}'}(\Diamond_{\leq k}\Gamma)$ since the probability value is obtained as the sum over all the cylinders obtained from paths constructed in the unfolding. Moreover, $\mathcal{M}' \models_a^1 \mathcal{I}$ because $\mathcal{M} \models_a^1 \mathcal{I}$ and the unfolding does not introduce any aggregation.
(Step 2) From $\mathcal{M}'$ we construct $\mathcal{M}_{\leq}$. Let $\mathcal{R}$ be the relation witnessing $\mathcal{M}' \models_a^1 \mathcal{I}$. If $\mathcal{R}$ satisfies the conditions of Proposition 1 we choose $\mathcal{M}' = \mathcal{M}_{\leq}$. Otherwise, for each state t of $\mathcal{M}$ such that $(t, s) \in \mathcal{R}$ and $|\mathcal{R}^{-1}(s') \cap \mathtt{succ}(t)| > 1$ for some $s' \in \mathtt{succ}(s)$ we proceed as follows. Let $t' \in \mathcal{R}^{-1}(s') \cap \mathtt{succ}(t)$ be the successor of t that minimizes the probability of reaching Γ within the reward bound up to t', i.e., $\mathbb{P}_{t'}^{\mathcal{M}'}(\Diamond_{\leq k'}\Gamma)$ where $k' = k - \mathcal{W}(\sigma)(|\sigma|)$ and σ is the finite path from the initial state of $\mathcal{M}'$ to t'. Then we redirect all the probability mass that was from t to $\mathcal{R}^{-1}(s') \cap \mathtt{succ}(t)$ to the single state t', "disconnecting" the set of states $(\mathcal{R}^{-1}(s') \cap \mathtt{succ}(t)) \setminus \{t'\}$ from t.
Let $\mathcal{M}_{\leq}$ be the MRM obtained from the above procedure. Note that $\mathcal{M}_{\leq}$ satisfies the conditions of Proposition 1 and $\mathbb{P}^{\mathcal{M}_{\leq}}(\Diamond_{\leq k}\Gamma) \leq \mathbb{P}^{\mathcal{M}'}(\Diamond_{\leq k}\Gamma)$.

As for the construction of $\mathcal{M}_{\geq}$, (Step 1) is done in the same way while in (Step 2) t' is chosen as the one that maximizes $\mathbb{P}_{t'}^{\mathcal{M}'}(\Diamond_{\leq k'}\Gamma)$. □

Lemma 3. *Let $\mathcal{I}$ be an IMRM, $\mathcal{M} \in [\![\mathcal{I}]\!]_a$, $\Gamma \in AP$, and $k \in \mathbb{N}$. Then, there exist $\mathcal{M}_{\leq}, \mathcal{M}_{\geq} \in [\![\mathcal{I}]\!]_d$ such that $\mathbb{P}^{\mathcal{M}_{\leq}}(\Diamond_{\geq k}\Gamma) \leq \mathbb{P}^{\mathcal{M}}(\Diamond_{\geq k}\Gamma) \leq \mathbb{P}^{\mathcal{M}_{\geq}}(\Diamond_{\leq k}\Gamma)$.*

Proof (sketch). The proof proceeds in two steps analogously as for Lemma 2.

By Lemma 1 and Theorem 1 we can assume w.l.o.g. that $\mathcal{M} \models_a^1 \mathcal{I}$. We present the construction of $\mathcal{M}_{\leq}$ and then explain how to adapt it for $\mathcal{M}_{\geq}$.

(Step 1) We build an MRM $\mathcal{M}'$ from $\mathcal{M}$ by unfolding its structure. The unfolding of each path terminates as soon as its accumulated weight exceeded k. Then, the last state of each unfolded path, say t, is replaced with a once-and-for-all model $\mathcal{M}''$ of $\mathcal{I}$ with initial state t such that $\mathbb{P}_t^{\mathcal{M}''}(\Diamond \Gamma) \leq \mathbb{P}_t^{\mathcal{M}}(\Diamond \Gamma)$. The existence of $\mathcal{M}''$ is guaranteed by [5, Lemma 4].
Note that $\mathbb{P}^{\mathcal{M}'}(\Diamond_{\geq k} \Gamma) \leq \mathbb{P}^{\mathcal{M}}(\Diamond_{\geq k} \Gamma)$ since the probability value is obtained as the sum over all the cylinders obtained from paths constructed in the unfolding. Moreover, $\mathcal{M}' \models_a^1 \mathcal{I}$ because $\mathcal{M} \models_a^1 \mathcal{I}$ and the unfolding does not introduce any aggregation.
(Step 2) From $\mathcal{M}'$ we construct $\mathcal{M}_{\leq}$ following the same procedure used for (Step 2) in the proof of Lemma 2.

As for the construction of $\mathcal{M}_{\geq}$, in (Step 1) we need to choose $\mathcal{M}''$ such that $\mathbb{P}_t^{\mathcal{M}''}(\Diamond \Gamma) \geq \mathbb{P}_t^{\mathcal{M}}(\Diamond \Gamma)$ and (Step 2) is modified as done for Lemma 2. □

Theorem 2. *The IMDP semantics and the at-every-step semantics are equivalent w.r.t.* $\mathcal{L}_{\texttt{reach}}$.

Proof. Let $\mathcal{I}$ be an IMRM, $\Phi = \mathcal{P}_{\bowtie \lambda}(\Diamond_{\trianglelefteq k} \Gamma) \in \mathcal{L}_{\texttt{reach}}$ and $\mathcal{M} \in [\![\mathcal{I}]\!]_a$ such that $\mathcal{M} \models \Phi$. We proceed by cases.

If $\trianglelefteq = \leq$, we consider two sub-cases. If $\bowtie \in \{<, \leq\}$ then by Lemma 2 there exists $\mathcal{M}_{\leq} \in [\![\mathcal{I}]\!]_d$ such that $\mathbb{P}^{\mathcal{M}_{\leq}}(\Diamond_{\leq k} \Gamma) \leq \mathbb{P}^{\mathcal{M}}(\Diamond_{\leq k} \Gamma)$. Therefore $\mathcal{I} \models_d \Phi$. If $\bowtie \in \{\geq, >\}$ then by Lemma 2 there exists $\mathcal{M}_{\geq} \in [\![\mathcal{I}]\!]_d$ such that $\mathbb{P}^{\mathcal{M}_{\geq}}(\Diamond_{\leq k} \Gamma) \geq \mathbb{P}^{\mathcal{M}}(\Diamond_{\leq k} \Gamma)$. Hence $\mathcal{I} \models_d \Phi$.

If $\trianglelefteq = \geq$ we use the same arguments as before by using Lemma 3 in place of Lemma 2. □

6 Model-Checking Algorithms

In this section we turn our attention to model-checking different fragments of rPCTL. By the results of the previous section, each IMRM semantics requires its own treatment for the full logic rPCTL. For the important fragment of reachability queries in $\mathcal{L}_{\texttt{reach}}$, we need two algorithms, one for the once-and-for-all semantics and one for the IMDP semantics. In the following, we restrict ourselves to formulae with only upper bounds on path-formulae and similar to $\mathcal{L}^{\leq}_{\texttt{reach}}$, we denote by $\texttt{rPCTL}^{\leq}$ the set of all rPCTL formulae with only upper bounds on the path formulae.

For the once-and-for-all semantics we reduce the model-checking problem w.r.t $\texttt{rPCTL}^{\leq}$ to the (existential) model-checking problem for parametric Markov reward models (PMRMs) [1] with interval constraints on the parameters. Efficient procedures for model-checking PMRMs against various logics have received a lot of attention in recent years and are now supported by modern tools such as PRISM [24], PARAM [16] and PROPHESY [10]. For the IMDP semantics we exploit the fact that all rewards are strictly positive to devise a reduction to the model-checking problem for the once-and-for-all semantics. For the at-every-step semantics we leave the model-checking problem for fragments containing $\mathcal{L}^{\leq}_{\texttt{reach}}$ open. We proceed by treating each semantics in turn.

6.1 Once-and-for-all Semantics

For the fragment $\texttt{rPCTL}^{\leq}$ we present a reduction to the (existential) model-checking problem for PMRMs with interval constraints on the parameters. We first recall the definition of PRMMs and then present the reduction. PMRMs extend MRMs by allowing the transition probabilities to take values in a finite nonempty set P of *parameters*. Thus, for any finite nonempty set X, the function $\mu_P \colon X \to [0,1] \cup P$ is a parametric distribution. The set $\mathcal{D}_P(X)$ is then the set of all parametric distributions.

Definition 10 (Parametric MRM). *A* parametric *Markov Reward Model (PMRM) is defined as a tuple* $\mathcal{M}_P = (S, s_0, \rho, \pi_P, \ell^{\mathcal{M}})$ *where* S, s_0, ρ *and* $\ell^{\mathcal{M}}$ *are defined as for MRMs and for each* $s \in S$, $\pi_P \in \mathcal{D}_P(S)$ *is the* parametric *probability transition function.*

A given PMRM gives rise to a set of MRMs by interpreting the parameters as rational values and making sure that the resulting distribution are probability distributions (i.e. sum up to 1). Formally, a valuation function $\kappa \colon \mathbb{Q}_{>0} \cup P \to [0,1]$ is a function such that for all $r \in \mathbb{Q}_{>0}$, $\kappa(r) = r$, for all $p \in P$, $\kappa(p) > 0$ and for all states $s \in S$, $\sum_{s' \in S} \kappa(\pi_P(s)(s'))) = 1$. We abuse notation and for any PMRM $\mathcal{M}_P$ write $\kappa(\mathcal{M}_P)$ for the MRM induced by κ.

Existential Model-Checking. We consider the following decision problem for PMRMs: given a PMRM $\mathcal{M}_P$ and formula $\Phi \in \texttt{rPCTL}^{\leq}$, does there exist a valuation function κ such that $\kappa(\mathcal{M}_P) \models \Phi$?

The problem is extended with interval-constraints on the parameters as follows: for all $(s, s') \in S \times S$ let $I_{s,s'} = [l_{s,s'}, u_{s,s'}] \in \mathbb{I}$ be some interval. The parameter valuation function κ must then also satisfy the following constraints:

$$\bigwedge_{(s,s') \in S \times S} \kappa(\pi_P(s)(s')) \in I_{s,s'}.$$

The Reduction. Let $\mathcal{I} = (S, s_0, \Pi, R, \ell^{\mathcal{I}})$ be an IMRM and $\Phi \in \texttt{rPCTL}^{\leq}$ an arbitrary formula. We now construct a PMRM $\mathcal{M}_P$ and a set of interval constraints such that if there exists a valuation function κ where $\kappa(\mathcal{M}_P) \models \Phi$, while κ satisfies the given interval constraints, then $\mathcal{I} \models_o \Phi$.

Let $\mathcal{M}_P$ be the PMRM identical to $\mathcal{I}$ except that each interval $\Pi(s)(s')$ is replaced by a parameter $p_{s,s'}$. For each $p_{s,s'}$, the interval constraint that κ must satisfy, is given by $\Pi(s)(s')$ i.e. any parameter valuation function κ must satisfy the following interval constraints:

$$\bigwedge_{(s,s') \in S \times S} \kappa(p_{s,s'}) \in \Pi(s)(s').$$

Assume that there exists a valuation function κ such that $\kappa(\mathcal{M}_P) \models \Phi$ in addition to satisfying the above interval constraints. Without loss of generality, we assume that all states in $\kappa(\mathcal{M}_P)$ are reachable as all states of any MRM have

to be reachable. If that is not the case, one can simply remove all unreachable states as they do not influence satisfiability. By construction, it is clear that $\kappa(\mathcal{M}_P) \in [\![\mathcal{I}]\!]_o$ as κ induces a probability distribution for each state s that satisfies the interval-specifications $\Pi(s)$ given by $\mathcal{I}$, while preserving rewards and labels of each state.

Interpreting $\texttt{rPCTL}^{\leq}$ *on PMRMs.* In literature, papers on PCTL model-checking for PMRMs only consider step-bounded or unbounded until-formulae, in contrast to the reward-bounded until formulae in $\texttt{rPCTL}^{\leq}$. This is not a restriction since any PMRM that contains a state s with a reward greater than 1 can be replaced by a sequence of states with reward 1. Hence, any upper bound on the formula can be interpreted as a step-bound in this (larger) model. In the same way, it is possible to "unroll" the model to reduce (step)-bounded reachability to unbounded reachability [23, Remark 4]. Thus, any technique for model-checking PCTL where the until is step-bounded or unbounded on PMRMs can be used for $\texttt{rPCTL}^{\leq}$ [3,10,16,23,24]. In the case of unbounded until, the model-checking problem for PMRMs is in PSPACE [19].

6.2 IMDP Semantics

Our approach for verifying properties with the IMDP semantics is based on the fact that the IMDP semantics is a simple extension of the once-and-for-all semantics, where one is allowed to choose a different probability distribution each time a state is visited. Recall that every reward in the model is strictly positive and we have concrete upper bounds on all path-formulae. Hence, even if one is allowed to choose a different distribution each time a state is visited, for the purpose of verifying Φ, we can bound the number of times a different probability distribution needs to be chosen for any IMDP implementation that satisfies Φ. Hence, one can do a bounded unfolding of the IMRM that preserves interval specifications, to encode all possible implementations that may satisfy Φ. The unfolding itself is an IMRM, where the set of states is the set of all non-empty finite paths, S^+, bounded by a given depth k. Interval-preservation is ensured by letting the transitioning between any two such states be defined by the transitioning between their two last states in the original IMRM.

Definition 11 (IMRM k-unfolding). *For any IMRM* $\mathcal{I} = (S, s_0, \Pi, R, \ell^{\mathcal{I}})$ *and* $k \in \mathbb{N}$, *let* $\mathcal{I}{\downarrow_k} = (S_k, s_0, \Pi_k, R_k, \ell^{\mathcal{I}\downarrow_k})$ *be the* interval specification preserving *k-unfolding of* $\mathcal{I}$, *defined as follows*[2]*:*

- $S_k = \{\sigma \in S^+ \mid \mathcal{W}(\sigma)(|\sigma|) \leq k, \quad \forall_{0 \leq i < |\sigma|}. \Pi(\sigma[i], \sigma[i+1]) \neq [0,0]\}$.
- *For all* $\sigma \in S_k \cup \{\epsilon\}$[3] *and* $s, s_1 \in S$, $\Pi_k(\sigma s, \sigma s s_1) = \Pi(s, s_1)$.
- *For any path* $\sigma = s_1, \ldots, s_n$ *in* S_k, $R_k(\sigma) = R(s_n)$ *and* $\ell^{\mathcal{I}\downarrow_k}(\sigma) = \ell^{\mathcal{I}}(s_n)$.

[2] Technically, self-loops must be added for states that represent maximal paths w.r.t k in order for the unfolding to be a proper IMRM.

[3] Where ϵ is the empty string.

As any MRM $\mathcal{M}$ can be seen as an IMRM with singleton intervals, we abuse notation and write $\mathcal{M}{\downarrow_k}$ for the k-unfolding of the MRM $\mathcal{M}$.

The following two lemmas prove two key properties of our unfolding definition. The first lemma states that for an IMRM $\mathcal{I}$ with initial state s_0, if any successor $s_0 s'$ of s_0 in the k-unfolding of $\mathcal{I}$ satisfies a given formula Φ, then this can be verified by changing the initial state to s' and performing a $(k - R(s_0))$-unfolding of $\mathcal{I}$ where $R(s_0)$ is the reward assigned by $\mathcal{I}$ to s_0. The second lemma states that whenever an MRM is an IMDP implementation of an IMRM $\mathcal{I}$ then the k-unfolding of $\mathcal{M}$ is an once-and-for-all implementation of the k-unfolding of $\mathcal{I}$. This implies that if from any formula $\Phi \in \texttt{rPCTL}^{\leq}$ we can define a $k \in \mathbb{N}$ such that the k-unfolding of $\mathcal{I}$ includes all the paths needed for verifying Φ, we can reduce the model-checking problem using the IMDP semantics to model-checking using the once-and-for-all semantics on the k-unfolding of $\mathcal{I}$, $\mathcal{I}{\downarrow_k}$.

Lemma 4. *For any two IMRMs defined as* $\mathcal{I}^{s_0} = (S, s_0, \Pi, R, \ell^{\mathcal{I}})$ *and* $\mathcal{I}^{s'} = (S, s', \Pi, R, \ell^{\mathcal{I}})$ *with* $s' \in \texttt{succ}(s_0)$, $k \geq R(s_0)$, $\Phi \in \texttt{rPCTL}^{\leq}$ *and semantics* $x \in \{o, d, a\}$, *it holds that*

$$\mathcal{I}^{s_0}{\downarrow_k}, s_0 s' \models_x \Phi \implies \mathcal{I}^{s'}{\downarrow_{k-R(s_0)}}, s' \models_x \Phi.$$

Proof. The lemma follows easily from the definition of unfolding and rPCTL semantics. The condition $k \geq R(s_0)$ ensures that $s_0 s'$ is a state in $\mathcal{I}^{s_0}$.

Lemma 5. *For any IMRM* $\mathcal{I}$, *MRM* $\mathcal{M}$ *and* $k \in \mathbb{N}$, *if* $\mathcal{M} \in [\![\mathcal{I}]\!]_d$ *then* $\mathcal{M}{\downarrow_k} \in [\![\mathcal{I}{\downarrow_k}]\!]_o$.

Remark 3. Strictly speaking, Lemma 5 only holds up to isomorphism as $\mathcal{M}$ by the IMDP semantics may contain states not in $\mathcal{I}$. In this case, the states of $\mathcal{M}{\downarrow_k}$ is not a subset of the states of $\mathcal{I}{\downarrow_k}$ as required by the once-and-for-all semantics.

For any formula $\Phi \in \texttt{rPCTL}^{\leq}$ we define the reward-depth denoted by $K(\Phi) \in \mathbb{N}$, on the structure of Φ. For a probabilistic reward-bounded reachability objective of the form $\Phi = \mathcal{P}_{\bowtie\lambda}(\Diamond_{\leq k} \Gamma)$, $K(\Phi) = k$ implies that only paths with an accumulated reward of at most k is of interest. Hence, a k-unfolding of $\mathcal{I}$ is sufficient for the purpose of verifying Φ.

Definition 12 (Reward-depth). *For every property* $\Phi \in \texttt{rPCTL}^{\leq}$, *the* reward-depth, $K(\Phi) \in \mathbb{N}$ *is defined inductively on the structure of* Φ:

$$\begin{aligned}
&K(\mathit{true}) = 0\\
&K(a) = 0\\
&K(\neg\Phi) = K(\Phi)\\
&K(\Phi_1 \wedge \Phi_2) = \max\{K(\Phi_1), K(\Phi_2)\}\\
&K(\mathcal{P}_{\bowtie\lambda}(\mathsf{X}_{\leq k}\Phi)) = k + K(\Phi)\\
&K(\mathcal{P}_{\bowtie\lambda}(\Phi_1 \mathsf{U}_{\leq k}\Phi_2)) = k + \max\{K(\Phi_1), K(\Phi_2)\}
\end{aligned}$$

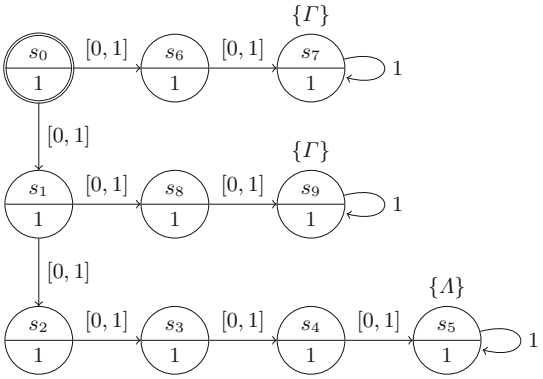

Fig. 6. IMRM $\mathcal{I}_4$

Example 7. Consider the IMRM $\mathcal{I}_3$ in Fig. 3 and formula $\Phi = \mathcal{P}_{>0.8}(\Diamond_{\leq 3}\Gamma)$. By definition, $K(\Phi) = 3 + \max\{K(\text{true}), K(\Gamma)\} = 3$. Notice that $\mathcal{I}_3 \models_d \Phi$ if and only if $k \geq 3$, with the witnessing implementation being the MRM in Fig. 4. Consider now the IMRM $\mathcal{I}_4$ in Fig. 6 and the property $\Phi' = \mathcal{P}_{\geq\lambda_1}(\Phi_1 \, U_{\leq 2} \, \Phi_2)$ where $\Phi_1 = \mathcal{P}_{\geq\lambda_2}(\Diamond_{\leq 2}\, \Gamma)$, $\Phi_2 = \mathcal{P}_{\geq\lambda_3}(\Diamond_{\leq 3}\, \Lambda)$ and $\lambda_1, \lambda_2, \lambda_3 \in [0, 1]$. For any semantics $x \in \{o, d, a\}$ it is clear that $\mathcal{I}_4{\downarrow_k} \not\models_x \Phi'$ if $k < 5$, irrespective of the concrete values for λ_1, λ_2 and λ_3, as the path $s_0, s_1, s_2, s_3, s_4, s_5$ must be preserved. By definition, $K(\Phi') = 5$ i.e. if one performs an unfolding of $\mathcal{I}$ with a reward-depth less than $K(\Phi')$, one cannot hope to find any implementation satisfying any concrete instantiation of Φ'.

As indicated by Example 7, $K(\Phi)$ is the reward-depth required for the verification of Φ. Hence, unfolding to a reward-depth greater than $K(\Phi)$ should not influence the satisfaction of Φ. The following lemma proves this monotonicity property.

Lemma 6 (Monotonicity). *For any MRM $\mathcal{M}$, formula $\Phi \in \texttt{rPCTL}^{\leq}$, $k \geq K(\Phi)$ and $\varepsilon > 0$,*

$$\mathcal{M}{\downarrow_k} \models \Phi \implies \mathcal{M}{\downarrow_{k+\varepsilon}} \models \Phi.$$

The next lemma states that if an MRM $\mathcal{M}$ satisfies Φ, then the $K(\Phi)$-unfolded model $\mathcal{M}{\downarrow_{K(\Phi)}}$ also satisfies Φ i.e. unfolding to a reward-depth of at least $K(\Phi)$ is sufficient to verify Φ.

Lemma 7. *For any MRM $\mathcal{M} \in [\![I]\!]_d$ and formula $\Phi \in \texttt{rPCTL}^{\leq}$,*

$$\mathcal{M} \models \Phi \implies \mathcal{M}{\downarrow_{K(\Phi)}} \models \Phi.$$

We now present the main theorem of this section, stating that rPCTL model-checking for IMRMs with the IMDP semantics can be reduced to model-checking using the once-and-for-all semantics on an IMRM constructed by unfolding to the reward-depth required by the given formula of interest.

Theorem 3. *For IMRM $\mathcal{I}$ and formula $\Phi \in \texttt{rPCTL}^{\leq}$,*

$$\mathcal{I} \models_d \Phi \implies \mathcal{I}{\downarrow_{K(\Phi)}} \models_o \Phi.$$

Proof. We assume $\mathcal{I} \models_d \Phi$, hence $\exists \mathcal{M} \in [\![I]\!]_d . \mathcal{M} \models \Phi$. By Lemma 7, $\mathcal{M} \models \Phi \implies \mathcal{M}{\downarrow_{K(\Phi)}} \models \Phi$ and by Lemma 5 we get $\mathcal{M} \in [\![I]\!]_d \implies \mathcal{M}{\downarrow_{K(\Phi)}} \in [\![I{\downarrow_{K(\Phi)}}]\!]_o$. Hence, $\mathcal{I}{\downarrow_{K(\Phi)}} \models_o \Phi$ as required. □

Complexity. For any IMRM $\mathcal{I} = (S, s_0, \Pi, R, \ell^{\mathcal{I}})$ let $R_{\min} = \min_{s \in S} R(s)$ be the smallest reward present in $\mathcal{I}$. The unfolded model, $\mathcal{I}{\downarrow_{K(\Phi)}}$ is then a model of size $\mathcal{O}\left(|S|^{\lfloor \frac{K(\Phi)}{R_{\min}} \rfloor}\right)$, as each state of $\mathcal{I}{\downarrow_{K(\Phi)}}$ is a leaf of the underlying $K(\Phi)$-bounded unfolding of $\mathcal{I}$ which is a tree with branching factor $\mathcal{O}(|S|)$ and height $\mathcal{O}\left(\lfloor \frac{K(\Phi)}{R_{\min}} \rfloor\right)$.

Remark 4. By Theorem 2, for any $\Phi \in \mathcal{L}^{\leq}_{\texttt{reach}}$ the approach presented in Sect. 6.2 is valid also for model checking Φ w.r.t. the at-every-step semantics.

7 Conclusion and Future Work

We investigated model-checking questions relative to IMRMs specifications interpreted according to three semantics: once-and-for-all, interval Markov decision process, and at-every-step. This work builds on the results of [5] on interval Markov chains by introducing an additional ingredient: rewards. We showed that by introducing rewards the one-at-for-all semantics is no longer expressive enough to answer (existential) reachability queries with respect to the other two semantics. Nevertheless, the IMDP semantics and the at-every-step semantics are still logically equivalent with respect to the reward-bounded reachability fragment of rPCTL.

We then presented how to algorithmically solve the model-checking problem for IMRMs by proposing different reductions to the model-checking problem for parametric Markov reward models (PMRMs). First, we presented a reduction to the model-checking problem of PMRMs for model checking IMRMs interpreted over the once-and-for-all semantics. Then, for the IMDP semantics, we presented a reduction to the model-checking problem for the once-and-for-all semantics, via a finite unfolding of the model. Crucial for our reduction is that the state rewards are positive. Notably, this reduction can also be used also to answer reward-bounded reachability queries for IMRMs interpreted according to the at-every-step semantics.

As future work, we plan to further investigate the model-checking problem with respect to the at-every-step semantic interpretations of IMRMs.

References

1. Alur, R., Henzinger, T.A., Vardi, M.Y.: Parametric real-time reasoning. In: Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing, San Diego, CA, USA, 16–18 May 1993, pp. 592–601 (1993). https://doi.org/10.1145/167088.167242

2. Andersen, H.R., Stirling, C., Winskel, G.: A compositional proof system for the modal mu-calculus. In: Proceedings of the Ninth Annual Symposium on Logic in Computer Science (LICS '94), Paris, France, 4–7 July 1994, pp. 144–153 (1994). https://doi.org/10.1109/LICS.1994.316076
3. Bacci, G., Hansen, M., Larsen, K.G.: Model checking constrained Markov reward models with uncertainties. In: Parker, D., Wolf, V. (eds.) QEST 2019. LNCS, vol. 11785, pp. 37–51. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-30281-8_3
4. Baier, C., Katoen, J.: Principles of Model Checking. MIT Press, Cambridge (2008)
5. Bart, A., Delahaye, B., Fournier, P., Lime, D., Monfroy, E., Truchet, C.: Reachability in parametric interval Markov chains using constraints. Theor. Comput. Sci. **747**, 48–74 (2018). https://doi.org/10.1016/j.tcs.2018.06.016
6. Benedikt, M., Lenhardt, R., Worrell, J.: LTL model checking of interval Markov chains. In: Piterman, N., Smolka, S.A. (eds.) TACAS 2013. LNCS, vol. 7795, pp. 32–46. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-36742-7_3
7. Benveniste, A., et al.: Contracts for system design. Found. Trends Electron. Design Autom. **12**(2–3), 124–400 (2018). https://doi.org/10.1561/1000000053
8. Chatterjee, K., Sen, K., Henzinger, T.A.: Model-checking w-regular properties of interval Markov chains. In: Amadio, R. (ed.) FoSSaCS 2008. LNCS, vol. 4962, pp. 302–317. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78499-9_22
9. Chen, T., Han, T., Kwiatkowska, M.Z.: On the complexity of model checking interval-valued discrete time Markov chains. Inf. Process. Lett. **113**(7), 210–216 (2013). https://doi.org/10.1016/j.ipl.2013.01.004
10. Dehnert, C.: PROPhESY: A PRObabilistic ParamEter SYnthesis tool. In: Kroening, D., Păsăreanu, C.S. (eds.) CAV 2015. LNCS, vol. 9206, pp. 214–231. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-21690-4_13
11. Delahaye, B.: Consistency for parametric interval Markov chains. In: 2nd International Workshop on Synthesis of Complex Parameters, SynCoP 2015, London, United Kingdom, 11 April 2015, pp. 17–32 (2015). https://doi.org/10.4230/OASIcs.SynCoP.2015.17
12. Delahaye, B., Larsen, K.G., Legay, A., Pedersen, M.L., Wąsowski, A.: Decision problems for interval Markov chains. In: Dediu, A.-H., Inenaga, S., Martín-Vide, C. (eds.) LATA 2011. LNCS, vol. 6638, pp. 274–285. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-21254-3_21
13. Delahaye, B., Larsen, K.G., Legay, A., Pedersen, M.L., Wasowski, A.: Consistency and refinement for interval Markov chains. J. Log. Algebr. Program. **81**(3), 209–226 (2012). https://doi.org/10.1016/j.jlap.2011.10.003
14. Delahaye, B., Lime, D., Petrucci, L.: Parameter synthesis for parametric interval Markov chains. In: Jobstmann, B., Leino, K.R.M. (eds.) VMCAI 2016. LNCS, vol. 9583, pp. 372–390. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49122-5_18
15. van Glabbeek, R.J., Smolka, S.A., Steffen, B., Tofts, C.M.N.: Reactive, generative, and stratified models of probabilistic processes. In: Proceedings of the Fifth Annual Symposium on Logic in Computer Science (LICS '90), Philadelphia, Pennsylvania, USA, 4–7 June 1990, pp. 130–141 (1990). https://doi.org/10.1109/LICS.1990.113740
16. Hahn, E.M., Hermanns, H., Wachter, B., Zhang, L.: PARAM: a model checker for parametric Markov models. In: Touili, T., Cook, B., Jackson, P. (eds.) CAV 2010. LNCS, vol. 6174, pp. 660–664. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-14295-6_56

17. Hansson, H., Jonsson, B.: A framework for reasoning about time and reliability. In: Proceedings of the Real-Time Systems Symposium - 1989, Santa Monica, California, USA, December 1989, pp. 102–111. IEEE Computer Society (1989). https://doi.org/10.1109/REAL.1989.63561. https://ieeexplore.ieee.org/xpl/conhome/268/proceeding
18. Hansson, H., Jonsson, B.: A logic for reasoning about time and reliability. Formal Asp. Comput. **6**(5), 512–535 (1994). https://doi.org/10.1007/BF01211866
19. Hutschenreiter, L., Baier, C., Klein, J.: Parametric Markov chains: PCTL complexity and fraction-free gaussian elimination. In: Proceedings Eighth International Symposium on Games, Automata, Logics and Formal Verification, GandALF 2017, Roma, Italy, 20–22 September 2017, pp. 16–30 (2017). https://doi.org/10.4204/EPTCS.256.2
20. Jonsson, B.: Modular verification of asynchronous networks. In: Proceedings of the Sixth Annual ACM Symposium on Principles of Distributed Computing, Vancouver, British Columbia, Canada, 10–12 August 1987, pp. 152–166 (1987). https://doi.org/10.1145/41840.41853
21. Jonsson, B.: A fully abstract trace model for dataflow and asynchronous networks. Distrib. Comput. **7**(4), 197–212 (1994). https://doi.org/10.1007/BF02280834
22. Jonsson, B., Larsen, K.G.: Specification and refinement of probabilistic processes. In: Proceedings of the Sixth Annual Symposium on Logic in Computer Science (LICS '91), Amsterdam, The Netherlands, 15–18 July 1991, pp. 266–277 (1991). https://doi.org/10.1109/LICS.1991.151651
23. Junges, S., et al.: Parameter synthesis for Markov models. CoRR abs/1903.07993 (2019). http://arxiv.org/abs/1903.07993
24. Kwiatkowska, M., Norman, G., Parker, D.: PRISM 4.0: verification of probabilistic real-time systems. In: Gopalakrishnan, G., Qadeer, S. (eds.) CAV 2011. LNCS, vol. 6806, pp. 585–591. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22110-1_47
25. Larsen, K.G., Skou, A.: Bisimulation through probabilistic testing. In: Conference Record of the Sixteenth Annual ACM Symposium on Principles of Programming Languages, Austin, Texas, USA, 11–13 January 1989, pp. 344–352 (1989). https://doi.org/10.1145/75277.75307
26. Larsen, K.G., Skou, A.: Bisimulation through probabilistic testing. Inf. Comput. **94**(1), 1–28 (1991). https://doi.org/10.1016/0890-5401(91)90030-6
27. Larsen, K.G., Thomsen, B.: A modal process logic. In: Proceedings of the Third Annual Symposium on Logic in Computer Science (LICS '88), Edinburgh, Scotland, UK, 5–8 July 1988, pp. 203–210. IEEE Computer Society (1988). https://doi.org/10.1109/LICS.1988.5119. https://ieeexplore.ieee.org/xpl/conhome/203/proceeding
28. Petrucci, L., van de Pol, J.: Parameter synthesis algorithms for parametric interval Markov chains. In: Baier, C., Caires, L. (eds.) FORTE 2018. LNCS, vol. 10854, pp. 121–140. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-92612-4_7
29. Sen, K., Viswanathan, M., Agha, G.: Model-checking Markov chains in the presence of uncertainties. In: Hermanns, H., Palsberg, J. (eds.) TACAS 2006. LNCS, vol. 3920, pp. 394–410. Springer, Heidelberg (2006). https://doi.org/10.1007/11691372_26
30. Winskel, G.: A complete proof system for SCCS with modal assertions. In: Maheshwari, S.N. (ed.) FSTTCS 1985. LNCS, vol. 206, pp. 392–410. Springer, Heidelberg (1985). https://doi.org/10.1007/3-540-16042-6_22