



Model Checking Constrained Markov Reward Models with Uncertainties

Giovanni Bacci^(✉), Mikkel Hansen, and Kim Guldstrand Larsen

Department of Computer Science, Aalborg University, Aalborg, Denmark
{giovbacci,mhan,kg1}@cs.aau.dk

Abstract. We study the problem of analysing Markov reward models (MRMs) in the presence of imprecise or uncertain rewards. Properties of interests for their analysis are (i) probabilistic bisimilarity, and (ii) specifications expressed as probabilistic reward CTL formulae.

We consider two extensions of the notion of MRM, namely (a) *constrained Markov reward models*, i.e., MRMs with rewards parametric on a set variables subject to some constraints, and (b) *stochastic Markov reward models*, i.e., MRMs with rewards modelled as real-valued random variables as opposed to precise rewards. Our approach is based on quantifier elimination for *linear* real arithmetic. Differently from existing solutions for parametric Markov chains, we avoid the manipulation of rational functions in favour of a symbolic representation of the set or parameter valuations satisfying a given property in the form of a quantifier-free first-order formula in the linear theory of the reals.

Our work finds applications in model repair, where parameters need to be tuned so as to satisfy the desired specification, as well as in robustness analysis in the presence of stochastic perturbations.

Keywords: Parameter synthesis · Markov chains · Markov reward models · Model checking

1 Introduction

The wide-spread diffusion of cyber-physical systems (CPSs) poses the challenge of handling their growing complexity, while meeting requirements on correctness, predictability, performance without compromising time- and cost-to-market. Their analysis requires one to address a number of non-functional properties related to the quantitative aspects that are typical of such systems.

Finite-state Markovian models are popular modelling formalisms for the quantitative analysis of systems with probabilistic behaviours. Among these, Markov reward models (MRMs) were proposed as a natural extension of the usual notion of discrete-time Markov chain with (real-valued) state rewards.

Interesting properties of MRMs may be expressed by means of quantitative extensions of CTL, such as Probabilistic Reward CTL (PRCTL), *cf.* e.g., [1].

Work supported by the Advanced ERC Grant nr. 867096 (LASSO) and the Innovation Fund Denmark Center DiCyPS (nr. 864701).




Fig. 1. Model of a Ph.D. student in CS.

As a running example, consider the MRM in Fig. 1 modelling the stress level of a Ph.D. student in computer science. It has seven states, namely s_1, \dots, s_7 , annotated with the propositions *thk* (thinking), *ths* (thesis), *wr* (writing paper), *tc* (tool coding), and *pe* (paper evaluation). The level of stress of the student is modelled by associating with each state a reward that represents the stress that the student accumulates in the state at each time unit. After spending some time thinking, the student starts developing a tool with probability $p_c = 0.2$, or writing a paper with probability $p_w = 0.5$ (40% of the times for a journal and 60% for a conference), otherwise the student submits the thesis with probability 0.1. Once the tool is mature enough, the student starts writing a paper about it with probability p_w (70% of which for a journal, otherwise for a conference). When a paper has been completed it is submitted for evaluation with probability $p_s = 0.3$. The paper may be rejected or accepted (moving resp. to s_7 or s_5) according to some acceptance rate. At any moment before the thesis is completed, the student may move back to s_1 with probability $p_t = 0.2$ to think on how to proceed with his/her Ph.D.

We can use PRCTL formulae for specifying properties such as “the average level of stress accumulated until the thesis is submitted doesn’t exceed 10”, “the student completes the thesis without passing through a state with stress level higher than 2 with probability greater than or equal to 0.9”, or “the probability of eventually start coding and, subsequently, submitting the thesis with expected accumulated stress within $[x, y]$ is less than or equal to z ”.

In the above MRM we fixed specific reward values. This is an unrealistic simplification since the level of stress at each state may vary depending on different factors, possibly external to the student. The same argument applies to CPSs, that typically rely on sensor measurements which are inherently imprecise, but also to other systems, e.g., resource-management protocols that depends on stochastic assumptions on the future workload.

Typically, there are two ways for dealing with uncertain measurements: (i) determine the precision of the instrument and associate an error ϵ with each measurement, or (ii) perform estimation statistics (e.g., by recursive Bayesian estimation [16]) and use random variables to model each measurement.

In this paper we address the problem of analysing MRMs in presence of imprecise rewards. We present two extensions of the notion of MRM: (i) *constrained Markov reward models* (CMRMs), i.e., MRMs with rewards and transition probabilities parametric on a set variables subject to some constraints expressed as a first-order formula in the linear theory of the reals; (ii) *stochastic Markov reward models* (SMRMs), i.e., MRMs with rewards modelled as real-valued random variables as opposed to exact rewards.

Intuitively, a CMRM models a family of MRMs arising by plugging in concrete values for the parameters provided that they satisfy its constraints. Analogously, an SMRM models a probability distribution over a family of MRMs.

We are interested in the analysis of these models with respect to (i) probabilistic bisimilarity, and (ii) specifications expressed as PRCTL formulae.

On the one hand, the analysis of CMRMs is done by inferring constraints over its parameters characterising the valuations satisfying the property then, verify the robustness of the model within the given precision. On the other hand, analysing an SMRM consists in measuring the likelihood that an MRM obtained by plugging in real-valued outcomes of its random variables satisfies a given property.

Our Contribution: In relation with the analysis of CMRMs our contribution is twofold. First, we show that the computation of the set of parameter valuations ensuring that some states are probabilistic bisimilar with each other can be done using quantifier elimination in the linear fragment of the theory of the reals. Secondly, we demonstrate that for CMRMs having parameters only on state rewards, also the PRCTL model checking problem can be solved within the linear fragment of the theory of the reals. This allows one to employ SMT solvers and quantifier elimination procedures specialised on linear constraints [14], avoiding to use more inefficient and less scalable solvers for bilinear inequalities [17] as suggested e.g., in the case of interval Markov chains [15].

As for the analysis of SMRMs, we describe how one can apply the results for CMRMs described above for estimating the probability of satisfying a given property by employing Monte Carlo simulation techniques.

Related Work: The analysis of CMRMs falls into the research area of parameter synthesis for parametric Markov chains (pMCs) [4, 6–10, 12] and Interval Markov chains (IMC) [2, 3, 15]. Most of these works [4, 5, 7, 8, 10] consider pMCs without state rewards and focus mainly on computing closed form solutions for reachability probabilities in pMCs as rational functions. The same approach was recently extended to the computation of closed form solutions for expected rewards and long-run average rewards [6, 9].

In the context of model checking pMCs and IMCs, quantifier elimination for the first-order theory of the reals has been mainly applied as a theoretical tool for proving complexity upper-bounds of the model checking problem for IMCs [2, 3, 15] and pMCs [9]. Notably, the model checker PROPhESY [5] relies on SMT solving via the existential theory of the reals to determine approximate covering of the parameter space.

2 Preliminaries and Notation

We denote by \mathbb{R} , \mathbb{Q} , and \mathbb{N} respectively the sets of real, rational, and natural numbers. We denote by Σ^n , Σ^* and, Σ^ω respectively the set of words of length $n \in \mathbb{N}$, finite length, and infinite length, built over the finite alphabet Σ .

The dot product of two vectors $\mathbf{a} = (a_1, \dots, a_n)$ and $\mathbf{b} = (b_1, \dots, b_n)$ in \mathbb{R}^n is defined as $\mathbf{a} \cdot \mathbf{b} = \sum_{i=1}^n a_i b_i$. For $c \in \mathbb{R}$, we use $\mathbf{c} \in \mathbb{R}^n$ to denote the constant vector (c, \dots, c) . For a function $f: I \rightarrow \mathbb{R}$ with finite domain with cardinality $|I| = n$ we may use f to denote also the I -indexed vector $(f(i))_{i \in I} = (f(i_1), \dots, f(i_n))$.

Measure Theory. Let Ω be a set. A family $\Sigma \subseteq 2^\Omega$ is called a σ -algebra if it contains the empty set and is closed under complement and countable unions, in this case (Ω, Σ) is called a *measurable space* and elements of Σ *measurable sets*. When $\Sigma = 2^\Omega$, the measure space (Ω, Σ) is *discrete*.

A *measure* on (Ω, Σ) is a σ -additive function $\mu: \Sigma \rightarrow \mathbb{R}$, i.e. a map satisfying $\mu(\bigcup_{i \in I} E_i) = \sum_{i \in I} \mu(E_i)$ for any countable family of pairwise disjoint measurable sets $(E_i)_{i \in I}$, in this case (Ω, Σ, μ) is called a *measure space*. If additionally μ satisfies $\mu(\Omega) = 1$, it is called a *probability measure* and (Ω, Σ, μ) a *probability space*. We denote by $\mathcal{D}(\Omega)$ the set of discrete probability distributions on Ω . For $x \in \Omega$, the *Dirac distribution* concentrated at x is the distribution $1_x \in \mathcal{D}(\Omega)$ defined, for arbitrary $y \in \Omega$, as $1_x(y) = 1$ if $x = y$, 0 otherwise.

For measurable spaces (Ω, Σ) and (Γ, Θ) , a map $f: \Omega \rightarrow \Gamma$ is *measurable* if for all $E \in \Theta$, $f^{-1}(E) = \{x \mid f(x) \in E\} \subseteq \Sigma$. Given a measurable map $f: \Omega \rightarrow \Gamma$ and a measure μ on (Ω, Σ) we define the measure $\mu[f]$ on (Γ, Θ) , called the *push forward of μ under f* , as $\mu[f](E) = \mu(f^{-1}(E))$ for $E \in \Theta$.

A real-valued *random variable* $X: \Omega \rightarrow \mathbb{R}$ is a measurable function from a probability space (Ω, Σ, P) to the Borel space $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$. Intuitively, X can be seen as the outcome of some experiment, e.g. measuring some sensor value. Given a “test” $A \in \mathcal{B}(\mathbb{R})$, we write $P[X \in A]$ for the probability that X has a value in A , i.e. $P[X \in A] = P[X](A)$. A random variable is associated with its *cumulative distribution function* (CDF) $F_X: \mathbb{R} \rightarrow [0, 1]$ defined as $F_X(x) = P[X \in (-\infty, x]]$; and a *probability density function* (PDF) f_X , a non-negative Lebesgue-integrable function satisfying $P[X \in [a, b]] = \int_a^b f_X(x) dx$. The *expected value* of X , written $E[X]$, is intuitively understood as the long-run average of repetitions of the experiment X , formalised by the Lebesgue integral $\int_\Omega X dP$ (corresponding to $\int_{\mathbb{R}} f_X(x) dx$ when X admits density function f_X).

3 Markov Reward Models

In this section we recall the definitions of Markov reward models (MRMs), *probabilistic bisimulation*, and *probabilistic reward CTL* (PRCTL).

In what follows we fix a finite set of atomic propositions \mathcal{AP} .

Definition 1. A Markov chain is a tuple $\mathcal{M} = (S, \tau, \ell)$ consisting of a finite nonempty set of states S , a transition probability function $\tau: S \rightarrow \mathcal{D}(S)$, and a labelling function $\ell: S \rightarrow 2^{\mathcal{AP}}$ mapping states to atomic propositions.

Intuitively, if \mathcal{M} is in state s it moves to state s' with probability $\tau(s)(s')$. In this sense, \mathcal{M} can be thought as a state-machine that generates paths in S^ω . A *path* is an infinite sequence of states $\pi = s_1 s_2 s_3 \dots \in S^\omega$; for $i \geq 1$ we denote by $\pi[i]$ the i -th state of π , i.e., $\pi[i] = s_i$, and $\pi|^i$ the prefix of length i of π , i.e., $\pi|^i = s_1 \dots s_i$ and $\pi|^0 = \varepsilon$.

We denote by $G_{\mathcal{M}} = (S, \rightarrow)$ the underlying graph of \mathcal{M} , where $s, s' \in S$ are connected by a directed edge, written $s \rightarrow s'$, if and only if $\tau(s)(s') > 0$. We indicate by \rightarrow^* , the transitive and reflexive closure of \rightarrow .

In order to associate probabilities to events, we adopt the classical cylinder set construction (cf. [1, Ch10]). For $w \in S^*$, the cylinder set of w is the set of all paths having prefix w , i.e., $\text{cyl}(w) = wS^\omega$. Given an *initial probability distribution* $\iota \in \mathcal{D}(S)$, we define the probability space $(S^\omega, \Sigma_{\mathcal{M}}, Pr_{\iota}^{\mathcal{M}})$, where $\Sigma_{\mathcal{M}} = \sigma(\{\text{cyl}(w) \mid w \in S^*\})$ is the smallest σ -algebra that contains all the cylinder sets, and $Pr_{\iota}^{\mathcal{M}}$ is the unique probability measure such that, for all $w = s_1 \dots s_n \in S^*$, $Pr_{\iota}^{\mathcal{M}}(\text{cyl}(w)) = \iota(s_1) \cdot \prod_{0 < i < n} \tau(s_i)(s_{i+1})$.

When $\iota = 1_s$ we write $Pr_s^{\mathcal{M}}$, or just Pr_s when \mathcal{M} is clear from the context.

Definition 2. A Markov reward model is a tuple $\mathcal{R} = (S, \tau, \ell, \rho)$ where (S, τ, ℓ) is a MC, and $\rho : S \rightarrow \mathbb{R}$ is the reward function assigning a reward to each state.

A Markov reward model generates paths in S^ω according to its underlying Markov chain; in addition, whenever a transition is performed, say from s to s' , the system is rewarded by $\rho(s)$. It is worth noting that the reward is given after leaving the current state.

In the following, it may be convenient to represent the reward function ρ and the probability transition distributions $\tau(s)$ (for $s \in S$) as S -indexed vectors.

Example 1. Consider the MRM depicted in Fig. 1. Its reward function ρ and the distribution $\tau(s_1)$ are respectively represented in vector notation as

$$\rho = (0, 0, 1.2, 1.5, -3, 2, 3), \quad \text{and} \quad \tau(s_1) = (p_t, 0.1, 0.6p_w, 0.4p_w, 0, p_c, 0).$$

The next definition extends the classic notion of probabilistic bisimulation for Markov chains by Larsen and Skou [13] to the case of MRMs.

Definition 3 (Bisimulation). Let $\mathcal{R} = (S, \tau, \ell, \rho)$ be an MRM. An *equivalence relation* $R \subseteq S \times S$ is a probabilistic bisimulation if whenever $(s, s') \in R$ then, $\ell(s) = \ell(s')$, $\rho(s) = \rho(s')$, and for all R -equivalence class $C \in S/R$, $\tau(s)(C) = \tau(s')(C)$.

Two states $s, s' \in S$ are probabilistic bisimilar, written $s \sim_{\mathcal{R}} s'$, if they are related by some probabilistic bisimulation.

Example 2. Consider the MRM depicted in Fig. 1. As it is presented, none of its states are bisimilar with each other because each pair of states differs on the rewards or the labels. If we consider instead an MRM with underlying chain as in Fig. 1 and rewards $\rho = (0, 0, 1, 1, 3, 2, 3)$, now we have that $s_5 \sim s_7$ and, consequently, $s_3 \sim s_4$ because $\tau(s_3)(\{s_5, s_7\}) = p_s = \tau(s_4)(\{s_5, s_7\})$. \square

Before presenting probabilistic reward CTL, we need to introduce the concept of expected cumulative reward for reaching a set of states $B \subseteq S$.

Let $\hat{\rho}_B: S^\omega \rightarrow \mathbb{R}$ be the random variable representing the reward accumulated along a prefix of the path belonging to $(S \setminus B)^*B$. This is formalised as $\hat{\rho}_B(\pi) = \sum_{i=1}^{n-1} \rho(\pi[i])$ if $\pi|n \in (S \setminus B)^*B$ for some $n \in \mathbb{N}$, otherwise 0.

We denote by $E[\hat{\rho}_B|s]$ the expected value of $\hat{\rho}_B$ with respect to the probability distribution Pr_s . Following [1, Ch10], $E[\hat{\rho}_B|s]$ can be computed as the (unique) solution r_s of the following system of linear equations

$$r_s = \begin{cases} 0 & \text{if } s \in B \text{ or } s \not\rightarrow^* B \\ \rho(s) + (\sum_{t \in S} \tau(s)(t) \cdot r_t) & \text{otherwise.} \end{cases} \quad (1)$$

We are now ready to present Probabilistic Reward CTL. PRCTL allows for *state formulae* describing properties about states in a MRM, and *path formulae* describing properties about paths in a MRM. State formulae Φ, Ψ and path formulae φ are constructed over the following abstract syntax:

$$\Phi, \Psi ::= \# \mid a \mid \rho \bowtie r \mid \neg\Phi \mid \Phi \wedge \Psi \mid \mathbb{P}_J(\varphi) \mid \mathbb{E}_R(\Phi), \quad (\text{STATE FORMULAE})$$

$$\varphi ::= \mathcal{X}\Phi \mid \Phi\mathcal{U}\Psi \mid \Phi\mathcal{U}_R^n\Psi \quad (\text{PATH FORMULAE})$$

where $a \in \mathcal{AP}$, $r \in \mathbb{Q}$, $n \in \mathbb{N}$, $\bowtie \in \{=, <, >\}$, and $J \subseteq [0, 1]$ and $R \subseteq \mathbb{R}$ are intervals with rational bounds.

Given a MRM $\mathcal{R} = (S, \tau, \ell, \rho)$, a state $s \in S$, and a path $\pi \in S^\omega$, we denote by $\mathcal{R}, s \models \Phi$ (resp. $\mathcal{R}, \pi \models \varphi$) the fact that the state s satisfies the state formula Φ (resp. the path π satisfies the path formula φ). $Sat(\Phi)$ denotes the set of all states satisfying the property Φ , i.e. $Sat(\Phi) = \{s \in S \mid \mathcal{R}, s \models \Phi\}$. Formally, the *satisfiability relation* \models is inductively defined as:

$\mathcal{R}, s \models \#$		always holds
$\mathcal{R}, s \models a$	iff	$a \in \ell(s)$
$\mathcal{R}, s \models \rho \bowtie r$	iff	$\rho(s) \bowtie r$
$\mathcal{R}, s \models \neg\Phi$	iff	$\mathcal{R}, s \not\models \Phi$
$\mathcal{R}, s \models \Phi \wedge \Psi$	iff	$\mathcal{R}, s \models \Phi$ and $\mathcal{R}, s \models \Psi$
$\mathcal{R}, s \models \mathbb{P}_J(\varphi)$	iff	$Pr_s(\{\pi \in S^\omega \mid \mathcal{R}, \pi \models \varphi\}) \in J$
$\mathcal{R}, s \models \mathbb{E}_R(\Phi)$	iff	$E[\hat{\rho}_{Sat(\Phi)} s] \in R$
$\mathcal{R}, \pi \models \mathcal{X}\Phi$	iff	$\mathcal{R}, \pi[2] \models \Phi$
$\mathcal{R}, \pi \models \Phi\mathcal{U}\Psi$	iff	there exists $j \geq 1$ such that $\mathcal{R}, \pi[j] \models \Psi$, and $\mathcal{R}, \pi[j'] \models \Phi$ for all $1 \leq j' < j$
$\mathcal{R}, \pi \models \Phi\mathcal{U}_R^n\Psi$	iff	there exists $1 \leq j \leq n$ such that $\mathcal{R}, \pi[j] \models \Psi$ and for all $1 \leq k \leq j$, $\sum_{i=1}^{k-1} \rho(\pi[i]) \in R$ and, for all $1 \leq h < j$, $\mathcal{R}, \pi[h] \models \Phi$.

As usual, we can derive the logical operators ff , \vee , and \rightarrow as follows: $\text{ff} \stackrel{\text{def}}{=} \neg \#$, $\Phi \vee \Psi \stackrel{\text{def}}{=} \neg(\neg\Phi \wedge \neg\Psi)$, and $\Phi \rightarrow \Psi \stackrel{\text{def}}{=} \neg\Phi \vee \Psi$. Similarly, we can derive the temporal operators \diamond and \square as: $\diamond\Phi \stackrel{\text{def}}{=} \# \mathcal{U}\Phi$ and $\square\Phi \stackrel{\text{def}}{=} \neg\diamond(\neg\Phi)$.

Example 3. Consider the MRM \mathcal{R} depicted in Fig. 1. We can verify that “the average level of stress accumulated by the student until the thesis is submitted doesn’t exceed 23” by proving that $\mathcal{R}, s_1 \models \mathbb{E}_{\leq 23}(\text{ths})$ holds true. Analogously, we can check that the property “the student completes the thesis without passing through a state with stress level higher than 2 with probability greater than or equal to 0.5” does not hold true by proving $\mathcal{R}, s_1 \models \neg\mathbb{P}_{\geq 0.5}((\rho \leq 2) \mathcal{U} \text{ths})$. \square

4 Constrained Markov Reward Models

Constrained Markov reward models (CMRMs) model families of MRMs where both transition probabilities and state rewards are parametric on a set of real-valued variables subject to constraints expressed as a first-order formula in the linear theory of the reals.

Let $\mathbf{x} = (x_1, \dots, x_k)$ be a vector of real-valued parameters. We denote by \mathcal{E} the set of affine maps $f : \mathbb{R}^k \rightarrow \mathbb{R}$ of the form $f(\mathbf{x}) = \mathbf{a} \cdot \mathbf{x} + b$ with $\mathbf{a} = (a_1, \dots, a_k) \in \mathbb{Q}^k$ and $b \in \mathbb{Q}$, i.e. $f(x_1, \dots, x_k) = (\sum_{i=1}^k a_i x_i) + b$.

Definition 4. A constrained Markov reward model is a tuple $\mathcal{F} = (S, \tau, \ell, \rho, F)$ where S and ℓ are defined as for MCs, $\tau : S \rightarrow (S \rightarrow \mathcal{E})$ is a parametric transition function, $\rho : S \rightarrow \mathcal{E}$ is a parametric reward function, and $F(\mathbf{x})$ is a linear first-order formula s.t., for all $s \in S$, $F(\mathbf{x})$ implies $\tau(s) \geq \mathbf{0} \wedge \mathbf{1} \cdot \tau(s) = 1$.

Intuitively, a CMRM $\mathcal{F} = (S, \tau, \ell, \rho, F)$ defines a family of MRMs arising by plugging in concrete values for the parameters. A parameter valuation $\mathbf{v} \in \mathbb{R}^k$ is *admissible* (or *feasible*) if $F(\mathbf{v})$ holds true. By abuse of notation, we may write F to indicate the set of admissible valuations. Given $\mathbf{v} \in F$ we denote by $\mathcal{F}(\mathbf{v})$ the MRM associated with \mathbf{v} . In this respect, it will be convenient to think of \mathcal{F} as a function $\mathcal{F} : F \rightarrow \text{MRM}$. The semantics of \mathcal{F} , written $[\mathcal{F}]$, is defined as the image of \mathcal{F} , i.e. $[\mathcal{F}] = \{\mathcal{F}(\mathbf{v}) \mid \mathbf{v} \in F\}$.

Example 4. As already mentioned in the introduction, the stress level of a Ph.D. student may be influenced by several factors. For instance, we can define the following CMRM $\mathcal{F} = (S, \tau, \ell, \bar{\rho}, F)$ having as underlying Markov chain that of Fig. 1, parametric vector of rewards $\rho = (0, 0, h + c, h + j, r, 0, a)$, and constraints $F = (-5 \leq h \leq 0) \wedge (1 \leq c \leq j \leq 5) \wedge (-3 \leq a + r \leq 3) \wedge a \leq r$. Here, the parameter h models the help of the supervisor, whereas c and j represent the stress that accumulates while writing a conference paper c , or a journal paper j . The stress caused by having a paper accepted or rejected is modelled by the parameters a and r respectively.

It is easy to note that the MRM in Fig. 1 is an instance of \mathcal{F} . \square

Remark 1. Our notion of CMRM trivially extends that of IMC w.r.t. the so-called uncertain Markov chain interpretation [15]. The condition that transitions probabilities lie within some intervals is trivially expressed as a conjunction of linear inequalities. We shall also recall, that our definition is similar to that of *augmented* pMCs given by Hutschenreiter et al. [9]. Nevertheless, we do not require all admissible valuations to induce the same underlying graph; instead, we require linear constraints as opposed to polynomial constraints. \square

Given a CMRM \mathcal{F} we are interested in finding suitable symbolic representations of the set of parameter valuations that ensure the corresponding MRMs to enjoy some property. Properties of interests for this paper are bisimulation and PRCTL formulae. Specifically, given states $s, s' \in S$, and a PRCTL state formula Φ , we are interested in the set of parameter valuations for which s becomes bisimilar to s' or such that Φ holds at s . These are formalised as follows.

$$\llbracket s \sim_{\mathcal{F}} s' \rrbracket \stackrel{\text{def}}{=} \{ \mathbf{v} \in F \mid s \sim_{\mathcal{F}(\mathbf{v})} s' \}, \quad \llbracket \mathcal{F}, s \models \Phi \rrbracket \stackrel{\text{def}}{=} \{ \mathbf{v} \in F \mid \mathcal{F}(\mathbf{v}), s \models \Phi \} .$$

Example 5. Consider the CMRM \mathcal{F} defined in Example 4 and the PCTL formula $\Phi = \mathbb{E}_{[0,10]}(ths)$. The set of feasible parameter valuations satisfying Φ , i.e., $\llbracket \mathcal{F}, s \models \Phi \rrbracket$, can be symbolically represented by the following formula

$$\phi = 0 \leq 57a + 575c + 750h + 175j + 168r \leq \frac{1750}{3} \wedge F . \quad (2)$$

The above formula can be used for instance to determine how much effort the supervisor needs to provide to ensure that the formula Φ is satisfied. In this case, this can be done by maximising h under the constraints ϕ . \square

5 Analysing Bisimilarity on Constrained MRMs

In this section we address the problem of finding symbolic representations of the set of parameter valuations ensuring that two given states of a CMRM are probabilistic bisimilar.

Before diving in this problem, it will be convenient to present an alternative characterisation for probabilistic bisimilarity on MRMs. The characterisation is a variant of that of Jonsson and Larsen [11] based on the notion of *coupling*.

Given $\mu, \nu \in \mathcal{D}(X)$, we denote by $\Gamma(\mu, \nu)$ the set of *couplings* for (μ, ν) , i.e., probability distributions $\gamma \in \mathcal{D}(X \times X)$ such that, for all $x, y \in X$

$$\sum_{x' \in X} \gamma(x', y) = \nu(y) \quad \text{and} \quad \sum_{y' \in X} \gamma(x, y') = \mu(x) . \quad (3)$$

The probability measures μ and ν are respectively called left and right marginals of γ . A coupling can be seen as a redistribution of the “probability mass” from the right marginal to the left and *vice versa*.

Lemma 1. *$s \sim s'$, if and only if there exists a relation $R \subseteq S \times S$, such that $s R s'$ and, whenever $m R n$ the following conditions hold*

- (i) $\ell(m) = \ell(n)$, $\rho(m) = \rho(n)$, and
 (ii) there exists $\gamma \in \Gamma(\tau(m), \tau(n))$ such that $\{(u, v) \in S \times S \mid \gamma(u, v) > 0\} \subseteq R$.

The above lemma allows us to encode $\llbracket s \sim_{\mathcal{F}} s' \rrbracket$ as the following first-order formula in the linear theory of the reals.

$$\llbracket s \sim_{\mathcal{F}} s' \rrbracket \stackrel{\text{def}}{=} F \wedge \exists \mathbf{b}. b_{ss'} = 1 \wedge \bigwedge_{m, n \in S} \beta(m, n) \wedge \gamma(m, n) \wedge \sigma(m, n) \quad (4)$$

where the sub-formulae β , γ , and σ are defined as follows:

$$\beta(m, n) \stackrel{\text{def}}{=} \bigwedge_{m, n \in S} (b_{mn} = 0 \vee b_{mn} = 1) \quad (5)$$

$$\gamma(m, n) \stackrel{\text{def}}{=} (b_{mn} = 1) \rightarrow \exists \mathbf{c} \geq \mathbf{0}. \left(\begin{array}{l} \bigwedge_{u \in S} \sum_{v \in S} c_{uv} = \tau(m)(u) \wedge \\ \bigwedge_{v \in S} \sum_{u \in S} c_{uv} = \tau(n)(v) \wedge \\ \bigwedge_{u, v \in S} c_{uv} \leq b_{uv} \end{array} \right) \quad (6)$$

$$\sigma(m, n) \stackrel{\text{def}}{=} (b_{mn} = 1) \rightarrow (\ell(m) = \ell(n) \wedge \rho(m) = \rho(n)). \quad (7)$$

In the formula $\llbracket s \sim_{\mathcal{F}} s' \rrbracket$, $\mathbf{b} = (b_{mn})_{m, n \in S}$ represents a selection of a relation $R \in S \times S$ by means of binary variables (cf. (5)) such that $b_{m, n} = 1$ iff $(m, n) \in R$. According to Lemma 1, the selection needs to include the pair (s, s') and satisfy the conditions (i) and (ii). This is modelled by imposing $b_{ss'} = 1$, $\sigma(m, n)$ and $\gamma(m, n)$ for all $m, n \in S$. In particular, in (6), the variables $\mathbf{c} = (c_{uv})_{u, v \in S}$ model coupling for $(\tau(m), \tau(n))$ the condition (ii) is enforced by requiring $c_{u, v} \leq b_{u, v}$ for all $u, v \in S$.

By Lemma 1, the elimination of the existential quantifiers in $\llbracket s \sim_{\mathcal{F}} s' \rrbracket$ yields a Boolean formula with linear predicates on \mathbf{x} representing the set of valuations $\llbracket s \sim_{\mathcal{F}} s' \rrbracket$. Notably, the formula $\llbracket s \sim_{\mathcal{F}} s' \rrbracket$ is a first-order formula in the existential theory of the reals which involves only linear predicates. Therefore, quantifier elimination can be performed by using tools specialised for the linear theory of real-arithmetic, such as MJOLLNIR [14].

6 Model Checking Constrained MRMs

In this section we consider the model checking problem of CMRMs against PRCTL formulae for the class of constrained MRMs having parameters only on state rewards. For this class, we show that the set of parameter valuations satisfying given PRCTL formula can, again, be encoded as a first-order formula in the linear theory of real-arithmetic.

In this section assume that in the CMRM $\mathcal{P} = (S, \tau, \ell, \rho, F)$, parameters occur only in the state rewards, that is, $\tau(s)(u) \in [0, 1]$ for all $s, u \in S$.

For a state $s \in S$, and a PRCTL state formula Φ , we characterise $\llbracket \mathcal{P}, s \models \Phi \rrbracket$ by means of the set of valuations satisfying $F \wedge \llbracket s, \Phi \rrbracket$. The formula $\llbracket s, \Phi \rrbracket$ encodes the satisfiability of \mathcal{P} in s up-to the constraints F , and it is defined by induction on the structure of Φ .

For the Boolean fragment of PRCTL, the reduction is as one may expect.

$$\begin{aligned} \langle s, tt \rangle &= tt, & \langle s, a \rangle &= a \in \ell(s), & \langle s, \rho \bowtie r \rangle &= \rho(s) \bowtie r, \\ \langle s, \neg\Phi \rangle &= \neg\langle s, \Phi \rangle, & \langle s, \Phi \wedge \Psi \rangle &= \langle s, \Phi \rangle \wedge \langle s, \Psi \rangle. \end{aligned}$$

For all other formulae, it is convenient to use $\langle u, \diamond\Phi \rangle$ as short for $\bigvee_{u \rightarrow^* v} \langle v, \Phi \rangle$.

$$\begin{aligned} \langle s, \mathbb{P}_J(\mathcal{X}\Phi) \rangle &= \exists \mathbf{p}. \mathbf{1} \cdot \mathbf{p} \in J \wedge \bigwedge_{u \in S} ((p_u = \tau(s)(u) \wedge \langle u, \Phi \rangle) \vee (p_u = 0 \wedge \neg\langle u, \Phi \rangle)) \\ \langle s, \mathbb{P}_J(\Phi\mathcal{U}\Psi) \rangle &= \exists \mathbf{p}. p_s \in J \wedge \bigwedge_{u \in S} \left(\begin{array}{l} (p_u = 1 \wedge \langle u, \Psi \rangle) \vee (p_u = 0 \wedge \neg\langle u, \diamond\Psi \rangle) \vee \\ (p_u = 0 \wedge \neg\langle u, \Phi \rangle \wedge \neg\langle u, \Psi \rangle) \vee \\ (p_u = \tau(u) \cdot \mathbf{p} \wedge \langle u, \Phi \rangle \wedge \neg\langle u, \Psi \rangle \wedge \langle u, \diamond\Psi \rangle) \end{array} \right) \\ \langle s, \mathbb{P}_J(\Phi\mathcal{U}_R^n\Psi) \rangle &= \exists \mathbf{q}. \mathbf{1} \cdot \mathbf{q} \in J \wedge \bigwedge_{w \in S^n} \left(\begin{array}{l} (q_w = Pr_s(cyl(w)) \wedge \beta(w)) \vee \\ (q_w = 0 \wedge \neg\beta(w)) \end{array} \right) \\ \langle s, \mathbb{E}_R(\Phi) \rangle &= \exists \mathbf{r}. r_s \in R \wedge \bigwedge_{u \in S} ((r_u = 0 \wedge \alpha_u) \vee (r_u = \rho(u) + \tau(u) \cdot \mathbf{r} \wedge \neg\alpha_u)) \end{aligned}$$

where $\mathbf{p} = (p_s)_{s \in S}$, $\mathbf{r} = (r_s)_{s \in S}$, $\mathbf{q} = (q_w)_{w \in S^n}$, $\alpha_u = \langle u, \Phi \rangle \vee \neg\langle u, \diamond\Phi \rangle$, and

$$\beta(s_1 \dots s_n) = \bigvee_{1 \leq j \leq n} \left(\langle s_j, \Psi \rangle \wedge \bigwedge_{1 \leq k \leq j} \sum_{i=1}^{k-1} \rho(s_i) \in R \wedge \bigwedge_{1 \leq h < j} \langle s_h, \Phi \rangle \right).$$

In $\langle s, \mathbb{P}_J(\mathcal{X}\Phi) \rangle$, the variable p_u represents the probability of moving from s to u , where u satisfies the property Φ , therefore $\mathbf{1} \cdot \mathbf{p} = \sum_{u \in S}$ is the probability of moving in one step from s to a state satisfying Φ . Analogously, in $\langle s, \mathbb{P}_J(\Phi\mathcal{U}\Psi) \rangle$, the variable p_u models the probability of u satisfying the property $\Phi\mathcal{U}\Psi$: if u satisfies the property Φ then p_u is 1; if u satisfies neither Φ nor Ψ or cannot reach a state satisfying Ψ , then p_u is 0; otherwise, p_u amounts to the probability of moving in one step to some other state u' and, from there, satisfying $\Phi\mathcal{U}\Psi$.

In $\langle s, \mathbb{P}_J(\Phi\mathcal{U}_R^n\Psi) \rangle$, q_w is the probability of executing the trace w in exactly n steps starting from s where $\beta(w)$ models that fact that w satisfies $\Phi\mathcal{U}_R^n\Psi$. Therefore, $\mathbf{1} \cdot \mathbf{q}$ amounts to the probability of executing from s a path that satisfies $\Phi\mathcal{U}_R^n\Psi$.

Finally, in $\langle s, \mathbb{E}_R(\Phi) \rangle$, one can readily see that \mathbf{r} models the unique solution of the system of linear equations (1), where formula α_u capture the fact that u either satisfies the property Φ or cannot reach any state that does.

The following result states the correctness of the above characterisation.

Theorem 1. $\mathbf{v} \models F \wedge \langle s, \Phi \rangle$ iff $\mathbf{v} \in \llbracket \mathcal{P}, s \models \Phi \rrbracket$.

Example 6. Consider the CMRM \mathcal{F} and the formula Φ from Example 5. The encoding of the satisfiability of s_1 with respect to Φ , namely $\langle s, \Phi \rangle$, boils down to the following first order linear formula.

$$\begin{aligned}
\exists \mathbf{x}. \quad & x_1 \geq 0 \wedge x_1 \leq 10 \wedge x_1 = \frac{x_1}{5} + \frac{x_2}{10} + \frac{3x_3}{10} + \frac{x_4}{5} + \frac{x_6}{5} \wedge \\
& x_2 = 0 \wedge x_3 = c + h + \frac{x_1}{5} + \frac{x_3}{2} + \frac{9x_5}{100} + \frac{21x_7}{100} \wedge \\
& x_4 = h + j + \frac{x_1}{5} + \frac{x_3}{2} + \frac{3x_5}{100} + \frac{27x_7}{100} \wedge \\
& x_5 = a + x_1 \wedge x_6 = \frac{x_1}{5} + \frac{3x_3}{20} + \frac{7x_4}{20} + \frac{3x_6}{10} \wedge x_7 = r + x_1
\end{aligned}$$

By using the tool MJOLLNIR [14] we are able to eliminate the existential quantifiers on $\mathbf{x} = (x_1, \dots, x_7)$ obtaining the formula ϕ of Eq. (2). \square

7 Markov Models with Stochastic Rewards

It's common practice to model experimental measurements by means of real-valued random variables distributed according to well studied families of distributions (e.g., normal or student's T).

In this section we introduce the notion of *stochastic Markov reward models* (SMRMs) where state rewards are real-valued random variables. Then, we present a PRCTL model checking framework for SMRMs.

From here on we fix the probability space (Ω, Σ, P) representing the environment where the experiments are performed, and we use \mathbb{Y} to denote the set of real-valued random variables of the form $Y: \Omega \rightarrow \mathbb{R}$.

Definition 5. A stochastic Markov reward model is a tuple $\mathcal{J} = (S, \tau, \ell, \rho)$ where (S, τ, ℓ) is a Markov chain, and $\rho: S \rightarrow \mathbb{Y}$ is a reward function assigning a real-valued random variable to each state.

An SMRM $\mathcal{J} = (S, \tau, \ell, \rho)$ can be intuitively interpreted as a measurable function $\mathcal{J}: \Omega \rightarrow \text{MRM}_{\mathcal{J}}$, where $\mathcal{J}(\omega)$ is the MRM having (S, τ, ℓ) as underlying Markov chain and $(\rho(s)(\omega))_{s \in S}$ as vector of rewards. Such interpretation justifies the intuition of \mathcal{J} being an experiment whose outcomes is a MRM.

The above intuition is formalised as follows. We denote by $\text{MRM}_{\mathcal{J}}$ the set of all MRMs having the same underlying Markov chain as \mathcal{J} . We construct the σ -algebra $\Sigma_{\mathcal{J}}$ as the family of sets $A \subseteq \text{MRM}_{\mathcal{J}}$ whose corresponding set of rewards vectors is Borel measurable in \mathbb{R}^m ($m = |S|$). Formally,

$$A \in \Sigma_{\mathcal{J}} \quad \text{iff} \quad A \subseteq \text{MRM}_{\mathcal{J}} \text{ and } \{\rho(\mathcal{R}) \mid \mathcal{R} \in A\} \in \mathcal{B}(\mathbb{R}^m).$$

Accordingly, the semantics of \mathcal{J} is the probability space $(\text{MRM}_{\mathcal{J}}, \Sigma_{\mathcal{J}}, P[\mathcal{J}])$.

Given a SMRM \mathcal{J} , a state $s \in S$, and a PRCTL formula Φ , it comes natural to ask how likely is that a concrete instance of \mathcal{J} satisfies Φ at s , denoted by $P[\mathcal{J}, s \models \Phi]$. This model checking problem is formalised as follows

$$P[\mathcal{J}, s \models \Phi] \stackrel{\text{def}}{=} P[\mathcal{J}](\{\mathcal{R} \in \text{MRM}_{\mathcal{J}} \mid \mathcal{R}, s \models \Phi\}). \quad (8)$$

We study the above model checking problem for a subclass of SMRMs having random variables $(Y: \Omega \rightarrow \mathbb{R}) \in \mathcal{E}_{\mathbf{X}}$ of the form $Y(\omega) = \mathbf{a} \cdot \mathbf{X}(\omega) + b$, with $\mathbf{a} \in \mathbb{Q}^k$, $b \in \mathbb{Q}$ and, where $\mathbf{X} = (X_1, \dots, X_k)$ is a vector of pairwise *independent* real-valued random variables¹. Observe that, elements in $\mathcal{E}_{\mathbf{X}}$ may not be independent from each other.

Hereafter we consider the SMRM $\mathcal{J} = (S, \tau, \ell, \rho)$ with $\rho: S \rightarrow \mathcal{E}_{\mathbf{X}}$, and we use \mathcal{P} to refer to the CMRM obtained by replacing the random variables X_i in \mathcal{J} with the parameters x_i ($i = 1, \dots, k$).

For Eq. (8) to be well-defined the set $\{\mathcal{R} \in \text{MRM}_{\mathcal{J}} \mid \mathcal{R}, s \models \Phi\}$ needs to be a measurable event in $\Sigma_{\mathcal{J}}$. The following result ensures that.

Lemma 2. $\{\mathcal{R} \in \text{MRM}_{\mathcal{J}} \mid \mathcal{R}, s \models \Phi\} \in \Sigma_{\mathcal{J}}$.

Proof. By definition of \mathcal{P} we have that

$$\{\mathcal{R} \in \text{MRM}_{\mathcal{J}} \mid \mathcal{R}, s \models \Phi\} = \{\mathcal{P}(\mathbf{v}) \mid \mathbf{v} \in \llbracket \mathcal{P}, s \models \Phi \rrbracket\}$$

Therefore, by def. of $\Sigma_{\mathcal{J}}$ and measurability of affine transformations we have that the claim holds iff $\llbracket \mathcal{P}, s \models \Phi \rrbracket \in \mathcal{B}(\mathbb{R}^k)$. By Theorem 1, $\llbracket \mathcal{P}, s \models \Phi \rrbracket$ can be described by means of a Boolean formula with linear predicates. Since σ -algebras are closed under complement (i.e., negation), countable unions (i.e., disjunctions) and countable intersections (i.e., conjunctions) and, affine transformations are measurable, we conclude that $\{\mathcal{R} \in \text{MRM}_{\mathcal{J}} \mid \mathcal{R}, s \models \Phi\} \in \Sigma_{\mathcal{J}}$.

The following theorem characterises the model checking problem for the SMRM \mathcal{J} in terms of the model checking problem for the CMRM \mathcal{P} .

Theorem 2. $P[\mathcal{J}, s \models \Phi] = P[\mathbf{X} \in \llbracket \mathcal{P}, s \models \Phi \rrbracket]$.

Proof. The claim holds true according to the following equalities.

$$\begin{aligned} P[\mathcal{J}, s \models \Phi] &= P[\mathcal{J}](\{\mathcal{R} \in \text{MRM}_{\mathcal{J}} \mid \mathcal{R}, s \models \Phi\}) && \text{(by Eq. (8))} \\ &= P[\mathcal{P} \circ \mathbf{X}](\{\mathcal{R} \in \text{MRM}_{\mathcal{J}} \mid \mathcal{R}, s \models \Phi\}) && (\mathcal{J} = \mathcal{P} \circ \mathbf{X}) \\ &= P[\mathcal{P} \circ \mathbf{X}](\mathcal{P}(\llbracket \mathcal{P}, s \models \Phi \rrbracket)) && \text{(def. } \mathcal{P} \text{ and def. } \llbracket \mathcal{P}, s \models \Phi \rrbracket) \\ &= P((\mathcal{P} \circ \mathbf{X})^{-1}(\mathcal{P}(\llbracket \mathcal{P}, s \models \Phi \rrbracket))) && \text{(def. push-forward)} \\ &= P(\mathbf{X}^{-1}(\llbracket \mathcal{P}, s \models \Phi \rrbracket)) && ((\mathcal{P} \circ \mathbf{X})^{-1} = \mathbf{X}^{-1} \circ \mathcal{P}^{-1}) \\ &= P[\mathbf{X} \in \llbracket \mathcal{P}, s \models \Phi \rrbracket]. && \text{(def. push-forward)} \end{aligned}$$

□

By Theorem 2 we can estimate the value p of $P[\mathcal{J}, s \models \Phi]$ by applying Monte Carlo simulation techniques. For this, we sample n independent repetitions of \mathbf{X} , associating with each repetition a Bernoulli random variable B_i . A realisation b_i of B_i is 1 if the corresponding sampled value of \mathbf{X} lays in $\llbracket \mathcal{P}, s \models \Phi \rrbracket$, and

¹ In fact, the vector \mathbf{X} is a multivariate random variable $\mathbf{X}: \Omega \rightarrow \mathbb{R}^n$ with marginals $X_i: \Omega \rightarrow \mathbb{R}$ ($i = 1, \dots, n$).

0 otherwise. Finally, we estimate p by means of the observed relative success rate $\tilde{p} = (\sum_{i=1}^n b_i)/n$. The absolute error ε of the estimation can be bound with a certain degree of confidence $\delta \in (0, 1]$ by tuning the number of required simulations based on the Hoeffding's inequality $P(|\tilde{p} - p| \geq \varepsilon) \leq 2e^{-2n\varepsilon^2}$. Thus, we can determine the number of samples required to estimate p with absolute error ε and confidence δ by imposing $2e^{-2n\varepsilon^2} = 1 - \delta$, from which we obtain

$$n = \left\lceil -\frac{\ln(\delta/2)}{2\varepsilon^2} \right\rceil. \quad (9)$$

For example, we can estimate p with an error $\varepsilon = 0.01$ with a confidence of 95% (i.e., $\delta = 0.95$) by drawing $n = 18445$ samples.

Example 7. Let $\mathbf{X} = (H, C, J, A, R)$ be a vector of random variables respectively distributed as $H \sim \text{unif}(-5, 0)$, $C \sim \text{unif}(1, 2)$, $J \sim \mathcal{N}(2, 0.1)$, $A \sim \mathcal{N}(-3, 0.5)$, and $R \sim \mathcal{N}(-3, 0.5)$. We define the SMRM $\mathcal{J} = (S, \tau, \ell, \rho)$ having as underlying Markov chain that of Fig. 1, and the following random vector of rewards

$$\rho = (0, 0, H + C, H + J, R, 0, A),$$

which shall be understood as a stochastic version of the one presented in Example 4. By taking the formula $\Phi = \mathbb{E}_{[0,10]}(ths)$ from Example 5, we can estimate $P[\mathcal{J}, s_1 \models \Phi] \cong 0.156061$ with an error $\varepsilon = 0.005$ and confidence of 99.99% (i.e., $\delta = 0.0001$) by generating $n = 198070$ samples. \square

8 Conclusion

We described a framework for the analysis of Markov reward models in presence of uncertain rewards. To this end we propose two extensions of the notion of MRM: (a) *constrained Markov reward models*, having state rewards parametric on a set variables subject to some constraints, and (b) *stochastic Markov reward models*, having rewards modelled as random variables.

We demonstrated that the analysis CMRMs with respect to probabilistic bisimilarity and PRCTL formulae, can be reduced to perform quantifier elimination on first-order formulae in the linear fragment of the theory of the reals. Our reduction does not lead to an improvement on the theoretical complexity of the model checking problem for (augmented) parametric Markov chains (*cf.* [9, Theorem 4]). However, we believe that our reduction is important from the perspective of implementation in practice, because it allows one to employ SMT solvers or quantifier elimination procedures specialised on linear constraints such as MJOLLNIR [14]. It is worth noting that our reduction can be also applied with PRCTL formulas with parametric bounds, extending even further the applicability of our approach.

Finally, we provided a characterisation of the model checking problem for SMRMs in terms of the model checking problem for CMRMs. As we have shown, this result allows one to estimate the probability that a given SMRM satisfies a given specification by employing Monte Carlo simulation techniques.

All the calculations presented in the examples have been done by using a prototype implementation of the above described algorithms² coded in Mathematica [18]. For the quantifier elimination we employed MJOLLNIR instead of the built-in solution offered by Mathematica.

Our work finds applications in model repair, where parameters need to be tuned so as to satisfy a desired specification, as well as in robustness analysis in presence of stochastic perturbations.

References

1. Baier, C., Katoen, J.: Principles of Model Checking. MIT Press, Cambridge (2008)
2. 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
3. Chatterjee, K., Sen, K., Henzinger, T.A.: Model-checking ω -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
4. Daws, C.: Symbolic and parametric model checking of discrete-time Markov chains. In: Liu, Z., Araki, K. (eds.) ICTAC 2004. LNCS, vol. 3407, pp. 280–294. Springer, Heidelberg (2005). https://doi.org/10.1007/978-3-540-31862-0_21
5. Dehnert, C., Junges, S., Jansen, N., Corzilius, F., Volk, M., Brintjes, H., Katoen, J.-P., Ábrahám, E.: 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
6. Gainer, P., Hahn, E.M., Schewe, S.: Accelerated model checking of parametric Markov chains. In: Lahiri, S.K., Wang, C. (eds.) ATVA 2018. LNCS, vol. 11138, pp. 300–316. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01090-4_18
7. 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
8. Hahn, E.M., Hermanns, H., Zhang, L.: Probabilistic reachability for parametric markov models. STTT **13**(1), 3–19 (2011). <https://doi.org/10.1007/s10009-010-0146-x>
9. Hutschenreiter, L., Baier, C., Klein, J.: Parametric Markov chains: PCTL complexity and fraction-free Gaussian elimination. In: Bouyer, P., Orlandini, A., Pietro, P.S. (eds.) Proceedings 8th International Symposium on Games, Automata, Logics and Formal Verification, GandALF 2017. EPTCS, vol. 256, pp. 16–30 (2017). <https://doi.org/10.4204/EPTCS.256.2>
10. Jansen, N., Corzilius, F., Volk, M., Wimmer, R., Ábrahám, E., Katoen, J.-P., Becker, B.: Accelerating parametric probabilistic verification. In: Norman, G., Sanders, W. (eds.) QEST 2014. LNCS, vol. 8657, pp. 404–420. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10696-0_31
11. Jonsson, B., Larsen, K.G.: Specification and refinement of probabilistic processes. In: Proceedings of the Sixth Annual Symposium on Logic in Computer Science (LICS 1991), Amsterdam, The Netherlands, 15–18 July 1991, pp. 266–277. IEEE Computer Society (1991). <https://doi.org/10.1109/LICS.1991.151651>

² Available at <http://people.cs.aau.dk/~giobacci/tools.html>.

12. Lanotte, R., Maggiolo-Schettini, A., Troina, A.: Parametric probabilistic transition systems for system design and analysis. *Formal Asp. Comput.* **19**(1), 93–109 (2007). <https://doi.org/10.1007/s00165-006-0015-2>
13. Larsen, K.G., Skou, A.: Bisimulation through probabilistic testing. *Inf. Comput.* **94**(1), 1–28 (1991)
14. Monniaux, D.: Quantifier elimination by lazy model enumeration. In: Touili, T., Cook, B., Jackson, P. (eds.) *CAV 2010*. LNCS, vol. 6174, pp. 585–599. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-14295-6_51
15. 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
16. Srkk, S.: *Bayesian Filtering and Smoothing*. Cambridge University Press, New York (2013)
17. Tran, D.Q., Gumussoy, S., Michiels, W., Diehl, M.: Combining convex-concave decompositions and linearization approaches for solving BMIs, with application to static output feedback. *IEEE Trans. Automat. Contr.* **57**(6), 1377–1390 (2012). <https://doi.org/10.1109/TAC.2011.2176154>
18. Wolfram Research Inc.: *Mathematica*, Version 11.2, Champaign, IL (2017)