Approximate Probabilistic Bisimulation for Continuous-Time Markov Chains

Timm Spork[1(✉)], Christel Baier[1], Joost-Pieter Katoen[2], Sascha Klüppelholz[1], and Jakob Piribauer[1,3]

[1] Technische Universität Dresden, Dresden, Germany
{timm.spork,christel.baier,sascha.klueppelholz, jakob.piribauer}@tu-dresden.de
[2] RWTH Aachen University, Aachen, Germany
katoen@cs.rwth-aachen.de
[3] Universität Leipzig, Leipzig, Germany

Abstract. We introduce (ε, δ)-bisimulation, a novel type of approximate probabilistic bisimulation for continuous-time Markov chains. In contrast to related notions, (ε, δ)-bisimulation allows the use of different tolerances for the transition probabilities (ε, additive) and total exit rates (δ, multiplicative) of states. Fundamental properties of the notion, as well as bounds on the absolute difference of time- and reward-bounded reachability probabilities for (ε, δ)-bisimilar states, are established.

Keywords: Continuous-Time Markov Chains · Approximate Probabilistic Bisimulation · Quasi-Lumpability · Time-Bounded Reachability

1 Introduction

Continuous-time Markov chains (CTMCs) are a prominent probabilistic model in various application fields, e.g., reliability engineering, systems biology, modeling of chemical reactions, and performance evaluation. CTMCs are state-based models whose transitions yield a discrete probability distribution over states—as in discrete-time Markov chains (DTMCs)—while the state residence times are governed by exponential distributions. Various model-checking approaches for CTMCs exist [7,11,22,64] and are supported by tools such as PRISM [40] and Storm [34]. CTMC model checking is used to analyze, e.g., stochastic Petri nets [6], fault trees [8,62], biological systems [41,46], and chemical reactions [2,4].

The central issue in CTMC model checking is computing timed reachability probabilities: what is the probability to reach a set of goal states within a given deadline from a given start state? The reliability of a fault tree, or the probability that all molecules have been catalyzed within two days, are instances of this question. Computing timed reachability probabilities reduces to computing transient probabilities in a uniformized CTMC, i.e., a CTMC in which the state

R. Piskac and Z. Rakamarić (Eds.): CAV 2025, LNCS 15932, pp. 56–81, 2025.
https://doi.org/10.1007/978-3-031-98679-6_3

residence times are "normalized" [13]. This method is quite efficient, numerically stable, and scales to CTMCs with millions of states.

In practical applications, however, transition probabilities and state residence time distributions—defined by exit rates—are usually not known exactly. Component failure rates in fault trees are vulnerable to environmental conditions, and reaction rates of molecules are obtained experimentally. This raises the question whether CTMC model-checking results are robust w.r.t. perturbations of their stochastic aspects. The aim of this paper is to investigate to what extent transition probabilities and exit rates in a given CTMC can be altered while ensuring that timed reachability probabilities are preserved up to a small tolerance θ.

To this end we define the novel notion of (ε, δ)-bisimulation on CTMCs, investigate its fundamental properties and derive bounds for timed reachability probabilities. The results yield under which (absolute) ε-tolerance on transition probabilities and (relative) δ-tolerance on exit rates, timed reachability probabilities are close up to θ. This enables, e.g., to determine the maximal tolerances in components' failure rates while ensuring the fault tree's (i.e., overall systems') reliability. Our notion generalizes strong probabilistic bisimulation [19] (also known as lumpability) that preserves timed reachability probabilities exactly.

Let us illustrate the conceptual difference of perturbing exit rates and transition probabilities separately in (ε, δ)-bisimulations compared to existing notions such as τ-quasi-lumpability (also known as near-lumpability) [19,29,30] that consider *transition rates*, i.e., products of exit rates and transition probabilities.

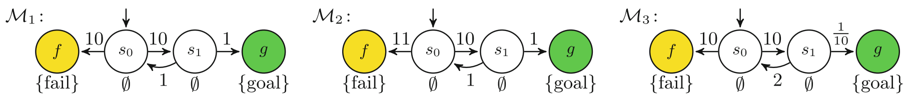

Fig. 1. Three CTMCs where the copies of the states can be related by a τ-quasi-lumpability iff $\tau \geq 1$. The numbers indicate the transition rates between states.

Example 1. Consider the CTMCs in Fig. 1. The value attached to an arrow from state s to state s' describes the *transition rate* $R(s, s')$ between s and s', given as the product of the total exit rate $E(s)$ of s and the probability $P(s, s')$ to move from s to s'. The colors of the states indicate their labels (also written next to the states). Note that the CTMCs $\mathcal{M}_1$ and $\mathcal{M}_2$ behave very similarly—in $\mathcal{M}_2$, the transition to f is only taken with a slightly higher rate, and the probability to reach the goal state g until time t is almost the same in the two models. The CTMCs $\mathcal{M}_1$ and $\mathcal{M}_3$, however, behave very differently—reaching g until t is much less likely in $\mathcal{M}_3$ because of the significantly lower rate $R(s_1, g)$. A τ-quasi-lumpability [19,29,30] is an equivalence relation such that related states have the same label and the total transition rates to move to any other equivalence class from related states differ by at most τ. If we want to relate the states of $\mathcal{M}_1$ with their copies in $\mathcal{M}_2$ with a τ-quasi-lumpability, τ has to be at least 1 due to the change of the transition rate from s_0 to f from 10 to 11. For

$\tau = 1$, however, we can also relate the states of $\mathcal{M}_1$ with their copies in $\mathcal{M}_3$. Hence, τ-quasi-lumpability does not allow us to capture the intuition that $\mathcal{M}_1$ and $\mathcal{M}_2$ behave similarly while $\mathcal{M}_1$ and $\mathcal{M}_3$ behave very differently.

The central idea behind (ϵ, δ)-bisimulations is to decouple the changes of exit rates and transition probabilities of related states. In $\mathcal{M}_1$ and $\mathcal{M}_2$ from Fig. 1, the total exit rate $E(s_0) = \sum_{s'} R(s_0, s')$ of s_0 changes by a factor of 1.05 from 20 to 21, and the probability $P(s_0, f) = \frac{R(s_0,f)}{E(s_0)}$ to transition from s_0 to f is 1/2 and 11/21, respectively. The absolute change of transition probabilities and the relative change of exit rates is small, so the chains behave similarly. For $\mathcal{M}_1$ and $\mathcal{M}_3$, however, $P(s_1, g)$ changes from 1/2 to 1/21, a huge difference. In this way, (ϵ, δ)-bisimulation can express (dis)similarities of CTMCs in a more fine-grained manner than notions that consider the transition rates between states.

Main Contributions. Our main contributions are the following:

- We introduce (ε, δ)-*bisimulation* for CTMCs, a novel type of approximate probabilistic bisimulation that allows an absolute tolerance of ε for the transition probabilities of related states, and a relative tolerance of δ for their total exit rates. We prove that the union of all (ε, δ)-bisimulations, called (ε, δ)-bisimilarity, is an (ε, δ)-bisimulation itself, that it is additive in ε and δ, and that it coincides with strong bisimilarity [15,19] iff $\varepsilon = \delta = 0$. Moreover, we discuss how quasi-lumpability [19,29,30] and (ε, δ)-bisimilarity relate, and show how to "split" (ε, δ)-bisimilarity into $(\varepsilon, 0)$- and $(0, \delta)$-bisimilarity, allowing an individual treatment of the parameters (Sect. 3).
- We derive bounds on the absolute difference of timed reachability probabilities of (ε, δ)-bisimilar states by uniformizing the CTMC and applying a bound from [18] for the difference in finite horizon reachability probabilities of ε-bisimilar states of DTMCs [25]. The bounds are tight if $\delta = 0$ (Sect. 4).
- We analyze the absolute difference of timed reachability probabilities between $(0, \delta)$-bisimilar states in more detail. Using so-called *Erlang CTMCs* [28] we show how to compute them *exactly* if $\varepsilon = 0$. Subsequently, we derive different bounds on the difference: an easy-to-compute one based on an Erlang CTMC of a specific length, and bounds that utilize a spectral decomposition of the transition probability matrix of the given CTMC (Sect. 5).
- We extend (some of) our results to reward-bounded reachability probabilities in CTMCs with nonnegative rewards, by utilizing a method from [12] that allows us to express them as timed reachability probabilities (Sect. 6).

Missing proofs can be found in the full version of the paper [59].

Related Work. In the discrete-time setting of DTMCs a lot of work has been done on approximate probabilistic bisimulations, mainly focusing on ε-bisimulations as introduced by Desharnais *et al.* [18,25]. Other notions include approximate probabilistic bisimulations with precision ε [1,3,27], up-to-(n, ε)-bisimulations

[16,25], or approximate versions of weak- and branching probabilistic bisimulation [5,60]. See, e.g., [60] for a comparison of several different notions.

For CTMCs, an approximate version of lumpability, called τ-*quasi*- or *near-lumpability*, is known [19,29,30]. It is best suited for the analysis of long-run or stationary properties of a chain, but does not provide good guarantees on its *transient* behavior. A more recent notion is *proportional lumpability* [48,49, 55,56]. Intuitively, states s, s' are proportionally lumpable w.r.t. a function κ that maps states to values in $\mathbb{R}_{>0}$ iff there is an equivalence R such that the transition rates of s and s' to any equivalence class of R differ by the constant factors $\kappa(s)$ resp. $\kappa(s')$. Proportional lumpability preserves the *exact* stationary distributions.

Uncertain continuous-time Markov chains (*UCTMC*) [20,21] are CTMCs with time-varying transition rates. At each time $t \geq 0$ and for all states s, s', the transition rate to move from s to s' must be contained in a fixed interval $[m(s, s'), M(s, s')] \subseteq \mathbb{R}_{\geq 0}$. Lumpabilities for UCTMCs are equivalences that require related states to have the same "extremal realizations", defined as the (time-invariant) CTMCs in which all transition rates $R(s, s')$ coincide with $m(s, s')$ resp. $M(s, s')$. These lumpabilities can be computed efficiently and are characterized both via value functions and satisfaction equivalence of formulas in a variant of the continuous stochastic logic CSL [9,13]. However, there does not seem to be a direct connection to our notion of (ε, δ)-bisimulation.

Another related line of research is *perturbation theory* for Markov chains. Given a CTMC $\mathcal{M}$ and a (slightly) perturbed $\mathcal{M}'$, the goal is to bound the difference of (some) performance measures of the models. See, e.g., [43–45,51,58] or the overview article [52]. The bounds presented in the literature are usually on the *total* error between the *stationary* distributions of $\mathcal{M}$ and $\mathcal{M}'$. Furthermore, they are oftentimes obtained under the assumption of *drift conditions* or rely on the *ergodicity* of the chain. This is in contrast to our work, which focuses on the *componentwise* difference in *transient* reachability probabilities up to time t.

2 Preliminaries

Distributions and Vectors. Let $S \neq \emptyset$ be finite. The set of *distributions* on S is $\mathrm{Distr}(S) = \{\mu \colon S \to [0, 1] \mid \sum_{s \in S} \mu(s) = 1\}$. The *support* of $\mu \in \mathrm{Distr}(S)$ is $\mathrm{supp}(\mu) = \{s \in S \mid \mu(s) > 0\}$. Given $A \subseteq S$, we set $\mu(A) = \sum_{a \in A} \mu(a)$. The *Dirac-distribution* 1_s satisfies $1_s(s') = 1$ if $s = s'$ and $1_s(s') = 0$ otherwise. We may associate $\mu \in \mathrm{Distr}(S)$ with a row vector $\boldsymbol{\mu} \in [0, 1]^{1 \times |S|}$ such that $\boldsymbol{\mu}[i] = \mu(i)$ for every $i \in S$. Vectors are written in bold face, and $\boldsymbol{\mu}^\top$ is the transpose of $\boldsymbol{\mu}$.

Relations. Given a relation $R \subseteq S \times S$, the *image* of $A \subseteq S$ under R is $R(A) = \{t \in S \mid \exists\, s \in A \colon (s, t) \in R\}$. A is called R-*closed* if $R(A) \subseteq A$. If R is an equivalence, we write $S/_R$ for the set of R equivalence classes. The R-closed sets of an equivalence R are precisely the (unions of) elements of $S/_R$.

Markov Chains. Fix a countable set AP of *atomic propositions*. A *discrete-time Markov chain* (*DTMC*) is a tuple $\mathcal{D} = (S, P, s_{init}, L)$, with S a finite set of *states*, $P\colon S \to \mathrm{Distr}(S)$ a *transition distribution function*, $s_{init} \in S$ a unique *initial state* and $L\colon S \to 2^{AP}$ a *labeling function*. $P(s, s') = P(s)(s')$ denotes the probability to move from s to s' in one step, and we write $\mathbf{P} \in [0,1]^{|S|\times|S|}$ for the *transition probability matrix* of $\mathcal{D}$ with entries $\mathbf{P}_{s,s'} = P(s, s')$. $\mathrm{Succ}(s) = \{s' \in S \mid P(s, s') > 0\}$ is the set of *successors* of s. Given $s \in S$, let $\mathcal{D}_s$ be the DTMC that is exactly like $\mathcal{D}$, but with initial state s. The *direct sum* $\mathcal{D} \oplus \mathcal{D}'$ of DTMCs $\mathcal{D}, \mathcal{D}'$ is the DTMC obtained from the disjoint union of $\mathcal{D}$ and $\mathcal{D}'$. The initial state of $\mathcal{D} \oplus \mathcal{D}'$ is not relevant for our purposes.

A *continuous-Time Markov Chain* (*CTMC*) is a tuple $\mathcal{M} = (S, P, E, s_{init}, L)$ such that $\mathcal{D}_{\mathcal{M}} = (S, P, s_{init}, L)$ is a DTMC, called the *embedded* or *underlying DTMC* of $\mathcal{M}$, and $E\colon S \to \mathbb{R}_{>0}$ is an *exit rate function*. We use the same notations for CTMCs as for DTMCs, and we let $\mathcal{M}, \mathcal{N}$ range over CTMCs. Note that we do not exclude the possibility of $P(s, s) > 0$, i.e., we allow self-loops in CTMCs. The residence time in a state s of $\mathcal{M}$ is negative exponentially distributed with rate $E(s)$, so the probability to take *any* outgoing transition of s until time $t \geq 0$ (including self-loops) is $1 - e^{-E(s)\cdot t}$, and the probability to take a transition from s to some specific state s' until t is $P(s, s') \cdot (1 - e^{-E(s)\cdot t})$. $\mathbf{Q} \in \mathbb{R}^{|S|\times|S|}$ denotes the (*infinitesimal*) *generator* of $\mathcal{M}$. The entries of $\mathbf{Q}$ are given as $\mathbf{Q}_{i,j} = P(i, j) \cdot E(i)$ if $i \neq j$ and $\mathbf{Q}_{i,i} = -\sum_{j\neq i} \mathbf{Q}_{i,j}$. For more information on CTMCs and DTMCs see, e.g., [14,36,39].

Paths and Probability Measures. Let $\mathcal{D}$ be a DTMC. A sequence $\sigma = s_0 s_1 \ldots \in S^\omega$ is an (*infinite*) *path* of $\mathcal{D}$ if $s_{i+1} \in \mathit{Succ}(s_i)$ for all $i \in \mathbb{N}$. $\sigma[i] = s_i$ is the state at position i of σ, and $\mathit{trace}(\sigma) = L(s_0)L(s_1)\ldots \in (2^{AP})^\omega$ is the *trace* of σ. The set of infinite paths is $\mathrm{Paths}(\mathcal{D})$. *Finite* paths $\sigma = s_0 s_1 \ldots s_k \in S^{k+1}$ and their traces are defined similarly. $\mathrm{Paths}^*(\mathcal{D})$ is the set of finite paths of $\mathcal{D}$.

Let $s \in S$. We consider the standard probability measure $\mathrm{Pr}_s^{\mathcal{D}}$ on subsets of $\mathrm{Paths}(\mathcal{D})$, defined via *cylinder sets* $\mathit{Cyl}(\rho) = \{\sigma \in \mathrm{Paths}(\mathcal{D}) \mid \rho \text{ is a prefix of } \sigma\}$ for $\rho \in \mathrm{Paths}^*(\mathcal{D})$. See, e.g., [14] for details. We write $\mathrm{Pr}^{\mathcal{D}}$ for $\mathrm{Pr}_{s_{init}}^{\mathcal{D}}$ and drop the superscript if $\mathcal{D}$ is clear from the context.

Now, let $\mathcal{M}$ be a CTMC. $\sigma = s_0 t_0 s_1 t_1 \ldots \in (S \cdot \mathbb{R}_{>0})^\omega$, where $\cdot$ denotes concatenation, is an (*infinite*) *timed path* of $\mathcal{M}$ if $s_{i+1} \in \mathrm{Succ}(s_i)$ for all $i \in \mathbb{N}$. $\sigma[i]$ and $\mathit{trace}(\sigma)$ are defined as for DTMCs. A *finite* timed path is a finite prefix of a timed path that ends in a state. We write $\mathrm{Paths}(\mathcal{M})$ (resp. $\mathrm{Paths}^*(\mathcal{M})$) for the set of infinite (resp. finite) timed paths of $\mathcal{M}$. The value $\sigma_i = t_i$ describes the residence time in state s_i along σ. Given some $t \in \mathbb{R}_{\geq 0}$, $\sigma@t$ is the state of σ at time t, i.e., $\sigma@t = \sigma[k]$ for k the smallest index such that $t \leq \sum_{i=0}^{k} t_i$ [13].

Let $k \in \mathbb{N}$, $s, s_0, \ldots, s_{k+1} \in S$ and $I_0, \ldots, I_k$ nonempty intervals in $\mathbb{R}_{\geq 0}$. We consider the standard probability measure $\mathrm{Pr}_s^{\mathcal{M}}$ on subsets of $\mathrm{Paths}(\mathcal{M})$, defined via *timed cylinder sets* $\mathit{Cyl}(s_0 I_0 \ldots s_k I_k s_{k+1}) = \{\sigma \in \mathrm{Paths}(\mathcal{M}) \mid \sigma[i] = s_i \text{ for } i \leq k+1 \text{ and} \sigma_j \in I_j \text{ for} j \leq k\}$. See, e.g., [13] for details. We use the same abbreviations as for DTMCs regarding $\mathrm{Pr}_s^{\mathcal{M}}$. For a finite untimed path $\pi = s_0 s_1 \ldots s_n$, let $\mathrm{Pr}^{\mathcal{M}}(\pi)$ denote the probability of $\mathcal{M}$ to follow any timed path with a state sequence prefixed by π.

Given $G \subseteq S$ and $t \geq 0$, $\Pr_s^{\mathcal{M}}(\Diamond^{\leq t} G)$ is the probability to reach in $\mathcal{M}$ from $s \in S$ a state in G after time at most t. For a DTMC $\mathcal{D}$, $\Pr_s^{\mathcal{D}}(\Diamond^{\leq k} G)$ is the probability to reach in $\mathcal{D}$ from s a state in G after at most $k \in \mathbb{N}$ steps.

Transient Probabilities. Let $n \in \mathbb{N}$. The probability of a DTMC $\mathcal{D}$ to be in state s' after n steps when starting in state s is $\pi_n^{\mathcal{D}}(s, s') = \mathbf{P}^n_{s,s'}$.

For CTMC $\mathcal{M}$, the values $\pi_t^{\mathcal{M}}(s)$ of the *transient probability distribution* $\pi_t^{\mathcal{M}} \in \text{Distr}(S)$ describe the probabilities of $\mathcal{M}$ to be in s at time t. Let $\boldsymbol{\pi}_t^{\mathcal{M}} \in [0,1]^{1\times|S|}$ denote the vector representation of $\pi_t^{\mathcal{M}}$, and let $\pi_t^{\mathcal{M}}(s, s') = \pi_t^{\mathcal{M}_s}(s')$ be the probability to be in s' after time t when starting in s. $\boldsymbol{\pi}_t^{\mathcal{M}}$ can be computed by solving the following system of forward Kolmogorov differential equations [38]:

$$\frac{d}{dt}\boldsymbol{\pi}_t^{\mathcal{M}} = \boldsymbol{\pi}_t^{\mathcal{M}} \cdot \mathbf{Q} \qquad \text{given } \boldsymbol{\pi}_0^{\mathcal{M}} = \mathbf{1}_{s_{init}}.$$

Let $q \geq \max_{s \in S} E(s)$. The DTMC $\text{unif}(\mathcal{M}, q) = (S, \overline{P}, s_{init}, L)$ with $\overline{P}(s, s') = \frac{P(s,s')\cdot E(s)}{q}$ if $s \neq s'$, and $\overline{P}(s,s) = 1 + \frac{P(s,s)\cdot E(s)}{q} - \frac{E(s)}{q}$ is the *uniformization* of $\mathcal{M}$ w.r.t. q. $\boldsymbol{\pi}_t^{\mathcal{M}}$ can also be computed by the *uniformization method* via [32,35]

$$\boldsymbol{\pi}_t^{\mathcal{M}} = \sum\nolimits_{k=0}^{\infty} e^{-q\cdot t} \cdot \frac{(q\cdot t)^k}{k!} \cdot \mathbf{1}_{s_{init}} \cdot \overline{\mathbf{P}}^k.$$

We omit the superscript from $\pi_t^{\mathcal{M}}$ (resp. $\pi_n^{\mathcal{D}}$) if no confusion can arise.

Bisimulation. A *(probabilistic) bisimulation* $R \subseteq S \times S$ for DTMC $\mathcal{D}$ is an equivalence such that for all $(s, s') \in R$ it holds that $L(s) = L(s')$ and $P(s, C) = P(s', C)$ for every $C \in {}^{S}/_{R}$ [42]. For CTMC $\mathcal{M}$, an equivalence R is a *strong (probabilistic) bisimulation* if it is a bisimulation on $\mathcal{D}_{\mathcal{M}}$, and additionally satisfies $E(s) = E(s')$ for all $(s, s') \in R$ [15,19]. States s and s' of $\mathcal{M}$ are *strongly (probabilistic) bisimilar*, written $s \sim s'$, if there is a strong bisimulation R on $\mathcal{M}$ with $(s, s') \in R$. CTMCs $\mathcal{M}$ and $\mathcal{N}$ are strongly (probabilistic) bisimilar, written $\mathcal{M} \sim \mathcal{N}$, if $s_{init}^{\mathcal{M}} \sim s_{init}^{\mathcal{N}}$ in $\mathcal{M} \oplus \mathcal{N}$.

3 (ε, δ)-Bisimulation on CTMCs

If not specified otherwise, we always assume that $\varepsilon, \delta \geq 0$. In this section, after a short recap on ε-bisimulation for DTMCs (Sect. 3.1), we formally define (ε, δ)-bisimulation for CTMCs (Sect. 3.2) and establish fundamental properties of this notion (Sect. 3.3).

3.1 ε-Bisimulation in the Discrete-Time Setting

In the discrete-time setting of DTMCs, the most well-known and well-studied [18,25,37] notion of approximate probabilistic bisimulation are *ε-bisimulations.* They were introduced by Desharnais *et al.* [25] for *labeled Markov processes* [23,24] and have since been adjusted to other models like DTMC [18,37].

Definition 2 **([18,25]).** *Let $\mathcal{D}$ be a DTMC. A reflexive*[1] *and symmetric relation $R \subseteq S \times S$ is an ε-*bisimulation *if for all $(s,t) \in R$ and any $A \subseteq S$*

$$\text{(i) } L(s) = L(t) \quad \textit{and} \quad \text{(ii) } P(s, A) \leq P(t, R(A)) + \varepsilon.$$

*States s and t are ε-*bisimilar*, denoted $s \sim_\varepsilon t$, if there is an ε-bisimulation R with $(s,t) \in R$. DTMCs $\mathcal{D}_1$ and $\mathcal{D}_2$ are ε-bisimilar, denoted $\mathcal{D}_1 \sim_\varepsilon \mathcal{D}_2$, if $s_{init}^{\mathcal{D}_1} \sim_\varepsilon s_{init}^{\mathcal{D}_2}$ in $\mathcal{D}_1 \oplus \mathcal{D}_2$.*

The tolerance ε describes by how much the transition probabilities of related states may differ. If $\varepsilon \approx 1$, many states can be related, but their behavior might be significantly different, whereas an ε close to 0 allows us to only relate states with almost the same behavior, decreasing the number of relatable states.

The satisfaction of condition (ii) of Definition 2 can be characterized by the existence of suitable *weight functions* that describe how to split the successor probabilities of related states such that the condition holds. This approach is used in, e.g., [33,37,61]. We will later use the following characterization (where we slightly abuse notation and write $\Delta_{s,t}(s', t')$ instead of $\Delta_{s,t}(s')(t')$).

Lemma 3 **([60]).** *A reflexive and symmetric relation $R \subseteq S \times S$ that only relates states with the same label is an ε-bisimulation iff for all $(s,t) \in R$ there is a map $\Delta_{s,t} \colon \mathrm{Succ}(s) \to \mathrm{Distr}(\mathrm{Succ}(t))$ such that*

1. *for all $t' \in \mathrm{Succ}(t)$ we have $P(t,t') = \sum_{s' \in \mathrm{Succ}(s)} P(s,s') \cdot \Delta_{s,t}(s',t')$, and*
2. *$\sum_{s' \in \mathrm{Succ}(s)} P(s,s') \cdot \Delta_{s,t}(s', R(s') \cap \mathrm{Succ}(t)) \geq 1 - \varepsilon$.*

Furthermore, step-bounded reachability probabilities $\Pr_*(\Diamond^{\leq n} g)$ of a goal state g for ε-bisimilar states s, s' are related as follows.

Theorem 4 **([18]).** *Let $\mathcal{D}$ be a DTMC, $k \in \mathbb{N}$ and $s \sim_\varepsilon s'$. Then*

$$|\Pr_s(\Diamond^{\leq k} g) - \Pr_{s'}(\Diamond^{\leq k} g)| \leq 1 - (1 - \varepsilon)^k.$$

3.2 Definition of (ε, δ)-Bisimulation

In contrast to the discrete-time case, the behavior of a CTMC $\mathcal{M}$ is not only influenced by the transition probabilities, but also by the exit rates of the states. Thus, it is natural to consider two different tolerance values ε and δ for the probabilities resp. the exit rates. As the values of $P(s, \cdot)$ are in $[0, 1]$ for any $s \in S$, similar to the discrete-time case an additive tolerance ε is suitable.

The rates $E(s)$, however, can take on any (finite) positive value. This makes the use of *additive* tolerances δ for the rates difficult. Consider, for example, a CTMC with states s, s' such that $E(s) = 1$ and $E(s') = 100$, and assume that the exit rates of related states are allowed to differ by at most 10%. For s, an additive δ would need a value of 0.1 (i.e., all states with a rate in $[0.9, 1.1]$ are

[1] In contrast to [18,25] we require reflexivity of ε-bisimulations. This assumption is rather natural (a state should always simulate itself) and does not affect $\sim_\varepsilon$.

considered to behave almost the same as s), while $\delta = 10$ would be necessary for s'. The former allows virtually no tolerance for $E(s')$ (only 0.001% instead of the desired 10%), while the latter allows a tolerance of up to 1000% for $E(s)$.

To circumvent this issue we propose the use of a *multiplicative* tolerance δ for the exit rates of related states, so that the tolerable difference is relative to the actual rates. For the rest of the paper, let $\ln(\cdot)$ denote the natural logarithm.

Definition 5 *Let $\mathcal{M}$ be a CTMC and let $R \subseteq S \times S$ be a reflexive and symmetric relation. R is an (ε, δ)-bisimulation if for all $(s, s') \in R$ it holds that*

1. $L(s) = L(s')$ (labeling condition)
2. $|\ln(E(s)) - \ln(E(s'))| \leq \delta$ (δ-condition)
3. *for all* $A \subseteq S$: $P(s, A) \leq P(s', R(A)) + \varepsilon$. ($\varepsilon$-condition)

States $s, s' \in S$ are (ε, δ)-bisimilar*, written $s \sim_{\varepsilon,\delta} s'$, if $(s, s') \in R$ for an (ε, δ)-bisimulation R. CTMCs $\mathcal{M}_1$ and $\mathcal{M}_2$ are* (ε, δ)-bisimilar*, written $\mathcal{M}_1 \sim_{\varepsilon,\delta} \mathcal{M}_2$, if $s_{init}^{\mathcal{M}_1} \sim_{\varepsilon,\delta} s_{init}^{\mathcal{M}_2}$ in $\mathcal{M}_1 \oplus \mathcal{M}_2$.*

Any (ε, δ)-bisimulation on $\mathcal{M}$ induces an ε-bisimulation on the embedded DTMC $\mathcal{D}_{\mathcal{M}}$, i.e., $s \sim_{\varepsilon,\delta}^{\mathcal{M}} t$ implies $s \sim_{\varepsilon}^{\mathcal{D}_{\mathcal{M}}} t$. Hence, there are weight functions $\Delta_{s,t}$ as in Lemma 3 (w.r.t. the probabilities of $\mathcal{M}$) for (ε, δ)-bisimilar states s, t.

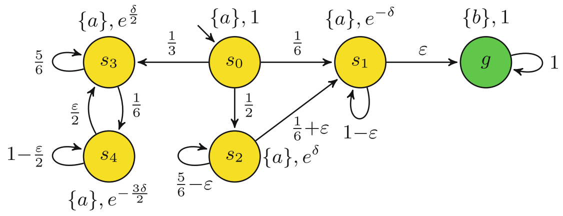

Fig. 2. The CTMC used in Example 6.

Example 6. Let $\varepsilon < \frac{1}{2}$ and $\delta > 0$. In Fig. 2, $\sim_{\varepsilon,\delta}$ is the reflexive and symmetric closure of $\{(s_0, s_2), (s_0, s_3), (s_1, s_4), (s_2, s_3)\}$. States s_1 and s_2 are not (ε, δ)-bisimilar as they violate the δ-condition: $|\ln(E(s_1)) - \ln(E(s_2))| = 2\delta$. Moreover, $s_0 \nsim_{\varepsilon,\delta} s_1$ since, although the pair of states satisfies the δ-condition, the ε-condition is violated: $P(s_0, \{s_2\}) = \frac{1}{2} > \varepsilon = P(s_1, \sim_{\varepsilon,\delta}(\{s_2\})) + \varepsilon$.

Note that the steady state distributions of (ε, δ)-bisimilar states can differ significantly: starting from s_0 the probability to reach g (and stay there forever) is 2/3, from s_2 it is 1 and from s_3 it is 0, even though s_0, s_2, and s_3 are pairwise (ε, δ)-bisimilar.

3.3 Fundamental Properties of $\sim_{\varepsilon,\delta}$

We now establish some fundamental properties of $\sim_{\varepsilon,\delta}$ that are generally desirable for approximate probabilistic bisimulations [25, 60].

Theorem 7. *Let $\mathcal{M}$ be a CTMC and $s, s', s'' \in S$.*

1. *$\sim_{\varepsilon,\delta}$ is the largest (ε, δ)-bisimulation on $\mathcal{M}$.*
2. *If $s \sim_{\varepsilon_1,\delta_1} s'$ and $s' \sim_{\varepsilon_2,\delta_2} s''$ then $s \sim_{\varepsilon_1+\varepsilon_2,\delta_1+\delta_2} s''$.*
3. *$s \sim s'$ iff $s \sim_{0,0} s'$.*

Theorem 7 shows that $\sim_{\varepsilon,\delta}$ is itself an (ε, δ)-bisimulation, and that it is the largest such relation (item 1). Furthermore $\sim_{\varepsilon,\delta}$ is additive in both ε and δ (item 2) and coincides with strong bisimilarity $\sim$ iff $\varepsilon = \delta = 0$ (item 3).

By slightly adjusting a procedure proposed in [25] that constructs $\sim_\varepsilon$ for DTMCs by solving flow networks á la [10,57], $\sim_{\varepsilon,\delta}$ can be computed in time polynomial in the number of states of $\mathcal{M}$.

Corollary 8. *For a given CTMC $\mathcal{M}$, $\sim_{\varepsilon,\delta}$ can be computed in time $\mathcal{O}(|S|^7)$.*

Furthermore, (ε, δ)-bisimilarity in $\mathcal{M}$ implies τ-bisimilarity (for a $\tau \geq 0$ that depends on ε and δ) in the uniformization $\mathrm{unif}(\mathcal{M}, q)$ of $\mathcal{M}$ for $q \geq \max_{s \in S} E(s)$.

Lemma 9. *$s \sim^{\mathcal{M}}_{\varepsilon,\delta} s'$ implies $s \sim^{\mathrm{unif}(\mathcal{M},q)}_{\tau} s'$, where $\tau = e^\delta \cdot (1 + \varepsilon) - 1$.*

A notion from the literature related to (ε, δ)-bisimulations is that of τ-quasi-lumpability, which is also known as τ- or near-lumpability [19,29,30]. It is defined for partitions $\Omega = \{\Omega_1, \ldots, \Omega_m\}$ of the state space S of $\mathcal{M}$. More precisely, Ω is a τ-quasi-lumpability for some $\tau \geq 0$ if for all $1 \leq i, j \leq m$ and all $s, s' \in \Omega_i$ it holds that $|P(s, \Omega_j) \cdot E(s) - P(s', \Omega_j) \cdot E(s')| \leq \tau$ [48]. The partitions of the state space induced by *transitive* (ε, δ)-bisimulations are quasi-lumpabilities.

Proposition 10. *Let R be a transitive (ε, δ)-bisimulation and $q \geq \max_{s \in S} E(s)$. The partition induced by R, i.e., the set $S/_R$, is a $q \cdot (e^\delta \cdot (1+\varepsilon) - 1)$-lumpability.*

In [59] we construct, for given $\varepsilon \in (0, 1), \delta > 0$ and $\tau > 0$, a CTMC $\mathcal{M}(\varepsilon, \delta, \tau)$ with states s, s' such that $\{s, s'\}$ is a block of a τ-quasi lumpability on $\mathcal{M}(\varepsilon, \delta, \tau)$ but the pair (s, s') does neither satisfy the ε- nor the δ-condition.

We finish the section by showing how to "split" (ε, δ)-bisimilarity into $(\varepsilon, 0)$- and $(0, \delta)$-bisimilarity. More precisely, given $\mathcal{M} \sim_{\varepsilon,\delta} \mathcal{N}$ we construct CTMCs $\mathcal{M}', \mathcal{N}'$ with the same graph structure and such that $\mathcal{M} \sim_{\varepsilon,0} \mathcal{M}' \sim_{0,\delta} \mathcal{N}' \sim \mathcal{N}$. This decomposition makes it possible to treat the two parameters individually and, together with the additivity of $\sim_{\varepsilon,\delta}$, allows us to extend results shown for $(\varepsilon, 0)$- and $(0, \delta)$-bisimilar states or chains to (ε, δ)-bisimilar ones.

Theorem 11. *Let $\mathcal{M} \sim_{\varepsilon,\delta} \mathcal{N}$. Then there are CTMCs $\mathcal{M}'$ and $\mathcal{N}'$ with the same graph structure and such that $\mathcal{M} \sim_{\varepsilon,0} \mathcal{M}' \sim_{0,\delta} \mathcal{N}' \sim \mathcal{N}$.*

Proof sketch. We describe how to construct $\mathcal{M}'$ and $\mathcal{N}'$. The models share the state space $S' = S^{\mathcal{M}} \times S^{\mathcal{N}}$, with initial state $(s^{\mathcal{M}}_{init}, s^{\mathcal{N}}_{init})$. The label function L' of both models is $L'((s,t)) = L^{\mathcal{N}}(t)$, and the exit rate functions are defined via

$$E^{\mathcal{M}'}((s,t)) = \begin{cases} E^{\mathcal{M}}(s), & \text{if } s \sim_{\varepsilon,\delta} t \\ E^{\mathcal{N}}(t), & \text{if} s \not\sim_{\varepsilon,\delta} t \end{cases} \qquad \text{and} \qquad E^{\mathcal{N}'}((s,t)) = E^{\mathcal{N}}(t).$$

Furthermore, the common transition probability function of $\mathcal{M}'$ and $\mathcal{N}'$ is

$$P'((s,t),(s',t')) = \begin{cases} \Delta_{s,t}(s',t')P^{\mathcal{M}}(s,s'), & \text{if } s \sim_{\varepsilon,\delta} t, s' \in \mathrm{Succ}(s), t' \in \mathrm{Succ}(t) \\ P^{\mathcal{M}}(s,s')P^{\mathcal{N}}(t,t'), & \text{if } s \nsim_{\varepsilon,\delta} t \\ 0, & \text{otherwise} \end{cases}$$

where, for a given pair of (ε,δ)-bisimilar states $(s,t) \in S^{\mathcal{M}} \times S^{\mathcal{N}}$, the function $\Delta_{s,t}\colon \mathrm{Succ}(s) \to \mathrm{Distr}(\mathrm{Succ}(t))$ is a weight function as in Lemma 3, which exists for all such (s,t) since $s \sim_{\varepsilon} t$ in the underlying DTMCs of $\mathcal{M}$ and $\mathcal{N}$. The proof proceeds by showing the desired relations between the models. □

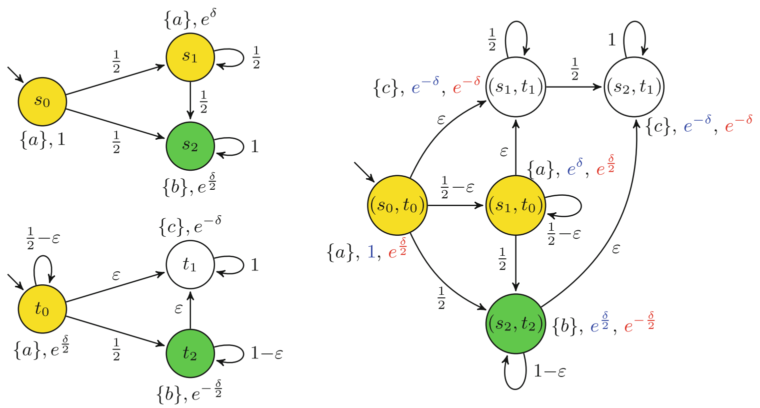

Fig. 3. The CTMCs used in Example 12.

Example 12. We illustrate the construction of Theorem 11 in Fig. 3. The CTMCs $\mathcal{M}$ (top left) and $\mathcal{N}$ (bottom left) are (ε,δ)-bisimilar, as all states with the same label are pairwise (ε,δ)-bisimilar in $\mathcal{M} \oplus \mathcal{N}$. The chains $\mathcal{M}'$ and $\mathcal{N}'$, which only differ in their exit rate functions $E^{\mathcal{M}'}$ (first, in blue) and $E^{\mathcal{N}'}$ (second, in red) are on the right of the figure. To construct $\mathcal{M}'$ and $\mathcal{N}'$ it is necessary to compute weight functions $\Delta_{s,t}$ for all $s \sim_{\varepsilon,\delta} t$. For example, a suitable choice for (s_0,t_0) is $\Delta_{s_0,t_0}(s_1,t_0) = 1-2\cdot\varepsilon$, $\Delta_{s_0,t_0}(s_1,t_1) = 2\cdot\varepsilon$, $\Delta_{s_0,t_0}(s_2,t_2) = 1$ and $\Delta_{s_0,t_0}(\cdot,\cdot) = 0$ otherwise. It is easy to prove that $\mathcal{M} \sim_{\varepsilon,0} \mathcal{M}' \sim_{0,\delta} \mathcal{N}' \sim \mathcal{N}$.

4 Bounding Timed Reachability Probabilities

Let $\mathcal{M}$ be a CTMC and $t \in \mathbb{R}_{\geq 0}$. We are interested in bounds for the absolute difference of the probabilities of (ε,δ)-bisimilar states s and s' to reach a unique *goal state* g until the deadline t, i.e., in bounds for $|\mathrm{Pr}_s(\Diamond^{\leq t} g) - \mathrm{Pr}_{s'}(\Diamond^{\leq t} g)|$.

For the rest of the paper we therefore assume that $\mathcal{M}$ has such a unique goal state g, which is w.l.o.g. [13] absorbing and uniquely labeled. If there are multiple, potentially non-absorbing goal states, these states can be collapsed to a single absorbing goal state without affecting the time-bounded reachability probabilities. The same pre-processing is also applied in the computation of satisfaction probabilities for time-bounded until-formulas of the *continuous stochastic logic CSL* [9,13]. Similarly, we assume that all states from which g is not reachable are collapsed to a single, uniquely labeled absorbing fail state. Hence, there are at most two absorbing states in $\mathcal{M}$ which are not (ε, δ)-bisimilar to any other state (because of the unique labels) and are eventually reached almost surely. Further, let $q = \max_{s \in S} E(s)$ be the smallest possible uniformization rate of $\mathcal{M}$.

By considering the DTMC $\text{unif}(\mathcal{M}, q)$ and applying the bound of Theorem 4 [18], which is possible because of the preservation of approximate bisimilarity in the uniformization (see Lemma 9), we derive an upper bound on the absolute difference of timed reachability probabilities of (ε, δ)-bisimilar states.

Proposition 13. *For* $s \sim_{\varepsilon,\delta} s'$*:* $|\Pr_s(\Diamond^{\leq t} g) - \Pr_{s'}(\Diamond^{\leq t} g)| \leq 1 - e^{-qt(e^{\delta}(1+\varepsilon)-1)}$.

Considering $\varepsilon = 0$ or $\delta = 0$ in Proposition 13 yields the following corollary.

Corollary 14. *Let* $s \sim_{\varepsilon,\delta} s'$*.*

1. *If* $\delta = 0$ *then* $|\Pr_s(\Diamond^{\leq t} g) - \Pr_{s'}(\Diamond^{\leq t} g)| \leq 1 - e^{-qt\varepsilon}$.
2. *If* $\varepsilon = 0$ *then* $|\Pr_s(\Diamond^{\leq t} g) - \Pr_{s'}(\Diamond^{\leq t} g)| \leq 1 - e^{-qt(e^{\delta}-1)}$.
3. *If* $\varepsilon = \delta = 0$ *then* $\Pr_s(\Diamond^{\leq t} g) = \Pr_{s'}(\Diamond^{\leq t} g)$.

The third result of Corollary 14 also directly follows from $s \sim_{0,0} s'$ iff $s \sim s'$, which was proved in the third item of Theorem 7, and the fact that $\sim$ exactly preserves transient (reachability) probabilities [13,19].

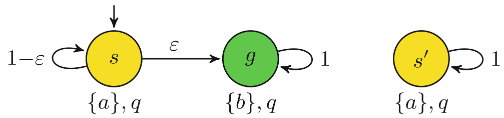

Fig. 4. A CTMC with $s \sim_{\varepsilon,0} s'$ and $|\Pr_s(\Diamond^{\leq t} g) - \Pr_{s'}(\Diamond^{\leq t} g)| = 1 - e^{-qt\varepsilon}$.

We now show that the bound of Proposition 13 is tight if $\delta = 0$.

Example 15. ([18]*).* In Fig. 4, $s \sim_{\varepsilon,0} s'$, $\Pr_{s'}(\Diamond^{\leq t} g) = 0$ and $\Pr_s(\Diamond^{\leq t} g) = 1 - e^{-qt\varepsilon}$, so $|\Pr_s(\Diamond^{\leq t} g) - \Pr_{s'}(\Diamond^{\leq t} g)| = |1 - e^{-qt\varepsilon} - 0| = 1 - e^{-qt\varepsilon}$ for all t.

A disadvantage of the bounds in Proposition 13 (and Corollary 14) is that they converge to 1 exponentially fast for $t \to \infty$, while the maximal total difference in timed reachability probabilities is usually much smaller for (ε, δ)-bisimilar states. Intuitively, the bounds converge to 1 because the bound of Theorem 4 used in their derivation always assumes the maximal possible error. However it does not consider, e.g., that for growing t some of the probability mass already reached the goal state and can thus not contribute to the total error anymore.

This disadvantage is particularly observable for $(0, \delta)$-bisimilar states, as in this case the actual error converges to 0 for $t \to \infty$ (see also Figs. 6 and 7 in Sect. 5). In Sect. 5, we derive explicit formulas and better bounds for the absolute difference in timed reachability probabilities of $(0, \delta)$-bisimilar states.

For now we consider the problem of computing, for given $\theta \in [0, 1)$ and time $t > 0$, values for ε and δ that guarantee the absolute difference in reachability probabilities of g until t for (ε, δ)-bisimilar states to be $\leq \theta$. By using the bound of Proposition 13 and solving $1 - e^{-qt(e^{\delta}(1+\varepsilon)-1)} \leq \theta$ w.r.t. ε and δ we obtain:

Theorem 16. *Let $\theta \in [0, 1)$, $t > 0$, and $q = \max_{p \in S} E(p)$. Then, for all ε, δ with $\varepsilon \in \left[0, \frac{1}{e^{\delta}} \cdot \left(\frac{q \cdot t - \ln(1-\theta)}{q \cdot t}\right) - 1\right]$ and $\delta \in \left[0, \ln\left(\frac{q \cdot t - \ln(1-\theta)}{(\varepsilon+1) \cdot q \cdot t}\right)\right]$, it holds that*

$$s \sim_{\varepsilon,\delta} s' \ \textit{ implies } \ |\Pr_s(\Diamond^{\leq t} g) - \Pr_{s'}(\Diamond^{\leq t} g)| \leq \theta.$$

As the bound of Proposition 13 is tight if $\delta = 0$ (see Example 15), the range of admissible values for ε in Corollary 16 is tight in this case as well, in the sense that for a given $\theta \in [0, 1)$ there is a CTMC $\mathcal{M}$ with states $s \sim_{\varepsilon,0} s'$ such that $|\Pr_s(\Diamond^{\leq t} g) - \Pr_{s'}(\Diamond^{\leq t} g)| \leq \theta$ iff $\varepsilon \leq -\frac{\ln(1-\theta)}{q \cdot t}$.

5 The Effect of Changing Rates Only

The goal of this section is to obtain more refined bounds for the absolute difference of timed reachability probabilities of $(0, \delta)$-bisimilar CTMCs. This setting is relevant in different scenarios like, e.g., queuing theory. Consider a CTMC $\mathcal{M}$ modeling a single server queue with a buffer of size n, in which tasks arrive with a constant rate τ and are served with a constant rate μ. If the buffer is full and a new task arrives, it has to be dropped and an absorbing fail state f is entered. The timed reachability probability $\Pr^{\mathcal{M}}(\Diamond^{\leq t} f)$ describes the probability of failure until t. Now let $\mathcal{M}'$ be like $\mathcal{M}$, but with an increased service rate $\mu' = \mu \cdot c$ for some $c > 1$. Then $\mathcal{M} \sim_{0,\delta} \mathcal{M}'$ for $\delta \geq \ln(c)$ and $|\Pr^{\mathcal{M}}(\Diamond^{\leq t} f) - \Pr^{\mathcal{M}'}(\Diamond^{\leq t} f)|$ describes how much higher the probability of failure until t is when using the slower server instead of the faster one, making it easy to compare their performance.

We first show how to compute the absolute difference in timed reachability probabilities of $(0, \delta)$-bisimilar CTMCs $\mathcal{M}$ and $\mathcal{M}'$ exactly, without relying on the uniformization method [32,35] or the Kolmogorov forward equations [38]. We then derive two bounds for the difference: an easy-to-compute constant value (that depends on t and δ), and a bound that yields better results if $t \to \infty$.

Remark 17. Throughout this section we assume that the CTMCs are uniform, i.e., that all states have the same exit rates. For *transitive* $(0, \delta)$-bisimulations R on non-uniformized CTMCs, the following procedure allows us to construct uniformized CTMCs for the analysis of time-bounded reachability probabilities: Let $\mathcal{M}$ and $\mathcal{N}$ be related by such an R. In all equivalence classes C of R, we change the exit rates in $\mathcal{M}$ to the minimal exit rate $E^{\mathcal{M}}_{\min}(C)$ present in C and all rates in $\mathcal{N}$ to $E^{\mathcal{M}}_{\min}(C) \cdot e^{\delta}$. As the transition probabilities of the chains are not affected by this change, and the rates of related states differ by a factor of at most e^{δ}, R is also a $(0, \delta)$-bisimulation on the transformed chains. Afterwards, $\mathcal{M}$ and $\mathcal{N}$ can be uniformized with rates q and $q \cdot e^{\delta}$, respectively, for a suitable q. This again preserves the fact that R is a $(0, \delta)$-bisimulation and results in uniform CTMCs. During the procedure, the difference $|\Pr^{\mathcal{M}}(\Diamond^{\leq t} g) - \Pr^{\mathcal{N}}(\Diamond^{\leq t} g)|$ in the original CTMCs is at most as big as in the resulting uniformized CTMCs. For more details on the transformation, see [59].

Our formula for $|\Pr^{\mathcal{M}}(\Diamond^{\leq t} g) - \Pr^{\mathcal{M}'}(\Diamond^{\leq t} g)|$ depends on the corresponding differences in CTMCs of a specific form, which we call *Erlang CTMCs* [28].

Definition 18. *Let $n \in \mathbb{N}$. The (n-)*Erlang CTMC *$\mathcal{E}_n$ is a CTMC with n non-goal states $s_0, \ldots, s_{n-1}$ and a unique goal state $s_n = g$, such that $L(s_i) = L(s_j)$ for $0 \leq i, j < n$, $L(g) \neq L(s_i)$ for any $i < n$, $E(s_i) = 1$ for $0 \leq i \leq n$ and $P(s_i, s_{i+1}) = 1$ for $0 \leq i < n$, as well as $P(g, g) = 1$.*

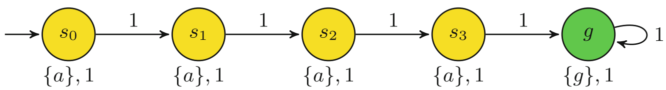

Fig. 5. The Erlang CTMC $\mathcal{E}_4$.

$\mathcal{E}_4$ is illustrated in Fig. 5. Given $\mathcal{M}$ and $c \in \mathbb{R}_{>0}$, let $c \cdot \mathcal{M}$ be the CTMC obtained from $\mathcal{M}$ by multiplying all rates with c. Then $\mathcal{M} \sim_{0,\delta} c{\cdot}\mathcal{M}$ for $\delta \geq \ln(c)$. For an absorbing, uniform CTMC $\mathcal{M}$ it is clear that an acceleration $c \cdot \mathcal{M}$ with $c = e^{\delta} \geq 1$ or a deceleration $1/c \cdot \mathcal{M}$ induces the maximal possible difference in timed reachability probabilities compared to $\mathcal{M}$ among all $\mathcal{M}'$ with $\mathcal{M} \sim_{0,\delta} \mathcal{M}'$. In the sequel, results are formulated for the comparison of $\mathcal{M}$ with $c \cdot \mathcal{M}$. By switching the roles of $\mathcal{M}$ and $c \cdot \mathcal{M}$, we obtain symmetric results for the deceleration.

Proposition 19. *Let $n \in \mathbb{N}$, $t \geq 0$, and $\mathcal{E}'_n = c \cdot \mathcal{E}_n$ for some $c \geq 1$. Then*

$$\mathrm{Diff}_t(\mathcal{E}_n) := |\Pr^{\mathcal{E}_n}(\Diamond^{\leq t} g) - \Pr^{\mathcal{E}'_n}(\Diamond^{\leq t} g)| = \sum_{k=0}^{n-1} \frac{t^k}{k!} \cdot \left(e^{-t} - c^k e^{-ct}\right).$$

If $c > 1$, $\mathrm{Diff}_t(\mathcal{E}_n)$ has a local maximum at $t = \frac{n \cdot \ln(c)}{c-1}$ that is global on $[0, \infty)$.

Remark 20. We could have also allowed arbitrary (but uniform) rates for $\mathcal{E}_n$ in Definition 18, i.e., $E(s) = r$ for all $s \in S$ and some $r > 0$. In this case,

$$\mathrm{Diff}_t(\mathcal{E}_n) = \sum_{k=0}^{n-1} \frac{(r \cdot t)^k}{k!} \cdot \left(e^{-r \cdot t} - c^k e^{-c \cdot r \cdot t}\right),$$

with a (local) maximum at $t^* = \frac{n \cdot \ln(c)}{r \cdot (c-1)}$ for $c > 1$. Note that $\mathrm{Diff}_{t^*}(\mathcal{E}_n)$ does not depend on r. Moreover, the assumption that $r = 1$ can be made w.l.o.g., as in the case of $r \neq 1$ the time-bounded reachability probabilities until t are the same as those of the corresponding CTMC with uniform exit rate 1 until $r \cdot t$.

Proposition 19 induces a way to exactly compute the absolute difference in timed reachability probabilities of a CTMC $\mathcal{M}$ and its acceleration $c \cdot \mathcal{M}$.

Theorem 21. *Let $\mathcal{M}$ be a CTMC with $E(s) = 1$ for all $s \in S$, and let g be a unique absorbing goal state. Let $\delta > 0$, $c = e^\delta > 1$, and $\mathcal{M}' = c \cdot \mathcal{M}$. Then*

$$\mathrm{Diff}_t(\mathcal{M}) = |\mathrm{Pr}^{\mathcal{M}'}(\Diamond^{\leq t} g) - \mathrm{Pr}^{\mathcal{M}}(\Diamond^{\leq t} g)| = \sum_{n=1}^{\infty} p_n \cdot \mathrm{Diff}_t(\mathcal{E}_n),$$

where p_n denotes the probability to reach g after exactly n (discrete) steps.

Proof sketch. Let Π_n contain all finite paths entering g after exactly n steps, and let $\mathrm{Traj}^*(\pi)$ contain all timed versions of $\pi \in \Pi_n$ that enter g in $(t, c \cdot t]$. Then $\mathrm{Diff}_t(\mathcal{M}) = \sum_{n=0}^{\infty} \sum_{\pi \in \Pi_n} \mathrm{Pr}^{\mathcal{M}}(\mathrm{Traj}^*(\pi)) = \sum_{n=0}^{\infty} \sum_{\pi \in \Pi_n} \mathrm{Pr}^{\mathcal{M}}(\mathrm{Traj}^*(\pi) \mid \pi) \cdot \mathrm{Pr}^{\mathcal{M}}(\pi)$ by Bayes' rule. As $\mathrm{Pr}^{\mathcal{M}}(\mathrm{Traj}^*(\pi) \mid \pi) = \mathrm{Diff}_t(\mathcal{E}_n)$ for all $\pi \in \Pi_n$, $\mathrm{Diff}_t(\mathcal{E}_0) = 0$ for all t, and $\sum_{\pi \in \Pi_n} \mathrm{Pr}^{\mathcal{M}}(\pi) = p_n$, the claim follows. □

In Theorem 21 (and all other results in this section) we could have also used any (uniform) exit rate $r > 0$ for the states of $\mathcal{M}$ since, as described in Remark 20, we can model different rates by scaling the time bound t.

Next we show how to obtain, for given t and δ, a maximal N (that depends on t, δ) such that $\mathrm{Diff}_t(\mathcal{M}) \leq \mathrm{Diff}_t(\mathcal{E}_N)$. Using the explicit form of $\mathrm{Diff}_t(\mathcal{E}_N)$ from Proposition 19, this yields an easy-to-compute upper bound for $\mathrm{Diff}_t(\mathcal{M})$.

Proposition 22. *Let $N = \left\lceil \frac{(e^\delta - 1) \cdot t}{\delta} \right\rceil$. Under the conditions of Theorem 21,*

$$\mathrm{Diff}_t(\mathcal{M}) \leq \mathrm{Diff}_t(\mathcal{E}_N) = \sum_{k=0}^{N-1} \frac{t^k}{k!} \cdot \left(e^{-t} - c^k e^{-ct}\right).$$

The bound of Proposition 22 is usually tighter than that for $(0, \delta)$-bisimilar states in Corollary 14, since it does not converge to 1 exponentially fast for increasing t. See Figs. 6 and 7 for some examples of the behavior of the bound.

Remark 23. Let $\mathcal{M}$ be as in Theorem 21 and let X be the random variable that describes the probability to reach g for the first time after exactly n steps. Then $p_n \leq \mathrm{Pr}(X \leq n)$, so the Markov inequality implies $p_n \leq \frac{\mathbb{E}(X)}{n}$ for every n, where $\mathbb{E}(X)$ is the expected value of X. Together with Theorem 21 this yields

$$\mathrm{Diff}_t(\mathcal{M}) = \sum_{n=1}^{\infty} p_n \cdot \mathrm{Diff}_t(\mathcal{E}_n) \leq \mathbb{E}(X) \cdot \sum_{n=1}^{\infty} \frac{1}{n} \cdot \mathrm{Diff}_t(\mathcal{E}_n).$$

This bound heavily depends on $\mathbb{E}(X)$, i.e., on the expected number of steps until reaching g. If $\mathbb{E}(X)$ is small the bound can be tighter than those of Propositions 13 and 22, while it can quickly become trivial if $\mathbb{E}(X)$ is large.

Next, we present a way to compute the values of p_n based on a spectral decomposition of the transition probability matrix $\mathbf{P}$ of CTMC $\mathcal{M}$. We start with the case of a diagonalizable $\mathbf{P}$, and afterwards deal with arbitrary $\mathbf{P}$ by using the Jordan canonical form [50]. Moreover, we use the so obtained explicit formulas for p_n to obtain upper bounds for $\mathrm{Diff}_t(\mathcal{M})$. In contrast to the results so far, these bounds converge to 0 for $t \to \infty$, just like the actual difference does.

The transition probability matrix $\mathbf{P}$ of $\mathcal{M}$ is diagonalizable if there is a diagonal matrix $\mathbf{D}$, with diagonal elements corresponding to the eigenvalues of $\mathbf{P}$ (repeated according to their multiplicities), and a regular matrix $\mathbf{S}$ such that $\mathbf{P} = \mathbf{S}\mathbf{D}\mathbf{S}^{-1}$ [50]. It is well-known that every eigenvalue λ_i of a stochastic matrix satisfies $|\lambda_i| \leq 1$, and that $\lambda_1 = 1$ is an eigenvalue of any such matrix. As we assume that an absorbing state is reached almost surely, 1 is the only eigenvalue of modulus 1, and it has a multiplicity $a_{\mathbf{P}}$ of at most 2 (as there are at most two absorbing states in $\mathcal{M}$). Hence, we can w.l.o.g. assume that for the m distinct eigenvalues of $\mathbf{P}$ we have $\lambda_1 = 1 > |\lambda_2| \geq \ldots \geq |\lambda_m|$, and that the diagonal of $\mathbf{D}$ is in descending order w.r.t. the absolute values of these eigenvalues. From now on we denote the second largest (absolute value wise) eigenvalue of $\mathbf{P}$ by λ.

Proposition 24 ([63]). *Let $\mathcal{M}$ be a CTMC such that $\mathbf{P} = \mathbf{S}\mathbf{D}\mathbf{S}^{-1}$ is diagonalizable. Let $k \in \mathbb{N}$ and $n = |S|$. Then $p_{k+1} = \sum_{j=a_{\mathbf{P}}+1}^{n} \mathbf{S}_{1,j} \cdot \mathbf{S}^{-1}_{j,n} \cdot (\lambda_j - 1) \cdot \lambda_j^k$.*

Combining Theorem 21 and Proposition 24 yields a new bound for $\mathrm{Diff}_t(\mathcal{M})$.

Proposition 25. *Under the conditions of Theorem 21 we have*

$$\mathrm{Diff}_t(\mathcal{M}) \leq (n - a_{\mathbf{P}}) \cdot C \cdot \sum_{k=1}^{\infty} |\lambda|^{k-1} \cdot \mathrm{Diff}_t(\mathcal{E}_k),$$

where $n = |S|$ and $C = \max_{i=a_{\mathbf{P}}+1,\ldots,n} |\mathbf{S}_{1,i} \cdot \mathbf{S}^{-1}_{i,n} \cdot (\lambda_i - 1)|$.

Proof sketch. By Proposition 24, $p_{k+1} = \sum_{j=a_{\mathbf{P}}+1}^{n} \mathbf{S}_{1,j} \cdot \mathbf{S}^{-1}_{j,n} \cdot (\lambda_j - 1) \cdot \lambda_j^k$, so by the triangle inequality $p_{k+1} \leq \sum_{j=a_{\mathbf{P}}+1}^{n} |\mathbf{S}_{1,j} \cdot \mathbf{S}^{-1}_{j,n} \cdot (\lambda_j - 1)| \cdot |\lambda_j|^k$. Because $|\lambda_j| \leq |\lambda|$ for all $j > a_{\mathbf{P}}$ we get $p_{k+1} \leq |\lambda|^k \cdot \sum_{j=a_{\mathbf{P}}+1}^{n} C = |\lambda|^k \cdot (n - a_{\mathbf{P}}) \cdot C$ for $C = \max_{i=a_{\mathbf{P}}+1,\ldots,n} |\mathbf{S}_{1,i} \cdot \mathbf{S}^{-1}_{i,n} \cdot (\lambda_i - 1)|$. The claim follows from Theorem 21. □

The bound obtained in Proposition 25 is tight: Consider a CTMC $\mathcal{M}$ with two states s and g, such that $E(s) = E(g) = 1$ and $P(s,s) = p, P(s,g) = 1 - p$ for a $p \in (0,1)$. Let $\mathcal{M}' = e^{\delta} \cdot \mathcal{M}$, so $\mathcal{M} \sim_{0,\delta} \mathcal{M}'$. The probability matrix $\mathbf{P}$ of $\mathcal{M}$ is diagonalizable with eigenvalues $\lambda_1 = 1$ and $\lambda_2 = \lambda = p$. Then $p_k = |\sum_{j=2}^{2} \mathbf{S}_{1,j} \cdot \mathbf{S}^{-1}_{j,2} \cdot (\lambda_j - 1) \cdot \lambda_j^{k-1}| = (1 - p) \cdot p^{k-1} = (n - 1) \cdot C \cdot \lambda^{k-1}$, so here the inequalities in the proof of Proposition 25 are equalities. Furthermore, since always $|\lambda| < 1$, the bounds converge to 0 for $t \to \infty$, just like $\mathrm{Diff}_t(\mathcal{M})$.

Example 26. Consider the CTMC $\mathcal{M}$ on the left of Fig. 6. The probability matrix $\mathbf{P}$ of $\mathcal{M}$ is diagonalizable, with second largest eigenvalue $\lambda = \frac{1}{2}$. The graphic on the right of the figure compares the different error bounds for $\mathrm{Diff}_t(\mathcal{M})$ when $\delta = 0.1$. We can observe that the bound from Proposition 13 based on the uniformization method quickly approaches 1. The bound from Proposition 22

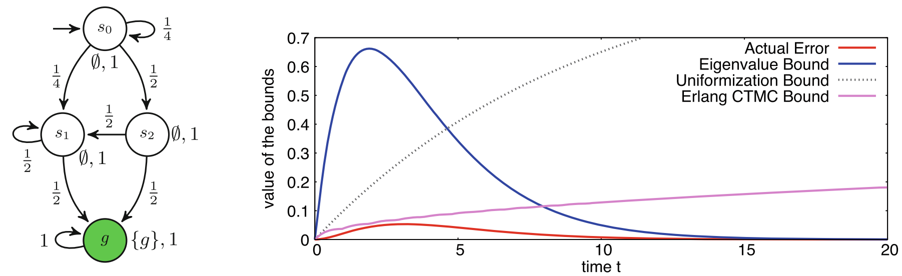

Fig. 6. A CTMC $\mathcal{M}$ with diagonalizable probability matrix (left) and a comparison of the different error bounds for $\mathrm{Diff}_t(\mathcal{M})$.

based on the maximal value of $\mathrm{Diff}_t(\mathcal{E}_n)$ performs well if t is small, but becomes imprecise if t grows. On the other hand, the bound from Proposition 25 that utilizes the diagonalization of $\mathbf{P}$ shows the opposite behavior. It performs bad for small values of t (as here the values of $\mathrm{Diff}(\mathcal{E}_k)$ are big for small k, for which $|\lambda|^k$ is not yet close to 0), while it converges to the actual error with increasing t.

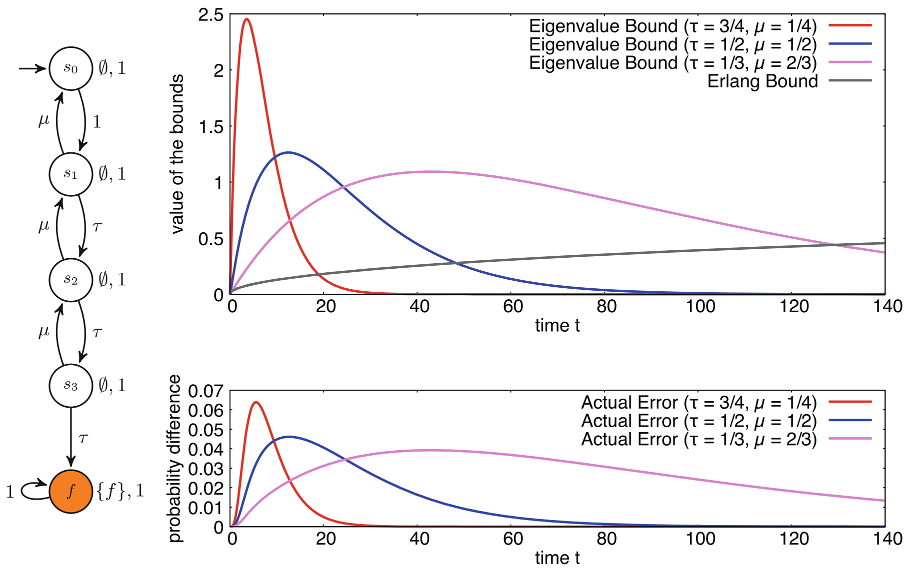

Fig. 7. The queue from Example 27 (left) and the bounds and errors for specific instances (right). The upper graphic depicts the bounds from Propositions 22 and 25 for the considered scenarios, the lower one the actual errors.

Example 27. Consider a single-server queue with capacity 4, modeled by the CTMC $\mathcal{M}$ on the left of Fig. 7. Tasks are completed with rate μ. If the queue is empty, the next task arrives with rate 1, and otherwise tasks arrive with rate τ. We assume that $\tau + \mu = 1$. If the queue is full, i.e., if it already contains four tasks, and a fifth one arrives it has to be dropped, leading to a fail-state f. This setting is close to the example discussed in the beginning of this section.

Let $\mathcal{M}' = e^{0.1} \cdot \mathcal{M}$. In particular, $\mathcal{M} \sim_{0,0.1} \mathcal{M}'$. We analyze the bounds from Propositions 22 and 25 w.r.t. $\mathcal{M}, \mathcal{M}'$ and f by considering three scenarios that are parameterized in τ and μ.[2] In the first scenario, $\tau_1 = \frac{3}{4}$ and $\mu_1 = \frac{1}{4}$, i.e., the arrival rate of tasks is three times higher than their service rate. In the second scenario, arrival and service rates coincide ($\tau_2 = \frac{1}{2} = \mu_2$), while in the third scenario the service rate is twice the arrival rate, i.e., $\tau_3 = \frac{1}{3}$ and $\mu_3 = \frac{2}{3}$. We denote the CTMC corresponding to scenario $i \in \{1, 2, 3\}$ by $\mathcal{M}_i$. Note that for each i the transition probability matrix $\mathbf{P}_i$ of $\mathcal{M}_i$ is diagonalizable.

The resulting bounds, as well as the actual errors, are depicted on the right-hand side of Fig. 7, where the same-colored (and similarly dashed) lines correspond to the same scenario (red and solid for the first, blue and dashed for the second and magenta and dash-dotted for the third). The dotted line in the upper graphic represents the bound from Proposition 22.

We observe that in each scenario the actual error and the bounds of Proposition 25 converge to 0 for $t \to \infty$, while the bound of Proposition 22 converges to 1 for increasing t. The convergence rate to 0 of the former bound can, however, be quite slow. In, e.g., the third scenario, the second largest eigenvalue of $\mathbf{P}_3$ has modulus 0.9778, causing a slow convergence of $|\lambda|^k$ to 0. On the other hand, in the first scenario $\mathcal{M}_1$ enters f quickly, which is reflected in a second largest eigenvalue of 0.7334 whose powers converge to 0 fast.

While the values of the bounds from Proposition 25 are, in particular for small t, significantly larger than the actual error, it is promising that their shape as a function of t closely resembles the actual difference. This scaling of the bounds by a large factor is caused by the constant C, which is obtained by using the triangle inequality and a maximum. Refining this constant in future work might lead to bounds following the actual error much more closely.

As the presented bounds work differently well for small and large t, the following formulation is useful.

Corollary 28. *Let C be as in Proposition 25, and N as in Proposition 22. Then*

$$\mathrm{Diff}_t(\mathcal{M}) \leq \min \left\{ \mathrm{Diff}_t(\mathcal{E}_N), (n - a_{\mathbf{P}}) \cdot C \cdot \sum_{k=1}^{\infty} |\lambda|^{k-1} \cdot \mathrm{Diff}_t(\mathcal{E}_k) \right\}.$$

Proposition 25 requires the transition probability matrix $\mathbf{P}$ of $\mathcal{M}$ to be diagonalizable. If this is not the case, we can use the *Jordan canonical form* (*JCF*) [50] instead, which looks as follows: Given an $n \times n$ matrix M over the field of complex numbers $\mathbb{C}$ with distinct eigenvalues $\lambda_1, \dots, \lambda_m$, given in descending order w.r.t. their absolutes values, the JCF of M is a decomposition of the form $M = \mathbf{SJS}^{-1}$ for matrices $\mathbf{S}, \mathbf{J} \in \mathbb{C}^{n \times n}$. $\mathbf{J}$ is a block diagonal matrix with so-called *Jordan blocks* $\mathbf{J}_i$, $i = 1, \dots, k$, $m \leq k \leq n$, on the diagonal. The i-th

[2] We omit the bound based on uniformization from Proposition 13 as we have already seen in Example 26 that it does not perform well for $(0, \delta)$-bisimilar models.

Jordan block $\mathbf{J}_i$ corresponds to one of M's eigenvalues τ and is a $r_i \times r_i$ matrix for some r_i that is at most the algebraic multiplicity of τ, and such that the sum over all r_i for the Jordan blocks belonging to the same eigenvalue τ equals its geometric multiplicity. The Jordan blocks for eigenvalue τ have τ everywhere on the diagonal and entry 1 directly above the diagonal. All other entries are 0. We w.l.o.g. assume that the Jordan blocks are in descending order w.r.t. the absolute value of the corresponding eigenvalues, and that multiple Jordan blocks for the same eigenvalue are in descending order w.r.t. their size. We write $q_{\mathbf{J}}$ for the total number of Jordan blocks of $\mathbf{J}$ and r_i for the size of the i-th Jordan block $\mathbf{J}_i$.

It is well-known that every matrix over $\mathbb{C}$ has a JCF [50], so the above decomposition exists for the transition probability matrix of *any* CTMC. We now show how to compute p_n, i.e., the probability to reach the goal state after exactly n discrete steps, in the case that $\mathbf{P}$ is not diagonalizable.

Proposition 29. *Denote by λ_i the eigenvalue corresponding to the i-th Jordan block $\mathbf{J}_i$ of $\mathbf{P} = \mathbf{SJS}^{-1}$. Let $N \in \mathbb{N}$, $z \in \{0, \ldots, q_{\mathbf{J}} - 1\}$ the number of Jordan blocks for eigenvalue 0 and $h_l = \sum_{i=1}^{l-1} r_i$. Then p_{N+1} equals*

$$\sum_{l=a_{\mathbf{P}}+1}^{q_{\mathbf{J}}-z} \sum_{j=1}^{r_l} \sum_{k=1}^{j} \mathbf{S}_{1,k+h_l} \cdot \mathbf{S}^{-1}_{j+h_l,n} \cdot \lambda_l^{N+k-j} \cdot \left(\lambda_l \cdot \binom{N+1}{j-k} - \binom{N}{j-k} \right) + R(N,z)$$

where, for $\mathbf{S}^*_{1,h_l+j} = \mathbf{S}_{1,h_l,+j}$ if $j > 0$ and $\mathbf{S}^*_{1,h_l+j} = 0$ if $j = 0$,

$$R(N,z) = \sum_{l=q_{\mathbf{J}}-z+1}^{q_{\mathbf{J}}} \sum_{j=0}^{r_l-(N+1)} (\mathbf{S}^*_{1,h_l+j} - \mathbf{S}_{1,h_l+j+1}) \cdot \mathbf{S}^{-1}_{h_l+N+j,n}.$$

In particular, the term $R(N, z)$ introduced in Theorem 29 equals 0 if $z = 0$ or if $N \geq \max_{q_{\mathbf{J}}-z+1 \leq i \leq q_{\mathbf{J}}} r_i$. Hence, if N is at least the size of the largest Jordan block corresponding to eigenvalue 0, the term $R(N, z)$ vanishes.

Remark 30. The decompositions of $\mathbf{P}$ proved in Propositions 24 and 29 also hold when the Markov chain contains more than 2 (reachable) absorbing states, i.e., if $a_{\mathbf{P}} > 2$. In an absorbing Markov chain, the multiplicity (both geometric and algebraic) of the eigenvalue 1 is the number of absorbing states, and thus each Jordan block for eigenvalue 1 has a size of 1, causing them to get canceled out in the computations of the explicit forms of p_{N+1}.

Combining Theorem 21 and Theorem 29 yields another upper bound on the absolute difference in timed reachability probabilities of $(0, \delta)$-bisimilar CTMCs.

Theorem 31. *In the setting of Theorem 29 let $|\lambda| > 0$. Let R be the maximal size of any Jordan block of $\mathbf{P}$, $r = \max_{i=a_{\mathbf{P}}+1,\ldots,q_{\mathbf{J}}-z} r_i$ the maximal size of any such block corresponding to an eigenvalue $\neq 0, 1$, $\lambda^*_l = \max\{|\lambda_l|, |1 - \lambda_l|\}$ for $l = a_{\mathbf{P}} + 1, \ldots, q_{\mathbf{J}} - z$ and $C = \sum_{l=a_{\mathbf{P}}+1}^{q_{\mathbf{J}}-z} \sum_{j=1}^{r_l} \sum_{k=1}^{j} |\mathbf{S}_{1,k+h_l} \cdot \mathbf{S}^{-1}_{j+h_l,n}| \cdot |\lambda^*_l|$. Then*

$$\mathrm{Diff}_t(\mathcal{M}) \leq \sum_{k=1}^{R-1} p_k \cdot \mathrm{Diff}_t(\mathcal{E}_k) + C \cdot \sum_{k=R}^{\infty} |\lambda|^{k-r} \cdot k^{r-1} \cdot \mathrm{Diff}_t(\mathcal{E}_k).$$

The bound of Proposition 31 is a conservative extension of the one for diagonalizable $\mathbf{P}$ from Proposition 25, as in this case $R = 1 = r$ and so the bound of Proposition 31 collapses to $\mathrm{Diff}_t(\mathcal{M}) \leq C \cdot \sum_{k=1}^{\infty} |\lambda|^{k-1} \cdot \mathrm{Diff}_t(\mathcal{E}_k)$.

Remark 32. In Proposition 31 we have to exclude the case $|\lambda| = 0$, i.e., the case that $\mathbf{P}$ only has eigenvalues 0 and 1. This happens, e.g., if $\mathbf{P}$ (and hence $\mathcal{M}$) is acyclic. In an acyclic CTMC the goal state g is, however, reached after *at most* as many steps as the length of the longest path in the chain. Therefore, if $\mathcal{M}$ is acyclic, $\mathrm{Diff}_t(\mathcal{M})$ can easily be computed exactly, for example by using the identity in Theorem 21 and computing the relevant values of $p_n = \mathbf{P}^{n+1} - \mathbf{P}^n$, or by applying methods like the ACE algorithm of [47] that allow computing the transient distribution functions of acyclic CTMCs directly.

Corollary 33. *Under the conditions and with the notation of Proposition 31 and Proposition 22,* $\mathrm{Diff}_t(\mathcal{M})$ *is bounded from above by*

$$\min \left\{ \mathrm{Diff}_t(\mathcal{E}_N), \sum_{k=1}^{R-1} p_k \cdot \mathrm{Diff}_t(\mathcal{E}_k) + C \cdot \sum_{k=R}^{\infty} |\lambda|^{k-r} \cdot k^{r-1} \cdot \mathrm{Diff}_t(\mathcal{E}_k) \right\}.$$

We finish the section by noting that the bounds for $(0, \delta)$-bisimilar CTMCs can also be used to derive bounds for the case that $\varepsilon > 0$. Given $\mathcal{M} \sim_{\varepsilon,\delta} \mathcal{N}$, by Theorem 11 there are CTMCs $\mathcal{M}'$ and $\mathcal{N}'$ with $\mathcal{M} \sim_{\varepsilon,0} \mathcal{M}' \sim_{0,\delta} \mathcal{N}' \sim \mathcal{N}$. Hence, the triangle inequality together with the fact that strong bisimilarity preserves timed reachability probabilities [13,19] implies $|\mathrm{Pr}^{\mathcal{M}}(\Diamond^{\leq t} g) - \mathrm{Pr}^{\mathcal{N}}(\Diamond^{\leq t} g)| \leq |\mathrm{Pr}^{\mathcal{M}}(\Diamond^{\leq t} g) - \mathrm{Pr}^{\mathcal{M}'}(\Diamond^{\leq t} g)| + |\mathrm{Pr}^{\mathcal{M}'}(\Diamond^{\leq t} g) - \mathrm{Pr}^{\mathcal{N}'}(\Diamond^{\leq t} g)|$. As $\mathcal{M} \sim_{\varepsilon,0} \mathcal{M}'$, the bound from Proposition 13 for the special case of $\varepsilon = 0$ is applicable to the first term, and one of the bounds for $(0, \delta)$-bisimilar chains can be applied to the second term. This yields another upper bound on the absolute difference in timed reachability probabilities for (ε, δ)-bisimilar CTMCs.

6 Bounds for Reward-Bounded Reachability Probabilities

In practice, CTMCs are often extended with *rewards* that allow to model, e.g., the accumulation of costs, the energy consumption, or different performance measures like the availability of the system [12,31]. We consider state-based reward functions $\rho\colon S \to \mathbb{R}$ that assign to every $s \in S$ a reward $\rho(s)$. Rewards are accumulated in the states, and the cumulative reward along the timed path $\pi = s_0 t_0 s_1 t_1 \ldots \in \mathrm{Paths}(\mathcal{M})$ of $\mathcal{M}$ until time t is given (for $\sigma@t = s_m$) as [12]

$$\rho(\sigma, t) = \sum_{j=0}^{m-1} t_j \cdot \rho(s_j) + \left(t - \sum_{j=0}^{m-1} t_j \right) \cdot \rho(s_m).$$

Similar to timed reachability probabilities $\mathrm{Pr}_s(\Diamond^{\leq t} g)$, logics like CSRL [12] allow to specify properties in terms of *reward-bounded* reachability probabilities $\mathrm{Pr}_s(\Diamond_{\leq r} g)$, which denote the probability to reach, from state s, a goal state g while accumulating at most reward $r \in \mathbb{R}$.

We extend the results of Sects. 4 and 5 to reward-bounded reachability probabilities. Here, we consider the special case of *nonnegative* rewards, i.e., we assume that $\rho(s) \geq 0$ for all $s \in S$. To accommodate for the addition of rewards to our model, we adjust the definition of (ε, δ)-bisimulations to require related states to have the same reward, i.e., $s \sim_{\varepsilon,\delta} s'$ is possible only if $\rho(s) = \rho(s')$. An extension to a notion of, say, $(\varepsilon, \delta, \mu)$-bisimulation, where μ describes an allowed error in the rewards in related states, is an interesting direction for future work.

First, assume $\rho(s) > 0$ for all s, a setting also analyzed in, e.g., [12]. There, based on ideas from [17], a transformation is proposed that allows to compute reward-bounded reachability probabilities in $\mathcal{M}$ via *time-bounded* reachability probabilities in a modified model $\widehat{\mathcal{M}}$. This model differs from $\mathcal{M}$ only in the exit rates $\widehat{E}(s) = \frac{E(s)}{\rho(s)}$ and rewards $\widehat{\rho}(s) = \frac{1}{\rho(s)}$ of states $s \in S$. A direct consequence of [12, Lem. 1] is that $\Pr_s^{\mathcal{M}}(\Diamond_{\leq r} g) = \Pr_s^{\widehat{\mathcal{M}}}(\Diamond^{\leq r} g)$.

Thus, given $s \sim_{\varepsilon,\delta} s'$ we can directly apply the bounds obtained in Sects. 4 and 5, provided that s and s' are still (ε, δ)-bisimilar in $\widehat{\mathcal{M}}$. However, this is indeed the case as we require $\rho(s) = \rho(s')$, and so the transformation of [12] simply scales the exit rates (and rewards) of s and s' by a common factor.

Proposition 34. *Let $s \sim_{\varepsilon,\delta} s'$, $r \geq 0$ and $\widehat{q} \geq \max_{p \in S} \widehat{E}(p)$. Then*

$$|\Pr_s(\Diamond_{\leq r} g) - \Pr_{s'}(\Diamond_{\leq r} g)| \leq 1 - e^{-\widehat{q} \cdot t \cdot (e^{\delta}(\varepsilon+1)-1)}.$$

If $\varepsilon = 0$, a similar result can be formulated for the bounds from Sect. 5.

Next, we allow $\rho(s) = 0$. As the reward of the goal state g does not matter for our purpose, we can w.l.o.g. assume $\rho(g) > 0$. The transformation from [12] is not applicable, as it could require division by 0. To circumvent this issue we first construct from $\mathcal{M}$ a CTMC $\mathcal{M}_{>0}$ without states with $\rho(s) = 0$. Intuitively, this is done by redirecting the incoming transitions of states with reward 0 to their successors, allowing us to remove any such state from the chain.

Proposition 35. *Let $\mathcal{M}$ be a CTMC. There is a CTMC $\mathcal{M}_{>0}$ with $S^{\mathcal{M}_{>0}} = S^{\mathcal{M}} \setminus \{x \mid \rho(x) = 0\}$ such that $\Pr_s^{\mathcal{M}}(\Diamond_{\leq r} g) = \Pr_s^{\mathcal{M}_{>0}}(\Diamond_{\leq r} g)$ for all $s \in S^{\mathcal{M}_{>0}}$.*

Hence, by first constructing $\mathcal{M}_{>0}$ from $\mathcal{M}$ and afterwards applying the transformation from [12] to $\mathcal{M}_{>0}$, we obtain the following corollary.

Corollary 36. *Proposition 34 also holds for CTMCs with nonnegative rewards.*

7 Conclusion and Future Work

We introduced (ε, δ)-bisimulations, a novel type of approximate probabilistic bisimulation for CTMCs. In contrast to notions from the literature such as quasi-lumpability [19,29,30], the separate bounds on changes in transition probabilities and changes in exit rates used in (ε, δ)-bisimulation result in a more flexible, fine-grained notion of behavioral similarity. We established fundamental properties of the notion, and analyzed the difference in timed reachability probabilities of (ε, δ)-bisimilar chains.

We obtained bounds on this difference by uniformizing the chain and applying known results from [18] for the discrete-time setting. Although tight if $\delta = 0$, the bounds usually do not perform well if $\delta > 0$. Hence, we investigated the special case of $(0, \delta)$-bisimilarity in more detail—also as a stepping stone towards a refined treatment of (ε, δ)-bisimilarity in future work. For $(0, \delta)$-bisimilar chains, we provided bounds based on the error in Erlang CTMCs $\mathcal{E}_n$, as well as based on a spectral analysis of the probability matrix $\mathbf{P}$ of the CTMC in question. Furthermore, we extended our results to reward-bounded reachability probabilities, provided that all rewards in the model are nonnegative.

From a theoretical point of view, the behavior of the bounds obtained from the spectral analysis of $\mathbf{P}$ is promising. They are tight and, as can be observed in Fig. 7, seem to evolve similar to the actual error. The bounds are, however, stretched by a (large) factor, which happens since, to obtain the constants C, the triangle inequality and maxima are used. This causes the constants to become significantly larger than necessary, leading to poor performance of the bounds in particular if t is small. Therefore, searching for smaller constants $C' \ll C$ in future work is in order. From a practical point of view, applicability of the bounds seems questionable. They tend to over-approximate the error and, depending on the second largest eigenvalue of $\mathbf{P}$, can take a long time (i.e., require large values of t) to converge to the actual error, even for small chains. Moreover, computing Jordan canonical forms (and diagonalizations) is computationally expensive and numerically unstable [26,53,54]. Consequently, the search for more practically relevant bounds is an important direction for future work.

Other open questions include an extension of (ε, δ)-bisimulation to a notion of $(\varepsilon, \delta, \mu)$-bisimulation that allows a tolerance of μ in the rewards of related states, an analysis of the achievable state space reduction in quotients w.r.t. (transitive) (ε, δ)-bisimulations, and a closer look at the (approximate) preservation of logical properties expressed by, e.g., CSL-formulas [9,13], between (ε, δ)-bisimilar states.

Acknowledgements. This work was partly funded by the DFG through the DFG grant 389792660 as part of TRR 248 (see https://perspicuous-computing.science) and the Cluster of Excellence EXC 2050/1 (CeTI, project ID 390696704, as part of Germany's Excellence Strategy).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Abate, A.: Approximation metrics based on probabilistic bisimulations for general state-space Markov processes: a survey. In: Proceedings of the First Workshop on Hybrid Autonomous Systems. Electronic Notes in Theoretical Computer Science, vol. 297, pp. 3–25 (2013). https://doi.org/10.1016/j.entcs.2013.12.002
2. Abate, A., Brim, L., Češka, M., Kwiatkowska, M.: Adaptive aggregation of Markov chains: quantitative analysis of chemical reaction networks. In: Kroening, D., Păsăreanu, C.S. (eds.) Computer Aided Verification. Lecture Notes in Computer Science (LNCS), vol. 9206, pp. 195–213. Springer International Publishing, Cham (2015). https://doi.org/10.1007/978-3-319-21690-4_12
3. Abate, A., Kwiatkowska, M., Norman, G., Parker, D.: Probabilistic model checking of labelled Markov processes via finite approximate bisimulations. In: van Breugel, F., Kashefi, E., Palamidessi, C., Rutten, J. (eds.) Horizons of the Mind. A Tribute to Prakash Panangaden: Essays Dedicated to Prakash Panangaden on the Occasion of His 60th Birthday, Lecture Notes in Computer Science (LNCS), vol. 8464, pp. 40–58. Springer International Publishing, Cham (2014). https://doi.org/10.1007/978-3-319-06880-0_2
4. Češka, M., Šafránek, D., Dražan, S., Brim, L.: Robustness analysis of stochastic biochemical systems. PLoS ONE **9**(4), e94553 (2014). https://doi.org/10.1371/journal.pone.0094553
5. Aldini, A.: A Note on the Approximation of Weak Probabilistic Bisimulation (2009)
6. Amparore, E.G., Beccuti, M., Donatelli, S.: (Stochastic) model checking in GreatSPN. In: Petri Nets. Lecture Notes in Computer Science (LNCS), vol. 8489, pp. 354–363. Springer (2014). https://doi.org/10.1007/978-3-319-07734-5_19
7. Amparore, E.G., Donatelli, S.: Efficient model checking of the stochastic logic CSL^{TA}. Perform. Eval. **123–124**, 1–34 (2018). https://doi.org/10.1016/j.peva.2018.03.002
8. Arnold, F., Belinfante, A., van der Berg, F.I., Guck, D., Stoelinga, M.: DFTCalc: a tool for efficient fault tree analysis. In: SAFECOMP. Lecture Notes in Computer Science (LNCS), vol. 8153, pp. 293–301. Springer (2013). https://doi.org/10.1007/978-3-642-40793-2_27
9. Aziz, A., Sanwal, K., Singhal, V., Brayton, R.: Model-checking continuous-time Markov chains. ACM Trans. Comput. Logic **1**(1), 162–170 (2000). https://doi.org/10.1145/343369.343402
10. Baier, C.: Polynomial time algorithms for testing probabilistic bisimulation and simulation. In: Alur, R., Henzinger, T.A. (eds.) Computer Aided Verification. Lecture Notes in Computer Science (LNCS), vol. 1102, pp. 50–61. Springer, Berlin, Heidelberg (1996). https://doi.org/10.1007/3-540-61474-5_57
11. Baier, C., Cloth, L., Haverkort, B.R., Kuntz, M., Siegle, M.: Model checking Markov chains with actions and state labels. IEEE Trans. Software Eng. **33**(4), 209–224 (2007). https://doi.org/10.1109/TSE.2007.36
12. Baier, C., Haverkort, B., Hermanns, H., Katoen, J.P.: On the logical characterisation of performability properties. In: Montanari, U., Rolim, J.D.P., Welzl, E. (eds.) Automata, Languages and Programming (ICALP 2000). Lecture Notes in Computer Science (LNCS), vol. 1853, pp. 780–792. Springer Berlin Heidelberg, Berlin, Heidelberg (2000). https://doi.org/10.1007/3-540-45022-X_65
13. Baier, C., Haverkort, B.R., Hermanns, H., Katoen, J.P.: Model-checking algorithms for continuous-time Markov chains. IEEE Trans. Softw. Eng. **29**(6), 524–541 (2003). https://doi.org/10.1109/TSE.2003.1205180

14. Baier, C., Katoen, J.P.: Principles of Model Checking. The MIT Press (2008)
15. Baier, C., Katoen, J.P., Hermanns, H., Wolf, V.: Comparative branching-time semantics for Markov chains. Inf. Comput. **200**(2), 149–214 (2005). https://doi.org/10.1016/j.ic.2005.03.001
16. Bartoletti, M., Murgia, M., Zunino, R.: Sound approximate and asymptotic probabilistic bisimulations for PCTL. Log. Meth. Comput. Sci. **19**(1) (2023). https://doi.org/10.46298/lmcs-19(1:22)2023
17. Beaudry, M.D.: Performance-related reliability measures for computing systems. IEEE Trans. Comput. **C-27**(6), 540–547 (1978). https://doi.org/10.1109/TC.1978.1675145
18. Bian, G., Abate, A.: On the relationship between bisimulation and trace equivalence in an approximate probabilistic context. In: Esparza, J., Murawski, A.S. (eds.) Foundations of Software Science and Computation Structures (FoSSaCS). Lecture Notes in Computer Science (LNCS), vol. 10203, pp. 321–337. Springer, Berlin, Heidelberg (2017). https://doi.org/10.1007/978-3-662-54458-7_19
19. Buchholz, P.: Exact and ordinary lumpability in finite Markov chains. J. Appl. Probab. **31**(1), 59–75 (1994). https://doi.org/10.2307/3215235
20. Cardelli, L., Grosu, R., Larsen, K.G., Tribastone, M., Tschaikowski, M., Vandin, A.: Lumpability for uncertain continuous-time Markov chains. In: Abate, A., Marin, A. (eds.) Quantitative Evaluation of Systems (QEST 2021), Lecture Notes in Computer Science (LNCS), vol. 12846, pp. 391–409. Springer International Publishing, Cham (2021). https://doi.org/10.1007/978-3-030-85172-9_21
21. Cardelli, L., Grosu, R., Larsen, K.G., Tribastone, M., Tschaikowski, M., Vandin, A.: Algorithmic minimization of uncertain continuous-time markov chains. IEEE Trans. Autom. Control **68**(11), 6557–6572 (2023). https://doi.org/10.1109/TAC.2023.3244093
22. Chen, T., Han, T., Katoen, J.P., Mereacre, A.: Quantitative model checking of continuous-time Markov chains against timed automata specifications. In: 24th IEEE Symposium on Logic In Computer Science (LICS), pp. 309–318. IEEE Computer Society (2009). https://doi.org/10.1109/LICS.2009.21
23. Desharnais, J., Edalat, A., Panangaden, P.: A logical characterization of bisimulation for labeled Markov processes. In: Proceedings of the Thirteenth Annual IEEE Symposium on Logic in Computer Science (LiCS), pp. 478–487 (1998). https://doi.org/10.1109/LICS.1998.705681
24. Desharnais, J., Edalat, A., Panangaden, P.: Bisimulation for labelled Markov processes. Inf. Comput. **179**(2), 163–193 (2002). https://doi.org/10.1006/inco.2001.2962
25. Desharnais, J., Laviolette, F., Tracol, M.: Approximate analysis of probabilistic processes: logic, simulation and games. In: Fifth International Conference on Quantitative Evaluation of Systems (QEST 2008), pp. 264–273 (2008). https://doi.org/10.1109/QEST.2008.42
26. Dey, P., Kannan, R., Ryder, N., Srivastava, N.: Bit complexity of Jordan normal form and polynomial spectral factorization. In: Tauman Kalai, Y. (ed.) 14th Innovations in Theoretical Computer Science Conference (ITCS 2023). Leibniz International Proceedings in Informatics (LIPIcs), vol. 251, pp. 42:1–42:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany (2023). https://doi.org/10.4230/LIPIcs.ITCS.2023.42
27. D'Innocenzo, A., Abate, A., Katoen, J.P.: Robust PCTL model checking. In: Proceedings of the 15th ACM International Conference on Hybrid Systems: Computation and Control (HSCC 2012), pp. 275–286. Association for Computing Machinery, New York, NY, USA (2012). https://doi.org/10.1145/2185632.2185673

28. Erlang, A.K.: Solution to some problems of the theory of probabilities of significance in automatic telephone exchange. Post Office Electric. Eng. J., 189–197 (1917)
29. Franceschinis, G., Muntz, R.R.: Bounds for quasi-lumpable Markov chains. Perform. Eval. **20**(1), 223–243 (1994). https://doi.org/10.1016/0166-5316(94)90015-9
30. Franceschinis, G., Muntz, R.R.: Computing bounds for the performance indices of quasi-lumpable stochastic well-formed nets. IEEE Trans. Softw. Eng. **20**(7), 516–525 (1994). https://doi.org/10.1109/32.297940
31. Gouberman, A., Siegle, M.: Markov reward models and Markov decision processes in discrete and continuous time: performance evaluation and optimization. In: Remke, A., Stoelinga, M. (eds.) Stochastic Model Checking. Rigorous Dependability Analysis Using Model Checking Techniques for Stochastic Systems: International Autumn School (ROCKS 2012). Lecture Notes in Computer Science (LNCS), vol. 8453, pp. 156–241. Springer Berlin Heidelberg, Berlin, Heidelberg (2014). https://doi.org/10.1007/978-3-662-45489-3_6
32. Grassmann, W.K.: Transient solutions in Markovian queueing systems. Computers & Operations Research **4**(1), 47–53 (1977). https://doi.org/10.1016/0305-0548(77)90007-7
33. Haesaert, S., Nilsson, P., Soudjani, S.: Formal multi-objective synthesis of continuous-state MDPs. In: 2021 American Control Conference (ACC), pp. 3428–3433 (2021). https://doi.org/10.23919/ACC50511.2021.9482873
34. Hensel, C., Junges, S., Katoen, J., Quatmann, T., Volk, M.: The probabilistic model checker STORM. Int. J. Softw. Tools Technol. Trans. (STTT) **24**, 589–610 (2022). https://doi.org/10.1007/s10009-021-00633-z
35. Jensen, A.: Markoff chains as an aid in the study of Markoff processes. Scand. Actuar. J. **1953**(sup1), 87–91 (1953). https://doi.org/10.1080/03461238.1953.10419459
36. Kemeny, J.G., Snell, J.L.: Finite Markov Chains. Undergraduate Texts in Mathematics (UTM), Springer, New York (1976)
37. Kiefer, S., Tang, Q.: Approximate bisimulation minimisation. In: Bojańczy, M., Chekuri, C. (eds.) 41st IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2021). Leibniz International Proceedings in Informatics (LIPIcs), vol. 213, pp. 48:1–48:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany (2021). https://doi.org/10.4230/LIPIcs.FSTTCS.2021.48
38. Kolmogorov, A.N.: Über die analytischen methoden in der wahrscheinlichkeitsrechnung. Math. Ann. **104**, 415–458 (1931). https://doi.org/10.1007/BF01457949
39. Kulkarni, V.G.: Introduction to Modeling and Analysis of Stochastic Systems. Springer Texts in Statistics (STS), Springer New York (2010). https://doi.org/10.1007/978-1-4419-1772-0
40. Kwiatkowska, M., Norman, G., Parker, D.: PRISM 4.0: verification of probabilistic real-time systems. In: Gopalakrishnan, G., Qadeer, S. (eds.) Proc. 23rd International Conference on Computer Aided Verification (CAV'11). Lecture Notes in Computer Science (LNCS), vol. 6806, pp. 585–591. Springer (2011). https://doi.org/10.1007/978-3-642-22110-1_47
41. Kwiatkowska, M.Z., Thachuk, C.: Probabilistic model checking for biology. In: Software Systems Safety, NATO Science for Peace and Security Series, D: Information and Communication Security, vol. 36, pp. 165–189. IOS Press (2014). https://doi.org/10.3233/978-1-61499-385-8-165

42. Larsen, K.G., Skou, A.: Bisimulation through probabilistic testing. Inf. Comput. **94**(1), 1–28 (1991). https://doi.org/10.1016/0890-5401(91)90030-6
43. Lin, N., Liu, Y.: Perturbation analysis for continuous-time Markov chains in a weak sense. J. Appl. Probab. **61**(4), 1278–1300 (2024). https://doi.org/10.1017/jpr.2024.20
44. Liu, Y.Y.: Perturbation analysis for continuous-time Markov chains. Sci. China Math. **58**(12), 2633–2642 (2015). https://doi.org/10.1007/s11425-015-5019-z
45. Liu, Y., Li, W.: Error bounds for augmented truncation approximations of Markov chains via the perturbation method. Adv. Appl. Probab. **50**(2), 645–669 (2018). https://doi.org/10.1017/apr.2018.28
46. Madsen, C., Zhang, Z., Roehner, N., Winstead, C., Myers, C.J.: Stochastic model checking of genetic circuits. ACM J. Emerg. Technol. Comput. Syst. **11**(3), 23:1–23:21 (2014). https://doi.org/10.1145/2644817
47. Maire, R.A., Reibman, A.L., Trivedi, K.S.: Transient analysis of acyclic Markov chains. Perform. Eval. **7**(3), 175–194 (1987). https://doi.org/10.1016/0166-5316(87)90039-3
48. Marin, A., Piazza, C., Rossi, S.: Proportional lumpability. In: André, É., Stoelinga, M. (eds.) Formal Modeling and Analysis of Timed Systems (FORMATS) 2019. Lecture Notes in Computer Science (LNCS), vol. 11750, pp. 265–281. Springer International Publishing, Cham (2019). https://doi.org/10.1007/978-3-030-29662-9_16
49. Marin, A., Piazza, C., Rossi, S.: Proportional lumpability and proportional bisimilarity. Acta Inf. **59**(2–3), 211–244 (2022). https://doi.org/10.1007/s00236-021-00404-y
50. Meyer, C.D.: Matrix Analysis and Applied Linear Algebra. Society for Industrial and Applied Mathematics, Philadelphia, PA (2000)
51. Mitrophanov, A.Y.: Sensitivity and convergence of uniformly ergodic Markov chains. J. Appl. Probab. **42**(4), 1003–1014 (2005). https://doi.org/10.1239/jap/1134587812
52. Mitrophanov, A.Y.: The arsenal of perturbation bounds for finite continuous-time Markov chains: a perspective. Mathematics **12**(11) (2024). https://doi.org/10.3390/math12111608
53. Moler, C., Van Loan, C.: Nineteen dubious ways to compute the exponential of a matrix. twenty-five years later. SIAM Rev. **45**(1), 3–49 (2003). https://doi.org/10.1137/S00361445024180
54. Pan, V.Y., Chen, Z.Q.: The complexity of the matrix eigenproblem. In: Proceedings of the thirty-first Annual ACM Symposium on Theory of Computing, pp. 507–516. ACM, Atlanta Georgia USA (1999). https://doi.org/10.1145/301250.301389
55. Piazza, C., Rossi, S.: Reasoning about proportional lumpability. In: Abate, A., Marin, A. (eds.) Quantitative Evaluation of Systems (QEST 2021). Lecture Notes in Computer Science (LNCS), vol. 12846, pp. 372–390. Springer International Publishing, Cham (2021). https://doi.org/10.1007/978-3-030-85172-9_20
56. Piazza, C., Rossi, S., Smuseva, D.: Efficient algorithm for proportional lumpability and its application to selfish mining in public blockchains. Algorithms **17**(4) (2024). https://doi.org/10.3390/a17040159
57. Segala, R., Lynch, N.: Probabilistic simulations for probabilistic processes. In: Jonsson, B., Parrow, J. (eds.) CONCUR '94: Concurrency Theory. Lecture Notes in Computer Science (LNCS), vol. 836, pp. 481–496. Springer, Berlin, Heidelberg (1994). https://doi.org/10.1007/978-3-540-48654-1_35

58. Shardlow, T., Stuart, A.M.: A perturbation theory for ergodic Markov chains and application to numerical approximations. SIAM J. Numer. Anal. **37**(4), 1120–1137 (2000). https://doi.org/10.1137/S0036142998337235
59. Spork, T., Baier, C., Katoen, J.P., Klüppelholz, S., Piribauer, J.: Approximate Probabilistic Bisimulation for Continuous-Time Markov Chains (2025). https://doi.org/10.48550/arXiv.2505.15587
60. Spork, T., Baier, C., Katoen, J.P., Piribauer, J., Quatmann, T.: A spectrum of approximate probabilistic bisimulations. In: Majumdar, R., Silva, A. (eds.) 35th International Conference on Concurrency Theory (CONCUR 2024). Leibniz International Proceedings in Informatics (LIPIcs), vol. 311, pp. 37:1–37:19. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany (2024). https://doi.org/10.4230/LIPIcs.CONCUR.2024.37
61. Tracol, M., Desharnais, J., Zhioua, A.: Computing distances between probabilistic automata. In: Massink, M., Norman, G. (eds.) Proceedings of the Ninth Workshop on Quantitative Aspects of Programming Languages (QAPL 2011), Saarbrücken, Germany, April 1-3, 2011. EPTCS, vol. 57, pp. 148–162 (2011). https://doi.org/10.4204/eptcs.57.11
62. Volk, M., Junges, S., Katoen, J.P.: Fast dynamic fault tree analysis by model checking techniques. IEEE Trans. Ind. Info. **14**(1), 370–379 (2018). https://doi.org/10.1109/TII.2017.2710316
63. Watterson, G.A.: Markov chains with absorbing states: a genetic example. Ann. Math. Stat. **32**(3), 716–729 (1961). https://doi.org/10.1214/aoms/1177704967
64. Zhang, L., Jansen, D.N., Nielson, F., Hermanns, H.: Automata-based CSL model checking. Log. Meth. Comput. Sci. **8**(2) (2011). https://doi.org/10.1007/978-3-642-22012-8_21