

Perturbation Analysis in Verification of Discrete-Time Markov Chains^{*}

Taolue Chen¹, Yuan Feng², David S. Rosenblum³, and Guoxin Su³

¹ Department of Computer Science, Middlesex University London

² Centre for Quantum Computation and Intelligent Systems,
University of Technology, Sydney

³ Department of Computer Science, School of Computing,
National University of Singapore

Abstract. Perturbation analysis in probabilistic verification addresses the *robustness* and *sensitivity* problem for verification of stochastic models against qualitative and quantitative properties. We identify two types of perturbation bounds, namely *non-asymptotic bounds* and *asymptotic bounds*. Non-asymptotic bounds are exact, pointwise bounds that quantify the upper and lower bounds of the verification result subject to a given perturbation of the model, whereas asymptotic bounds are closed-form bounds that approximate non-asymptotic bounds by assuming that the given perturbation is sufficiently small. We perform perturbation analysis in the setting of Discrete-time Markov Chains. We consider three basic matrix norms to capture the perturbation distance, and focus on the computational aspect. Our main contributions include algorithms and tight complexity bounds for calculating both non-asymptotic bounds and asymptotic bounds with respect to the three perturbation distances.

1 Introduction

Probabilistic verification techniques, and in particular probabilistic model checking, have been successfully applied to a variety of domains ranging from wireless communication protocols to dynamic power management schemes, and to systems biology and quantum cryptography. Mature probabilistic model checking tools such as PRISM [22] support verification of most existing stochastic models, e.g., Discrete-time Markov Chains (DTMCs), Markov Decision Processes (MDPs), and stochastic games, against a wide range of qualitative and quantitative properties.

When modelling real-life systems with stochastic models, one usually has to face the issue that these systems are governed by empirical or unknown distributions, such as the failure rate of some system component. As a result, measurements or experiments are employed to determine, for instance, transition probabilities (for discrete-time systems) or transition rates (for continuous-time systems). Those statistical quantities are

^{*} The work is partially supported by the Australian Research Council (Grant Nos. DP130102764 and FT100100218) and the Singapore Ministry of Education (Grant R-252-000-458-133). Y. Feng is also supported by the Overseas Team Program of the Academy of Mathematics and Systems Science, CAS and the CAS/SAFEA International Partnership Program for Creative Research Team. T. Chen is partially supported by the EU Erasmus Grant.

imprecise. In the worst case, a tiny but non-trivial change to some quantities in the model might lead to a misleading or even invalid verification result.

These issues motivate the following important problem for probabilistic verification: *If some of the quantities in the stochastic model are perturbed, what is the influence on verification of the model?* In other words, given a stochastic model, we need to measure the *robustness* and *sensitivity* of verification results. The purpose of perturbation analysis, as the central topic of the current paper, is to shed light on this problem.

A straightforward approach is to modify the model manually for each set of values of each perturbed quantity, and then perform model checking multiple times with a model checker such as PRISM. Such a solution is simple but unsatisfactory: It is resource-consuming while providing little information about the impact of model perturbations on verification. Instead, in this paper, we present a sound and rigorous approach to characterising the *maximal/minimal* variation that might occur to the verification outcome, with respect to a given perturbation of the model. This yields a measure for the sensitivity and robustness of these verification results. Such an analysis also potentially reduces the overall time of verifying a large number of similar stochastic models if only approximated results are required. More specifically, we pursue two types of perturbation bounds, namely *non-asymptotic bounds*⁴ and *asymptotic bounds*.

- Non-asymptotic bounds are pointwise bounds that quantify the maximum and the minimum of the verification result subject to a given perturbation of the model.
- Asymptotic bounds are closed-form, lightweight approximations of non-asymptotic bounds when the model perturbation is sufficiently small (i.e., close to 0).

The main task of perturbation analysis in the current paper is to compute these bounds, the formal definitions of which are presented in Section 2.

Contributions. In this paper, we focus on the computational aspect of perturbation analysis in DTMC verification. We consider three different perturbation distances for DTMCs based on three norms over stochastic matrices—the “entrywise” ∞ -norm, the induced ∞ -norm, and the “entrywise” 1-norm, which quantify the perturbation distance of DTMCs. These norms are widely adopted in literature (e.g. [8, 15, 31]), and somehow are easy to compute as they are “linear”. Henceforth, we refer to the three distances as Type I, II and III distances, respectively. Our key contributions, summarised in Table 1, include two aspects:

- We present algorithms to compute non-asymptotic bounds under Type I, II and III distances, respectively, and identify tight computational complexity bounds. For Type I and II distances, we present polynomial-time algorithms, while for Type III distance, we show that the computation (technically, the aligned decision problem) is in PSPACE and is SQUARE-ROOT-SUM hard.
- We provide a unified treatment for the asymptotic bounds of an arbitrary degree for distances of all three types. In particular, we show how to compute the *linear* and the *quadratic* asymptotic bounds. This subsumes the results reported previously [31] regarding linear asymptotic bounds.

⁴ The term is adopted from *non-asymptotic* analysis of random matrices and *non-asymptotic* information theory.

Table 1. Complexity of Computing Perturbation Bounds

Distance	Non-Asymptotic	Asymptotic	
		linear	quadratic
I & II	P	PL	NP
III	in PSPACE, SRS-hard		

The computation of non-asymptotic bounds is related to verification of stochastic models with uncertainty (e.g. Interval Markov Chains (IMCs)). Typically, two different semantics for IMCs are studied in the literature, namely *Uncertain Markov Chains* (UMCs) and *Interval Markov Decision Processes* (IMDPs) [28]. The non-asymptotic bounds adopt the UMC semantics. For Type I and II distances, since the UMC and IMDP semantics coincide, we apply a technique similar to the one by Puggelli *et al.* [24] to obtain polynomial-time algorithms for non-asymptotic bounds. However, for Type III distance, we can only obtain a PSPACE algorithm—we show a slightly better complexity upper bound, namely the complexity of the existential theory of reals. For the lower bound, we give a reduction from the well-known SQUARE-ROOT-SUM (SRS) problem. The exact complexity of the SRS problem, i.e., whether it is in P or even in NP, is open since 1976. This suggests that our PSPACE upper bound cannot be substantially improved without a breakthrough concerning this long-standing open problem.

The study of asymptotic bounds in probabilistic verification was initiated by Su *et al.* [30, 31], where linear and quadratic asymptotic bounds with respect to a single perturbation distance function are studied. Apart from giving a unified formulation of general asymptotic bounds under the three types of distances, the current paper also improves the complexity results reported in the previous work. Our main techniques for this are from multivariate calculus: We resort the problem to optimisation problems of multivariate polynomials under (virtually) linear constraints. For linear asymptotic bounds, this enables us to derive an analytical expression, whereas for quadratic asymptotic bounds, we exploit quadratic programming. We also identify complexity upper bounds for the two cases.

For simplicity, we focus on reachability in this paper. However, the presented techniques can be generalised for ω -regular properties and various performance properties, such as expected rewards and long-run average rewards, without substantial difficulty (and see our previous work [31] for ω -regular properties).

Related work. In general, perturbation theory for applied mathematics investigates solutions for mathematically formulated problems that involve parameters subject to perturbations [23]. There exists a line of research on perturbation analysis of DTMCs, the common goal of which is to find a suitable condition number for the distance of steady states and the distance of transition matrices between two DTMCs [27, 13, 29]. In these works, the condition number is defined as the supremum of the quotient of the deviation of the perturbed DTMCs and the allowed perturbation. The deviation is hence bounded universally for all DTMCs with respect to chosen norms of distance metrics. In the formal verification setting, the closest work is by Chatterjee [8], who studied the continuity and robustness of the value function in stochastic parity games with respect to

imprecision in the transition probabilities. This can be regarded as a (rough) perturbation analysis for stochastic games. However, non-asymptotic and asymptotic bounds are not considered there, nor is their computational aspect. It is an interesting direction for future work to extend our results to the game setting. Moreover, Desharnaisa *et al.* [15] addressed perturbation analysis for labelled Markov processes (LMPs). The authors defined a distance akin to the Type II distance and gave a bound on the difference between LMPs measured by a behaviour pseudo-metric with respect to the perturbation.

Most available verification results are on IMCs, (arguably) the simplest variant models of DTMCs with uncertainty. In particular, Sen *et al.* [28] proved that model checking IMC against probabilistic computational tree logic (PCTL) is NP-hard. More general results on IMCs against ω -regular properties are reported in [9]. Chen *et al.* [12] presented thorough results on the complexity of model checking IMCs against reachability and PCTL properties, under both the UMC and the IMDP semantics. Benedikt *et al.* [5] considered the LTL model checking problem for IMCs.

Other related work includes *parameter synthesis* for stochastic models [14, 21, 19, 20] and *model repair* [16, 4, 11]. In general, these studies attempt to identify some (or all) parameter configuration(s) in a parametric model such that a given property is satisfied. Hence, the approaches there are considerably different from ours.

Structure of the paper. The rest of the paper is structured as follows: Section 2 presents definitions of models, model distances, and non-asymptotic and asymptotic bounds. Section 3 presents results on computation of non-asymptotic bounds with respect to three types of distances. Section 4 presents results on computation of asymptotic bounds. Section 5 concludes the paper and outlines future work. An extended version of the paper contains proofs and more details [10].

2 Models, Distances and Perturbation Bounds

Given a finite set S , we use $\Delta(S)$ to denote the set of (discrete) *probability distributions* over S , i.e., functions $\mu : S \rightarrow [0, 1]$ with $\sum_{s \in S} \mu(s) = 1$.

Definition 1. A Discrete-time Markov Chain (DTMC) is a tuple $\mathcal{D} = (S, \alpha, \mathbf{P})$, where

- S is a finite set of states,
- $\alpha \in \Delta(S)$ is the initial distribution, and
- $\mathbf{P} : S \times S \rightarrow [0, 1]$ is a transition probability matrix such that for any state $s \in S$, $\sum_{s' \in S} \mathbf{P}(s, s') = 1$, i.e., $\mathbf{P}(s, \cdot) \in \Delta(S)$.

An (infinite) *path* in \mathcal{D} is a sequence $\pi = s_0 s_1 \dots$ such that $s_i \in S$ and $\mathbf{P}(s_i, s_{i+1}) > 0$ for each $i \geq 0$. Denote the i -th state of π (i.e., s_i) as $\pi[i]$, and the set of paths in \mathcal{D} as $\text{Paths}^{\mathcal{D}}$. The *probability distribution* $\text{Pr}^{\mathcal{D}}$ over $\text{Paths}^{\mathcal{D}}$ is defined in a standard way [3, Chapter 10].

Definition 2. A Markov Decision Process (MDP) is a tuple $\mathcal{M} = (S, \alpha, \mathcal{T})$, where

- S and α are defined the same as in Definition 1, and
- $\mathcal{T} : S \rightarrow \wp(\Delta(S))$ is the transition function s.t. $\mathcal{T}(s)$ is finite for each $s \in S$.

Without loss of generality, we assume that $\mathcal{T}(s) \neq \emptyset$ for each $s \in S$. At each state s of \mathcal{M} , a probability distribution μ (over S) is chosen *nondeterministically* from the set $\mathcal{T}(s)$. A successor state s' is then chosen according to μ with probability $\mu(s')$. An (infinite) path π in \mathcal{M} is a sequence of the form $s_0 \xrightarrow{\mu_1} s_1 \xrightarrow{\mu_2} \dots$ where $s_i \in S$, $\mu_{i+1} \in \mathcal{T}(s_i)$ and $\mu_{i+1}(s_{i+1}) > 0$ for each $i \geq 0$. A *finite* path is a prefix of an infinite path ending in a state. Let $Paths^{\mathcal{M}}$ be the set of finite paths. A *scheduler* $\sigma : Paths^{\mathcal{M}} \rightarrow \Delta(S)$ maps a finite path (the *history*) to a distribution over S such that for any finite path $\pi = s_0 \xrightarrow{\mu_1} \dots \xrightarrow{\mu_n} s_n$, $\sigma(\pi) \in \mathcal{T}(s_n)$. In particular, a *simple* scheduler σ chooses a distribution only based on the current state, and thus for each finite path π ending in s , $\sigma(\pi) = \sigma(s) \in \mathcal{T}(s)$. Note that we obtain a (possibly infinite-state) DTMC by fixing a scheduler in an MDP [3, 25]. In the sequel, we write \mathcal{M}_σ for such a DTMC given an MDP \mathcal{M} and a scheduler σ .

We often relax the definition of MDPs by allowing $\mathcal{T}(s)$ to be *infinite*. As long as $\mathcal{T}(s)$ is compact (for instance in the paper, $\mathcal{T}(s) \subseteq \mathbb{R}^{|S|}$ with respect to the Euclidean topology), most interesting properties for MDPs are carried over. This feature is made use of by existing work on IMDPs mentioned in the Introduction.

For the convenience of perturbation analysis, we also define a parametric variant of DTMCs [31]. When performing perturbation analysis for a DTMC in practice, it is usually required that some of transitions remain unchanged. To accommodate this, we specify a set of transitions $C \subseteq S \times S$ for a DTMC \mathcal{D} with state space S . The intuition behind this requirement is that only probabilities of transitions in C can be perturbed. The perturbed quantities are captured by a sequence of pair-wise distinct variables $\mathbf{x} = (x_1, \dots, x_k)$ with $k = |C|$.

Definition 3. *The parametric DTMC of \mathcal{D} on \mathbf{x} is a tuple $\mathcal{D}(\mathbf{x}) = (S, \alpha, \mathbf{P}, F)$ where F is a one-to-one mapping from C to the variable set $\{x_i\}_{1 \leq i \leq k}$.*

For simplicity, we denote by $\mathbf{P}(\mathbf{x})$ the parametric variation of \mathbf{P} with the (s, t) -entry being $\mathbf{P}(s, t) + F(s, t)$ if $(s, t) \in C$, and $\mathbf{P}(s, t)$ otherwise. We defer the specification of domains for variables from $\{x_i\}_{1 \leq i \leq k}$ in Section 2.2.

Reachability. For a given DTMC $\mathcal{D} = (S, \alpha, \mathbf{P})$, let $G \subseteq S$ be a set of *target* states. We consider the probability of reaching G . Formally, let $\diamond G = \{\pi \in Paths^{\mathcal{D}} \mid \pi[i] \in G \text{ for some } i \geq 0\}$. We are interested in $\Pr^{\mathcal{D}}(\diamond G)$. Let $S_0 = \{s \in S \mid \Pr^{\mathcal{D}}(s \models \diamond G) = 0\}$, and $S_\gamma = S \setminus (S_0 \cup G)$. Let $\tilde{\mathbf{P}}$ be the matrix obtained by restricting \mathbf{P} on S_γ . Then, $I - \tilde{\mathbf{P}}$ is invertible. Let \mathbf{b} be a vector on S_γ such that $\mathbf{b}[s] = \sum_{t \in G} \mathbf{P}(s, t)$ for each $s \in S_\gamma$. Let $\tilde{\alpha}$ be the restriction of α on S_γ . We have that $\Pr^{\mathcal{D}}(\diamond G) = \tilde{\alpha}^T (I - \tilde{\mathbf{P}})^{-1} \mathbf{b}$ [3, Chapter 10]. For an MDP \mathcal{M} with state space $S \supseteq G$, we can also define the *maximum* reachability probability $\sup_\sigma \Pr^{\mathcal{M}_\sigma}(\diamond G)$, which can be calculated efficiently by linear programming [25].

2.1 Distance between DTMCs

A DTMC (S, α, \mathbf{P}) induces a digraph in a standard way: The set of vertices of the digraph is S , and there is an edge from s to t iff $\mathbf{P}(s, t) > 0$. Given two DTMCs $\mathcal{D}_1 = (S, \alpha_1, \mathbf{P}_1)$ and $\mathcal{D}_2 = (S, \alpha_2, \mathbf{P}_2)$, we say that \mathcal{D}_1 and \mathcal{D}_2 are *structurally*

equivalent, denoted as $\mathcal{D}_1 \equiv \mathcal{D}_2$, if for each pair of states $s, t \in S$, $\mathbf{P}_1(s, t) > 0$ iff $\mathbf{P}_2(s, t) > 0$. Namely, \mathcal{D}_1 and \mathcal{D}_2 have the same underlying digraphs. We now identify three distances for two (structurally equivalent) DTMCs.

Definition 4 (Distance of DTMCs). Given two DTMCs \mathcal{D}_1 and \mathcal{D}_2 such that $\mathcal{D}_1 \equiv \mathcal{D}_2$, we define the distances d_I, d_{II} , and d_{III} as

- (1) $d_I(\mathcal{D}_1, \mathcal{D}_2) = \max_{s, t \in S} |\mathbf{P}_1(s, t) - \mathbf{P}_2(s, t)|$,
- (2) $d_{II}(\mathcal{D}_1, \mathcal{D}_2) = \max_{s \in S} \{ \sum_{t \in S} |\mathbf{P}_1(s, t) - \mathbf{P}_2(s, t)| \}$,
- (3) $d_{III}(\mathcal{D}_1, \mathcal{D}_2) = \sum_{s, t \in S} |\mathbf{P}_1(s, t) - \mathbf{P}_2(s, t)|$.

We call d_I, d_{II} , and d_{III} as Type I, Type II, and Type III distances respectively. We use \star to range over $\{I, II, III\}$, and d_\star to denote a generic distance definition.

Remark 1. In matrix theory, Type I is the distance induced from the ‘‘entrywise’’ ∞ -norm, Type II is induced from the ∞ -norm (which is an induced norm of the ∞ -norm for vectors), and Type III is induced from the ‘‘entrywise’’ 1-norm.

Let $C_s = \{t \in S \mid (s, t) \in C\}$ (C_s may be empty), and so $C = \bigsqcup_{s \in S} \{s\} \times C_s$. To simplify notations in the remainder of the paper, when given a DTMC, we fix an associated C . Accordingly, the distance definitions can be formulated as norms of variable vectors for parametric DTMCs. Recall that $\mathbf{x} = \{x_i\}_{1 \leq i \leq k}$ with $k = |C|$. Let $\mathbf{x}_s = (x_{s,t})_{t \in C_s}$ and so \mathbf{x} is a concatenation of \mathbf{x}_s for $s \in S$. Note that the distance between $\mathcal{D}(\mathbf{x})$ and \mathcal{D} is exactly the corresponding norm of \mathbf{x} which is defined as follows:

- (1) $\|\mathbf{x}\|_I = \max_{1 \leq i \leq k} |x_i|$,
- (2) $\|\mathbf{x}\|_{II} = \max_{s \in S} \{ \sum_{t \in C_s} |x_{s,t}| \}$,
- (3) $\|\mathbf{x}\|_{III} = \sum_{1 \leq i \leq k} |x_i|$.

2.2 Perturbation Bounds

For the purpose of perturbation analysis, we define two types of perturbation bounds, namely *non-asymptotic bounds* and *asymptotic bounds*, which are the main research object of this paper.

We write $\mathcal{D} \sim_C \mathcal{D}'$ if \mathcal{D} and \mathcal{D}' differ only for transitions in C , and let

$$[\mathcal{D}]_{\star, \delta} = \{ \mathcal{D}' \text{ is a DTMC} \mid \mathcal{D} \equiv \mathcal{D}', d_\star(\mathcal{D}, \mathcal{D}') \leq \delta \text{ and } \mathcal{D} \sim_C \mathcal{D}' \}.$$

Definition 5 (Non-Asymptotic bound). The upper and lower non-asymptotic bounds of a DTMC \mathcal{D} are defined as follows:

$$\rho_\star^+(\delta) = \sup \left\{ \Pr^{\mathcal{D}'}(\diamond G) \mid \mathcal{D}' \in [\mathcal{D}]_{\star, \delta} \right\} \text{ and } \rho_\star^-(\delta) = \inf \left\{ \Pr^{\mathcal{D}'}(\diamond G) \mid \mathcal{D}' \in [\mathcal{D}]_{\star, \delta} \right\}.$$

We use the subscript \star to emphasise the fact that $\rho_\star^+(\delta)$ and $\rho_\star^-(\delta)$ are dependent on the distance d_\star . We can present an alternative characterisation for non-asymptotic bounds with parametric DTMCs. Let

$$\mathbf{U}_{\star, \delta} = \{ \mathbf{x} \in \mathbf{R}^k \mid \mathcal{D}(\mathbf{x}) \in [\mathcal{D}]_{\star, \delta} \}.$$

Here, and in the sequel, we abuse the notation slightly to denote by $\mathcal{D}(\mathbf{x})$ the DTMC obtained by instantiating the variables in the corresponding parametric DTMC with the real vector \mathbf{x} . In particular, we have $\mathcal{D}(\mathbf{0}) = \mathcal{D}$ with $\mathbf{P}(\mathbf{0}) = \mathbf{P}$.

Note that, for each $\mathbf{x} \in \mathbf{U}_{\star, \delta}$, $\mathcal{D}(\mathbf{x})$ and \mathcal{D} are structurally equivalent and $I - \tilde{\mathbf{P}}(\mathbf{x})$ is invertible. We then write

$$p(\mathbf{x}) := \Pr^{\mathcal{D}(\mathbf{x})}(\diamond G) = \alpha^{\top} (I - \tilde{\mathbf{P}}(\mathbf{x}))^{-1} \mathbf{b}(\mathbf{x}). \quad (1)$$

There are alternative ways of generating or expressing $p(\mathbf{x})$ reported in [14, 19, 18]. Obviously, $p(\cdot)$ is a multivariate rational function on $\mathbf{U}_{\star, \delta}$ and thus is infinitely differentiable. It is then straightforward to observe that

$$\rho_{\star}^{+}(\delta) = \sup_{\mathbf{x} \in \mathbf{U}_{\star, \delta}} p(\mathbf{x}) \text{ and } \rho_{\star}^{-}(\delta) = \inf_{\mathbf{x} \in \mathbf{U}_{\star, \delta}} p(\mathbf{x}) \quad (2)$$

Also note that $\mathbf{U}_{\star, \delta}$ is convex and thus connected. Hence, by the continuity of $p(\cdot)$ and the Intermediate Value Theorem, for any value $y \in (\rho_{\star}^{-}(\delta), \rho_{\star}^{+}(\delta))$, there exists $\mathbf{x} \in \mathbf{U}_{\star, \delta}$ such that $p(\mathbf{x}) = y$.

Asymptotic bounds provide reasonably accurate approximations for $\rho_{\star}^{+}(\delta)$ and $\rho_{\star}^{-}(\delta)$ when $\delta > 0$ is close to 0. Let $r = \min_{(s,t) \in C} \{\mathbf{P}(s,t), 1 - \mathbf{P}(s,t)\} > 0$.

Definition 6 (Asymptotic bound). *An asymptotic bound of degree n for $\rho_{\star}^{+}(\cdot)$ (resp. $\rho_{\star}^{-}(\cdot)$) is a function $f_n^{+} : (0, r) \rightarrow \mathbf{R}$ (resp. $f_n^{-} : (0, r) \rightarrow \mathbf{R}$) such that*

$$f_n^{+}(\delta) - \rho_{\star}^{+}(\delta) = o(\delta^n) \text{ and } f_n^{-}(\delta) - \rho_{\star}^{-}(\delta) = o(\delta^n);$$

in other words,

$$\lim_{\delta \rightarrow 0} \frac{|f_n^{+}(\delta) - \rho_{\star}^{+}(\delta)|}{\delta^n} = 0 \text{ and } \lim_{\delta \rightarrow 0} \frac{|f_n^{-}(\delta) - \rho_{\star}^{-}(\delta)|}{\delta^n} = 0.$$

In words, Definition 6 states that, as δ tends to 0, the convergent rate of f_n^{+} (resp. f_n^{-}) to ρ_{\star}^{+} (resp. ρ_{\star}^{-}) is *at least* of order n or, equivalently, f_n^{+} (resp. f_n^{-}) approaches to $\rho_{\star}^{+}(\delta)$ (resp. $\rho_{\star}^{-}(\delta)$) at least as fast as any polynomial function on δ of degree n . We note that asymptotic perturbation bounds can be non-unique. We refer to asymptotic bounds of degree one as *linear asymptotic bounds* (*linear bounds* for short), and asymptotic bound of degree two as *quadratic asymptotic bounds* (*quadratic bounds* for short).

3 Computing Non-Asymptotic Bounds

In this section, we present algorithms for computing non-asymptotic bounds and analyse the complexity. An obvious fact about non-asymptotic bounds is given by the following proposition:

Proposition 1. *Given a DTMC \mathcal{D} and $\star \in \{\text{I, II, III}\}$, $\rho_{\star}^{+}(\cdot)$ and $\rho_{\star}^{-}(\cdot)$ are continuous functions in $(0, r)$.*

3.1 Type I and Type II Distances

In this section, we deal with non-asymptotic bounds under Type I and II distances given in Definition 4. In particular, we focus on Type I distance, while Type II distance can be dealt with in a similar way.

In general terms, our strategy is to reduce the computation of $\rho_I^+(\cdot)$ and $\rho_I^-(\cdot)$ to linear programming via MDPs. Let $\mathcal{D} = (S, \alpha, \mathbf{P})$. Consider an MDP $\mathcal{M} = (S, \alpha, \mathcal{T})$, where for each state $s \in S$,

$$\mathcal{T}(s) = \{\mu \in \Delta(S) \mid |\mu(s') - \mathbf{P}(s, s')| \leq \delta \text{ for any } s' \in S\}$$

We have the following proposition:

Proposition 2. $\rho_I^+(\delta) = \sup_{\sigma} \Pr^{\mathcal{M}_{\sigma}}(\diamond G)$ and $\rho_I^-(\delta) = \inf_{\sigma} \Pr^{\mathcal{M}_{\sigma}}(\diamond G)$.

Proposition 2 allows us to reduce the problem of computing $\rho_I^+(\delta)$ (resp. $\rho_I^-(\delta)$) to computing the maximum (resp. minimum) reachability probability for the MDP \mathcal{M} , and the latter is resorted to the standard linear programming technique. Note that here the MDP \mathcal{M} is merely a tool which can simplify the technical development, and that we are not considering verification of ‘‘perturbed’’ MDPs. *Below we only present an algorithm for $\rho_I^+(\delta)$, since an algorithm for $\rho_I^-(\delta)$ can be obtained in a dual manner.*

Recall that $\mathbf{x} = (x_{s,t})_{(s,t) \in C}$ is a concatenation of vectors $\mathbf{x}_s = (x_{s,t})_{t \in C_s}$ for each $s \in S$. Intuitively, $x_{s,t}$ captures the perturbed quantity at the (s, t) -entry of the transition probability matrix \mathbf{P} . We introduce a new vector of variables $\mathbf{y} = (y_s)_{s \in S}$. Intuitively, y_s captures the probability to reach G from state s . For each state s , $\Omega(s)$ is a set of vectors defined as:

$$\mathbf{x}_s \in \Omega(s) \text{ iff } \begin{cases} \sum_{t \in C_s} x_{s,t} = 0 \\ 0 \leq \mathbf{P}(s, t) + x_{s,t} \leq 1, \text{ for each } t \in C_s \\ -\delta \leq x_{s,t} \leq \delta, \text{ for each } t \in C_s \end{cases} \quad (3)$$

For simplicity, we also write

$$\Gamma(\mathbf{x}_s, \mathbf{y}) = \sum_{t \in C_s} (\mathbf{P}(s, t) + x_{s,t}) \cdot y_t + \sum_{t \notin C_s} \mathbf{P}(s, t) \cdot y_t.$$

Then, we consider the following (pseudo-) linear program (which can be derived directly from the MDP formulation [25, 6]):

$$\begin{aligned} & \text{minimise} && \sum_{s \in S} \alpha(s) y_s \\ & \text{subject to} && y_s \geq \max_{\mathbf{x}_s \in \Omega(s)} \Gamma(\mathbf{x}_s, \mathbf{y}) \quad s \notin G \\ & && y_s = 1 \quad s \in G \end{aligned} \quad (4)$$

Note that, for a fixed \mathbf{y} , $\max_{\mathbf{x}_s \in \Omega(s)} \Gamma(\mathbf{x}_s, \mathbf{y})$ in Problem (4) is itself a linear program where the constraint is given in (3). (It also follows that $\max_{\mathbf{x}_s \in \Omega(s)} \Gamma(\mathbf{x}_s, \mathbf{y})$ does exist although $\Omega(s)$ is infinite.) We denote its (Lagrange) *dual function* as $\min_{\lambda_s} \Gamma'(\lambda_s, \mathbf{y})$,

where λ_s is the Lagrange multiplier vectors (dual variables) for the linear program. Strong duality implies that

$$\max_{\mathbf{x}_s} \Gamma(\mathbf{x}_s, \mathbf{y}) = \min_{\lambda_s} \Gamma'(\lambda_s, \mathbf{y}) \quad (5)$$

Hence, Problem (4) becomes the following problem:

$$\begin{aligned} & \text{minimise} && \sum_{s \in S} \alpha(s) y_s \\ & \text{subject to} && y_s \geq \min_{\lambda_s} \Gamma'(\lambda_s, \mathbf{y}) \quad s \notin G \\ & && y_s = 1 \quad s \in G \end{aligned} \quad (6)$$

It is not hard to observe that Problem (6) is equivalent to the following problem:

$$\begin{aligned} & \text{minimise} && \sum_{s \in S} \alpha(s) y_s \\ & \text{subject to} && y_s \geq \Gamma'(\lambda_s, \mathbf{y}) \quad s \notin G \\ & && y_s = 1 \quad s \in G \end{aligned} \quad (7)$$

Note that Problem (7) is a linear program and is solvable in polynomial time.

By a similar argument (detailed in our extended paper [10]) we can demonstrate that ρ_{II}^+ can be computed in polynomial time. We conclude our results in this subsection by the following theorem:

Theorem 1. *Given any DTMC, $\rho_{\text{I}}^+(\cdot)$, $\rho_{\text{I}}^-(\cdot)$, $\rho_{\text{II}}^+(\cdot)$, $\rho_{\text{II}}^-(\cdot)$ can be computed in polynomial time.*

Remark 2. It is worth mentioning that Chen *et al.* [12] gave a thorough answer on the complexity of model checking IMCs against PCTL under both the UMC and the IMDP semantics. The main technique there is (a generalised version of) the ellipsoid algorithm for linear programming. Their approach can also be adopted here to tackle the problem for Type I and II distances. However, the technique exploited here (and by Pergelli *et al.* [24]) allows us to use off-the-shelf linear program solvers (e.g., Matlab), while the approach by Chen *et al.* requires more efforts in implementation. Furthermore, our extended paper [10] presents more practical algorithms based on a “value iteration” scheme from MDPs, which underpin the tool support.

3.2 Type III Distance

In this section, we focus on Type III distance given in Definition 4. We note that the technique employed in the previous subsection for Type I and II distances *cannot* be used here. Nevertheless, we still formulate the problem as an optimisation problem




Fig. 1. Examples of DTMCs: (a) $\mathcal{D}_{a,b}$ and (b) \mathcal{D}_0

(with the same optimisation variables as in Problem (4) in the previous subsection):

$$\begin{aligned}
& \text{maximise} && \sum_{s \in S} \alpha(s) y_s \\
& \text{subject to} && y_s = 1 && \text{for } s \in G \\
& && y_s = \sum_{t \in S} (\mathbf{P}(s, t) + x_{s,t}) \cdot y_t && \text{for } s \notin G \\
& && \sum_{(s,t) \in C} |x_{s,t}| \leq \delta && (8) \\
& && \sum_{t \in C_s} x_{s,t} = 0 && \text{for all } s \in S \\
& && 0 \leq \mathbf{P}(s, t) + x_{s,t} \leq 1 && \text{for each } (s, t) \in C
\end{aligned}$$

Clearly, Problem (8) is *not* a convex programming problem owing to the bilinear form $\sum_{t \in S} (\mathbf{P}(s, t) + x_{s,t}) \cdot y_t$. However, an obvious observation is that the corresponding decision problem, that is, deciding whether $\rho_{\text{III}}^+(\delta) \geq \theta$ for some given rational $\theta \in [0, 1]$, can be formulated in the existential theory of reals. Since the decision problem of the existential theory of reals is in PSPACE [7], a PSPACE complexity upper bound follows.

Proposition 3. *The problem to decide whether $\rho_{\text{III}}^+(\delta) \geq \theta$ for a given rational $\theta \in [0, 1]$ is in PSPACE.*

Below, we show that the computation of $\rho_{\text{III}}^+(\delta)$ is unlikely to admit a polynomial-time algorithm. Moreover, even the achievement of an NP upper bound (for its corresponding decision problem) is difficult. We illustrate this by an example.

Example 1. Consider a DTMC $\mathcal{D}_{a,b}$ depicted in Fig. 1(a). The transition matrix of $\mathcal{D}_{a,b}$ and its perturbed matrix are respectively as follows:

$$\begin{bmatrix} 0 & b & a & 1-a-b \\ 0.5 & 0 & 0 & 0.5 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \text{ and } \begin{bmatrix} 0 & b & a & 1-a-b \\ 0.5-x & 0 & 0 & 0.5+x \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

where $0 < a, b < a + b < 1$ and x ranges over $(-0.5, 0.5)$. Basic calculation reveals that the probability of reaching s_4 from s_1 is

$$1 - \frac{a}{1 - b(0.5 - x)} = 1 - \frac{a/b}{1/b - 0.5 + x}.$$

We construct another DTMC \mathcal{D}_0 , depicted in Fig. 1(b). \mathcal{D}_0 contains n “copies” of $\mathcal{D}_{a,b}$, each denoted by \mathcal{D}_{a_i,b_i} with $1 \leq i \leq n$, and global initial state s_0 that has $1/n$ probability to reach each initial state $s_{i,1}$ in each \mathcal{D}_{a_i,b_i} . Then, the probability of reaching states in $\{s_{i,4} \mid 1 \leq i \leq n\}$ from s_0 is captured by

$$p(x_1, \dots, x_n) = 1 - \frac{1}{n} \sum_{i=1}^n \frac{a_i/b_i}{1/b_i - 0.5 + x_i}.$$

Let $\rho_{\text{III}}^+(2\delta)$ be the *maximum* of $p(\cdot)$ subject to $\sum_{i=1}^n |x_i| \leq \delta$ and $x_i \in (-0.5, 0.5)$ for each $1 \leq i \leq n$. Essentially, to compute $\rho_{\text{III}}^+(2\delta)$, we need to *minimise*

$$1 - p(x_1, \dots, x_n) = \frac{1}{n} \sum_{i=1}^n \frac{a_i/b_i}{1/b_i - 0.5 + x_i}$$

subject to the same constraints. Since $1 - p(\cdot)$ is a decreasing function, it is not hard to see that, in order to minimise it, each x_i must be non-negative and thus $\sum_{i=1}^n x_i = \delta$.

By the Cauchy–Schwarz inequality,

$$\sum_{i=1}^n (1/b_i - 0.5 + x_i) \cdot \sum_{i=1}^n \frac{a_i/b_i}{1/b_i - 0.5 + x_i} \geq \left(\sum_{i=1}^n \sqrt{\frac{a_i}{b_i}} \right)^2.$$

Namely, $\sum_{i=1}^n \frac{a_i/b_i}{1/b_i - 0.5 + x_i} \geq \left(\sum_{i=1}^n \sqrt{\frac{a_i}{b_i}} \right)^2 / (\sum_{i=1}^n 1/b_i - 0.5n + \delta)$. The equality in the above inequality holds if and only if

$$x_i = \frac{\sqrt{a_i/b_i}}{\sum_{i=1}^n \sqrt{a_i/b_i}} \cdot \left(\sum_{i=1}^n (1/b_i - 0.5) + \delta \right) - (1/b_i - 0.5) \geq 0.$$

Clearly, in general $\rho_{\text{III}}^+(2\delta)$ is *not* a rational number, and neither are x_i ’s. This example suggests the insight that in general one would not be able to use linear programming to compute $\rho_{\text{III}}^+(\cdot)$, which is in a sharp contrast to $\rho_{\text{I}}^+(\cdot)$ and $\rho_{\text{II}}^+(\cdot)$.

With a generalisation of Example 1, we can show that the SQUARE-ROOT-SUM (SRS) problem can be reduced to deciding whether $\rho_{\text{III}}^+(\delta) \geq \theta$ for some given rational $\theta \in [0, 1]$. An instance of the SRS problem is the decision of $\sum_{i=1}^n \sqrt{t_i} \leq y$ for a given tuple (t_1, \dots, t_n, y) of natural numbers (greater than 1). The reduction is involved and is detailed in our extended paper [10].

Proposition 4. *Given a DTMC, deciding $\rho_{\text{III}}^+(\delta) \geq \theta$ for given δ and θ is SQUARE-ROOT-SUM hard.*

Using a similar construction, one can also show that computing $\rho_{\text{III}}^-(\cdot)$ is SRS hard. The SRS problem has been studied extensively, especially in computational geometry where the square root sum represents the sum of Euclidean distances between given pairs of points with integer/rational coordinates.⁵ Allender *et al.* [1] showed that this

⁵ For example, determining whether the length of a TSP tour of a set of points on the plane is bounded by a given threshold can be easily encoded as the SRS problem.

problem is decidable in the 4-th level of the Counting Hierarchy (an analogue of the polynomial-time hierarchy for counting classes); hence it is unlikely to be PSPACE-hard. But it remains open whether the problem can be decided in P or even in NP. Interesting examples that are related to formal verification can be studied by Etessami and Yannakakis [17], among others.

4 Computing Asymptotic Bounds

In this section, we consider the computation of asymptotic bounds. Recall that the reachability probability $p(\mathbf{x})$ (cf. (1)) is smooth, namely, infinitely differentiable. We present a unified characterisation for ρ_\star^+ and ρ_\star^- with $\star \in \{\text{I, II, III}\}$ using the Taylor expansion of $p(\mathbf{x})$. Define the following multi-variate index notations: Let $\iota = (\iota_1, \dots, \iota_k)$ be a vector of integers. Let

$$|\iota| = \iota_1 + \dots + \iota_k, \quad \iota! = \iota_1! \dots \iota_k!, \quad \mathbf{x}^\iota = x_1^{\iota_1} \dots x_k^{\iota_k}$$

$$\text{and } \nabla^\iota p(\mathbf{x}) = \frac{\partial^{|\iota|} p(\mathbf{x})}{\partial x_1^{\iota_1} \dots \partial x_k^{\iota_k}}.$$

Recall that $r = \min_{(s,t) \in C} \{\mathbf{P}(s,t), 1 - \mathbf{P}(s,t)\}$. For $n \in \mathbb{N}$, let $g_{\star,n}^+ : (0, r) \rightarrow \mathbf{R}$ such that $g_{\star,n}^+(\delta)$ is the solution of the following optimisation problem:

$$\begin{aligned} & \text{maximise} && \sum_{1 \leq |\iota| \leq n} \frac{\nabla^\iota p(\mathbf{0})}{\iota!} \mathbf{x}^\iota \\ & \text{subject to} && \mathbf{x} \in \mathbf{U}_{\star,\delta} \end{aligned} \tag{9}$$

Theorem 2. For each $\star \in \{\text{I, II, III}\}$, $g_{\star,n}^+(\cdot)$ is an asymptotic bound of degree n for ρ_\star^+ .

An asymptotic bound of degree n for ρ_\star^- can be obtained in a similar way as in Problem (9) by replacing *maximise* by *minimise*. We hence focus on the maximum case. The remainder of this section presents a method for computing $g_{\star,1}^+(\cdot)$ and $g_{\star,2}^+(\cdot)$, namely, the linear and quadratic bounds.

4.1 Linear Bounds

The linear bound of ρ_\star^+ , $g_{\star,1}^+(\cdot)$, can be obtained by instantiating Problem (9) with $n = 1$. We show that $g_{\star,1}^+(\delta) = \kappa_\star \delta$, where κ_\star is a solution of the following optimisation problem:

$$\begin{aligned} & \text{maximise} && \nabla p(\mathbf{0}) \cdot \mathbf{x} \\ & \text{subject to} && \sum_{t \in C_s} x_{s,t} = 0 \text{ for each } s \in S \\ & && \|\mathbf{x}\|_\star = 1 \end{aligned} \tag{10}$$

We write $\nabla p(\mathbf{0})$ as $\mathbf{h} = (h_1, \dots, h_k)$, which is of dimension $k = |C|$. Then \mathbf{h} can be computed according to the following proposition.

Proposition 5. For each $1 \leq i \leq k$,

$$h_i = \alpha^T [I - \tilde{\mathbf{P}}]^{-1} \tilde{\mathbf{P}}_i [I - \tilde{\mathbf{P}}]^{-1} \mathbf{b} + \alpha^T [I - \tilde{\mathbf{P}}]^{-1} \mathbf{b}_i$$

where $\tilde{\mathbf{P}}_i$ is the matrix on S_γ such that $\tilde{\mathbf{P}}_i(s, t) = 1$ if $\tilde{\mathbf{P}}(\mathbf{x})(s, t)$ contains x_i and 0 otherwise, and \mathbf{b}_i is the vector on S_γ such that $\mathbf{b}_i(s) = 1$ if $\mathbf{b}(\mathbf{x})(s)$ contains x_i and 0 otherwise.

When instantiated with the three types of norms (corresponding to the three types of distances, respectively), we obtain analytical solutions of Problem (10) for each $\star \in \{\text{I, II, III}\}$.

Proposition 6. The following statements hold:

- Let $\mathbf{h}_s = (h_{s,t})_{t \in C_s}$ (i.e., \mathbf{h} is a concatenation of \mathbf{h}_s for each $s \in S$). We sort each \mathbf{h}_s in non-decreasing order to get $\mathbf{h}'_s = (h'_{s,1}, \dots, h'_{s,k_s})$ where $k_s = |C_s|$. Then

$$\kappa_{\text{I}} = \sum_{s \in S} \left(\sum_{1 \leq i \leq \lfloor k_s/2 \rfloor} h'_{s,i} - \sum_{\lceil k_s/2 \rceil + 1 \leq i \leq k_s} h'_{s,i} \right);$$

- $\kappa_{\text{II}} = \sum_{s \in S} \frac{1}{2} (\max_{t \in C_s} \{h_{s,t}\} - \min_{t \in C_s} \{h_{s,t}\})$; and
- $\kappa_{\text{III}} = \max_{s \in S} \frac{1}{2} (\max_{t \in C_s} \{h_{s,t}\} - \min_{t \in C_s} \{h_{s,t}\})$

As κ_\star for each $\star \in \{\text{I, II, III}\}$ is nonnegative, Theorem 2 immediately implies the following theorem.

Theorem 3. For each $\star \in \{\text{I, II, III}\}$, $\kappa_\star \delta$ is a linear bound for ρ_\star^+ .

Essentially, computing κ_\star boils down to computing an inverse matrix, which can be done by Gaussian elimination. Hence we have

Proposition 7. The problem of computing linear bounds is in $\mathcal{O}(|\mathcal{D}|^3)$.

Remark 3. We can show that computing κ_\star can be done in GapL, which concerns logspace-bounded computation. In a nutshell, $\sharp L$ is defined, in analogy to $\sharp P$, to be the set of functions that count the number of accepting computation paths of a nondeterministic logspace-bounded Turing machine. The class GapL is defined by Allender and Ogihara [2] to be the closure of $\sharp L$ under subtraction. Furthermore, the decision version of computing κ_\star is in PL (probabilistic logspace), and the technical details can be found in the extended version of our paper [10].

4.2 Quadratic Bounds

Similar to the linear case, we can instantiate Problem (9) to obtain the quadratic bound $g_{\star,2}^+$. However, it is usually inefficient to solve Problem (9) for every given δ . Instead, we show that there exists a uniform direction vector for \mathbf{x} for all sufficiently small δ , along

which the quadratic bound is obtained. For this purpose, we consider an alternative optimisation problem:

$$\begin{aligned}
& \text{maximise} && \sum_{|t|=2} \frac{\nabla^t p(\mathbf{0})}{t!} \mathbf{x}^t \\
& \text{subject to} && \sum_{t \in C_s} x_{s,t} = 0 \text{ for each } s \in S \\
& && \|\mathbf{x}\|_* = 1 \text{ and } \mathbf{h} \cdot \mathbf{x} = \kappa_*
\end{aligned} \tag{11}$$

For each $\star \in \{\text{I, II, III}\}$, let v_\star be the solution of Problem (11). The following technical result states that the coefficient of the linear term of $g_{\star,2}^+$ is exactly κ_\star , and the coefficient of the quadratic term of $g_{\star,2}^+$ is v_\star (obtained by solving Problem (11)).

Theorem 4. *For each $\star \in \{\text{I, II, III}\}$, the function $v_\star \delta^2 + \kappa_\star \delta$ is a quadratic bound for ρ_\star^+ .*

To compute the quadratic bounds, we must solve a quadratic program, which is known to be NP-complete [26]. The following result is rather straightforward.

Proposition 8. *The problem of computing quadratic bounds is in TFNP (namely Total Function NP).*

5 Conclusion

In this paper, we have performed an in-depth study on perturbation analysis in the setting of DTMC verification. We defined non-asymptotic and asymptotic perturbation bounds and focused on their computation. In particular, we considered three fundamental matrix norms for stochastic matrices to quantify perturbations of DTMCs. With respect to these distances, we presented algorithms and complexity analysis for computing the non-asymptotic and asymptotic perturbation bounds.

An ongoing work is to generalise the results for continuous-time models, MDPs, stochastic games, etc. Moreover, in the current paper, we only consider structurally equivalent perturbations, but it would be interesting to see how to relax this assumption.

References

1. E. Allender, P. Bürgisser, J. Kjeldgaard-Pedersen, and P. B. Miltersen. On the complexity of numerical analysis. *SIAM J. Comput.*, 38(5):1987–2006, 2009.
2. E. Allender and M. Ogihara. Relationships among PL, #L, and the determinant. *ITA*, 30(1):1–21, 1996.
3. C. Baier and J.-P. Katoen. *Principles of Model Checking*. MIT Press, 2008.
4. E. Bartocci, R. Grosu, P. Katsaros, C. R. Ramakrishnan, and S. A. Smolka. Model repair for probabilistic systems. In *TACAS’11*, pages 326–340, 2011.
5. M. Benedikt, R. Lenhardt, and J. Worrell. LTL model checking of Interval Markov Chains. In *TACAS’13*, pages 32–46, Berlin, Heidelberg, 2013.

6. A. Bianco and L. de Alfaro. Model checking of probabalistic and nondeterministic systems. In *FSTTCS*, volume 1026 of *LNCS*, pages 499–513, 1995.
7. J. F. Canny. Some algebraic and geometric computations in PSPACE. In J. Simon, editor, *STOC*, pages 460–467. ACM, 1988.
8. K. Chatterjee. Robustness of structurally equivalent concurrent parity games. In L. Birkedal, editor, *FoSSaCS*, volume 7213 of *LNCS*, pages 270–285. Springer, 2012.
9. K. Chatterjee, K. Sen, and T. A. Henzinger. Model-checking omega-regular properties of Interval Markov Chains. In *FoSSaCS*, volume 4962 of *LNCS*, 302–317, 2008.
10. T. Chen, Y. Feng, D. S. Rosenblum, and G. Su. Perturbation analysis in verification of discrete-time Markov chains. Technical report, Middlesex University London, 2014. Available via www.cs.mdx.ac.uk/staffpages/taoluechen/pub-papers/concur14-full.pdf.
11. T. Chen, E. Hahn, T. Han, M. Kwiatkowska, H. Qu, and L. Zhang. Model repair for Markov decision processes. In *TASE'13*, pages 85–92, July 2013.
12. T. Chen, T. Han, and M. Z. Kwiatkowska. On the complexity of model checking interval-valued discrete time Markov chains. *Inf. Process. Lett.*, 113(7):210–216, 2013.
13. G. E. Cho and C. D. Meyer. Comparison of perturbation bounds for the stationary distribution of a Markov chain. *Linear Algebra Appl*, 335:137–150, 2000.
14. C. Daws. Symbolic and parametric model checking of discrete-time Markov chains. In *ICTAC'04*, pages 280–294, Berlin, Heidelberg, 2005.
15. J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden. Metrics for labelled Markov processes. *Theor. Comput. Sci.*, 318(3):323–354, 2004.
16. R. Donaldson and D. Gilbert. A model checking approach to the parameter estimation of biochemical pathways. In *CMSB'08*, pages 269–287, 2008.
17. K. Etessami and M. Yannakakis. Recursive Markov chains, stochastic grammars, and monotone systems of nonlinear equations. *J. ACM*, 56(1), 2009.
18. A. Filieri, C. Ghezzi, and G. Tamburrelli. Run-time efficient probabilistic model checking. In *ICSE'11*, pages 341–350, New York, NY, USA, 2011.
19. E. Hahn, H. Hermanns, and L. Zhang. Probabilistic reachability for parametric Markov models. *International Journal on Software Tools for Technology Transfer*, 13(1):3–19, 2011.
20. E. M. Hahn, T. Han, and L. Zhang. Synthesis for PCTL in parametric Markov decision processes. In *NASA Formal Methods*, volume 6617 of *LNCS*, pages 146–161, 2011.
21. T. Han, J.-P. Katoen, and A. Mereacre. Approximate parameter synthesis for probabilistic time-bounded reachability. In *RTSS*, pages 173–182. IEEE Computer Society, 2008.
22. M. Kwiatkowska, G. Norman, and D. Parker. PRISM 4.0: Verification of probabilistic real-time systems. In *CAV'11*, pages 585–591. Springer, 2011.
23. J. A. Murdock. *Perturbation: Theory and Method*. JOHN WILEY & SONS, INC., 1991.
24. A. Puggelli, W. Li, A. Sangiovanni-Vincentelli, and S. Seshia. Polynomial-time verification of PCTL properties of MDPs with convex uncertainties. In *CAV*, 527–542, 2013.
25. M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley, New York, 1994.
26. S. Sahni. Computationally related problems. *SIAM Journal on Computing*, 3:262–279, 1974.
27. P. J. Schweitzer. Perturbation theory and finite Markov chains. *Journal of Applied Probability*, 5(2):401–413, 1968.
28. K. Sen, M. Viswanathan, and G. Agha. Model-checking Markov chains in the presence of uncertainties. In *TACAS*, volume 3920 of *LNCS*, pages 394–410. Springer, 2006.
29. E. Solan and N. Vieille. Perturbed Markov chains. *J. Applied Prob*, 40(1):107–122, 2003.
30. G. Su and D. S. Rosenblum. Asymptotic bounds for quantitative verification of perturbed probabilistic systems. In *ICFEM*, volume 8144 of *LNCS*, pages 297–312. Springer, 2013.
31. G. Su and D. S. Rosenblum. Perturbation analysis of stochastic systems with empirical distribution parameters. In *ICSE*, pages 311–321. ACM, 2014.