

A Spectrum of Approximate Probabilistic Bisimulations

Timm Spork  

Technische Universität Dresden, Dresden, Germany

Christel Baier  

Technische Universität Dresden, Dresden, Germany

Joost-Pieter Katoen  

RWTH Aachen University, Aachen, Germany

Jakob Piribauer  

Technische Universität Dresden, Dresden, Germany

Universität Leipzig, Leipzig, Germany

Tim Quatmann  

RWTH Aachen University, Aachen, Germany

Abstract

This paper studies various notions of approximate probabilistic bisimulation on labeled Markov chains (LMCs). We introduce approximate versions of weak and branching bisimulation, as well as a notion of ε -perturbed bisimulation that relates LMCs that can be made (exactly) probabilistically bisimilar by small perturbations of their transition probabilities. We explore how the notions interrelate and establish their connections to other well-known notions like ε -bisimulation.

2012 ACM Subject Classification Theory of computation $\rightarrow$ Logic and verification; Theory of computation $\rightarrow$ Random walks and Markov chains

Keywords and phrases Markov chains, Approximate bisimulation, Abstraction, Model checking

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2024.37

Related Version *Full Version:* <http://arxiv.org/abs/2407.07584>

Funding *Christel Baier, Jakob Piribauer and Timm Spork:* This work was partly funded by the DFG Grant 389792660 as part of TRR 248 (Foundations of Perspicuous Software Systems) and the Cluster of Excellence EXC 2050/1 (CeTI, project ID 390696704, as part of Germany's Excellence Strategy).

Tim Quatmann: This research was funded by a KI-Starter grant from the Ministerium für Kultur und Wissenschaft NRW.

Acknowledgements We thank the reviewers for their helpful feedback, comments and suggestions. In particular, we thank the reviewer who pointed out the work in [31, 32] on approximate simulation relations which we were previously not aware of.

1 Introduction

Probabilistic model checking is widely used for the automatic verification of probabilistic models, like labeled Markov chains (LMC), against properties specified in (temporal) logics like PCTL* [11]. In practice, a big obstacle is the *state space explosion problem*: the number of states required to model a system can make its verification intractable [37, 11, 36].

To circumvent this issue, a well-established approach is the use of abstractions. For a given LMC $\mathcal{M}$, an abstraction $\mathcal{A}$ is a model derived from $\mathcal{M}$ that is (oftentimes) smaller than $\mathcal{M}$ and preserves some properties of interest. Instead of verifying a formula on $\mathcal{M}$, one does so on $\mathcal{A}$ and afterwards transfers the result back to the original model [11, 27, 38].



© Timm Spork, Christel Baier, Joost-Pieter Katoen, Jakob Piribauer, and Tim Quatmann;
licensed under Creative Commons License CC-BY 4.0

35th International Conference on Concurrency Theory (CONCUR 2024).

Editors: Rupak Majumdar and Alexandra Silva; Article No. 37; pp. 37:1–37:19

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

A prominent type of abstraction are probabilistic bisimulation quotients. They are constructed w.r.t. probabilistic bisimulations, a class of behavioral equivalences introduced by Larsen and Skou [41] as an extension of Milner's bisimulation [43] to probabilistic models. A probabilistic bisimulation is an equivalence R on the state space of an LMC $\mathcal{M}$ that only relates states that behave exactly the same, i.e., that have the same local properties, and transition to R -equivalence classes with equal probability. The coarsest probabilistic bisimulation $\sim$, called (probabilistic) bisimilarity, is the union of all probabilistic bisimulations in $\mathcal{M}$ [11]. The bisimilarity relation can be computed efficiently [9, 19, 48] and preserves PCTL* state formulas [8, 33]. Since verifying PCTL* on bisimulation quotients can significantly speed up the verification process [37], their use is a vital part of probabilistic model checkers such as, e.g., STORM [34].

Other notions of behavioral equivalence are *weak* and *branching* probabilistic bisimulations [42, 51, 10, 12, 18, 50], which were introduced with the intention to abstract from sequences of *internal actions* or *stutter steps* a model can perform. Intuitively, these notions can abstract from the possibility of a state to, for some time, only visit equally labeled states (weak) or stay in its own equivalence class (branching) [35]. It is well-known that weak and branching probabilistic bisimilarity, denoted $\approx^w$ and $\approx^b$, respectively, coincide for LMCs [10], and that they characterize satisfaction equivalence for a variant of PCTL* [24].

A problem with all of the above notions lies, however, in their lack of robustness against errors in the transition probabilities. The requirement of related states to have *exactly* the same transition probabilities to equivalence classes implies that even an infinitesimally small perturbation of any of these probabilities can cause two bisimilar states to become non-bisimilar, resulting in larger quotients [20, 52, 27]. This disadvantage was first observed in [30], where the use of *approximate* notions of bisimulation is suggested for its mitigation.

The literature proposes various types of approximate bisimilarity, the most well-known and well-studied one being ε -bisimilarity ($\sim_\varepsilon$) [25]. Other notions include approximate probabilistic bisimilarity with precision ε ($\equiv_\varepsilon$), or ε -APB for short [27, 1, 2], up-to- (n, ε) -bisimilarity ($\sim_\varepsilon^n$) [25, 13], or ε -lumpability of a given LMC [17, 29, 28]. Here, we propose definitions for approximate versions of weak ($\approx_\varepsilon^w$) and branching probabilistic bisimilarity ($\approx_\varepsilon^b$). Similar notions have, to the best of our knowledge, only been discussed sporadically in the context of noninterference under the term “weak bisimulation with precision ε ” [4, 7, 5, 6, 26, 3]. Moreover, we introduce ε -perturbed bisimilarity ($\simeq_\varepsilon$) which relates two LMCs if they can be made bisimilar by small perturbations of their transition probabilities. Implicitly, this relation arises in the work [38] on a type of abstraction called ε -quotients. With our definition, two LMCs are ε -perturbed bisimilar iff they have bisimilar ε -quotients.

All of the approximate notions have in common that they allow a small tolerance, say $\varepsilon > 0$, in the transition probabilities of related states, but differ in the specifics of where and how this tolerance is put to use. Broadly speaking, we can distinguish two groups of relations: while $\sim_\varepsilon$, $\equiv_\varepsilon$, $\sim_\varepsilon^n$ and $\approx_\varepsilon^w$ are additive in their tolerances and are closer to classic process relations, the notions underlying $\sim_\varepsilon^*$ and $\equiv_\varepsilon^*$, denoting *transitive* ε -bisimilarity and *transitive* ε -APB, respectively, as well as $\simeq_\varepsilon$ and $\approx_\varepsilon^b$ are better suited for the construction of abstractions since they are required to be equivalences. Collapsing the equivalence classes of such a relation into single states yields quotient models, which in some cases are such that formulas given in specific (fragments of) logics are (approximately) preserved between the original LMC and its quotient. However, it turns out that requiring transitivity can cause some unnatural behavior, like the possibility to distinguish probabilistically bisimilar LMCs and a lack of additivity. Furthermore, the induced bisimilarity relations, which are again defined as the union of all corresponding relations in the model $\mathcal{M}$ (e.g., $\approx_\varepsilon^b$ is the union of

■ **Table 1** Overview of the notions of approximate bisimulation we consider and some of their properties. Being suitable for “quotienting” is meant w.r.t. the underlying bisimulation relation.

Notion	Symbol for Union	Additive	Quotienting
ε -Bisimulation [25, 14]	$\sim_\varepsilon$	✓	×
ε -APB [27, 1, 2]	$\equiv_\varepsilon$		
Up-To- (n, ε) -Bisimulation [25, 13]	$\sim_\varepsilon^n$		
Weak ε -Bisimulation	$\sim_\varepsilon^w$		
Transitive ε -Bisimulation	$\sim_\varepsilon^*$	×	✓
Transitive ε -APB	$\equiv_\varepsilon^*$		
Branching ε -Bisimulation	$\sim_\varepsilon^b$		
ε -Perturbed Bisimulation [38]	$\simeq_\varepsilon$		

all branching ε -bisimulations in $\mathcal{M}$) might themselves not be of the respective type anymore (e.g., $\approx_\varepsilon^b$ is not necessarily a branching ε -bisimulation). This contrasts the non-transitive case, where the induced bisimilarity relations are always of the respective type. We summarize the relations we consider, together with some of their properties, in Table 1.

Main Contributions. The main contributions are as follows:


1. Starting with the classic notion of ε -bisimilarity, we show tightness of a bound from [32] on the absolute difference of unbounded reachability probabilities in ε -bisimilar states (Example 3.13).
2. We introduce ε -perturbed bisimilarity ($\simeq_\varepsilon$), a notion that relates two LMCs if they have bisimilar ε -quotients à la [38], i.e., if they can be made probabilistically bisimilar by small perturbations of their transition probabilities. We show that $\simeq_\varepsilon$ is strictly finer than (transitive) ε -bisimilarity $\sim_\varepsilon^{(*)}$ (Lemma 4.6 and Theorem 4.7) and that deciding both $\simeq_\varepsilon$ and $\sim_\varepsilon^{(*)}$ is NP-complete (Theorem 4.12). Furthermore, we characterize $\simeq_\varepsilon$ in terms of transitive ε -bisimulations satisfying a *centroid property* (Theorem 4.10) and discuss some anomalies of $\simeq_\varepsilon$: the relation is not always an ε -perturbed bisimulation itself, it is not additive in ε and it can distinguish bisimilar LMCs (Proposition 4.4).
3. We define approximate versions of weak ($\approx_\varepsilon^w$) and branching probabilistic bisimilarity ($\approx_\varepsilon^b$). Our definitions can be evaluated locally and coincide with the exact notions $\approx^b$ and $\approx^w$, respectively, if $\varepsilon = 0$. We discuss how $\approx_\varepsilon^w$ and $\approx_\varepsilon^b$ are related to one another, as well as to ε -bisimilarity (Propositions 5.4 and 5.5). Moreover, we extend the bounds for reachability probabilities of Theorem 3.11 to states related by $\approx_\varepsilon^w$ and $\approx_\varepsilon^b$ (Corollary 5.9 and Proposition 5.10), and prove that deciding $\approx_\varepsilon^b$ is NP-complete (Theorem 5.11).

Together with various known results from the literature and some easy observations, our results complete the relation between several notions of approximate probabilistic bisimulation, as summarized in Figure 1.

Structure. Section 2 presents preliminaries. Section 3 considers ε -bisimulations, ε -APBs and up-to- (n, ε) -bisimulations. Section 4 introduces and analyzes ε -perturbed bisimulations. Section 5 introduces weak and branching ε -bisimulations and establishes how they relate to ε -bisimulations. Section 6 summarizes our results and points out future work.

2 Preliminaries

Distributions. $\text{Distr}(S) = \{\mu: S \rightarrow [0, 1] \mid \sum_{s \in S} \mu(s) = 1\}$ is the set of *distributions* over countable $S \neq \emptyset$. $\mu \in \text{Distr}(S)$ has *support* $\text{supp}(\mu) = \{s \in S \mid \mu(s) > 0\}$, and for $A \subseteq S$ we set $\mu(A) = \sum_{s \in A} \mu(s)$. The L_1 -distance of $\mu, \nu \in \text{Distr}(S)$ is $\|\mu - \nu\|_1 = \sum_{s \in S} |\mu(s) - \nu(s)|$.



■ **Figure 1** The relationship of different approximate probabilistic bisimulations.

Labeled Markov chains. Fix a countable set AP of *atomic propositions*. A *labeled Markov chain* (LMC) $\mathcal{M} = (S, P, s_{init}, l)$ has a countable set of states $S \neq \emptyset$, a *transition distribution function* $P: S \rightarrow \text{Distr}(S)$, a unique *initial state* s_{init} , and a *labeling function* $l: S \rightarrow 2^{AP}$. We use $\mathcal{M}$ and $\mathcal{N}$ to range over LMCs. For $s \in S$, let $L(s) = \{t \in S \mid l(s) = l(t)\}$. $\mathcal{M}$ is *finitely branching* if $|\text{supp}(P(s))| < \infty$ for all $s \in S$, and $\mathcal{M}$ is *finite* if $|S| < \infty$. The *direct sum* $\mathcal{M} \oplus \mathcal{N}$ is the LMC obtained from the disjoint union of $\mathcal{M}$ and $\mathcal{N}$. The initial state of $\mathcal{M} \oplus \mathcal{N}$ is not relevant for our purposes.

For $s, t \in S$, $P(s)(t)$ denotes the probability to move from s to t in a single step. We write $\text{Succ}(s) = \text{supp}(P(s))$ for the set of direct successors of s . $\pi = s_0 s_1 \dots \in S^\omega$ is an (*infinite*) *path* of $\mathcal{M}$ if $s_{i+1} \in \text{Succ}(s_i)$ for all $i \in \mathbb{N}$. $\pi[i] = s_i$ is the state at position i of π , and $\text{trace}(\pi) = l(s_0)l(s_1)\dots \in (2^{AP})^\omega$ is the *trace* of π . The set of infinite paths is $\text{Paths}(\mathcal{M})$. *Finite paths* $\pi = s_0 s_1 \dots s_k \in S^{k+1}$ for some $k \in \mathbb{N}$ and their traces are defined analogously.

Let $s \in S$. We consider the standard probability measure $\text{Pr}_s^{\mathcal{M}}$ on sets of infinite paths of LMCs, defined via *cylinder sets* $\text{Cyl}(\rho) = \{\pi \in \text{Paths}(\mathcal{M}) \mid \rho \text{ is a prefix of } \pi\}$ of finite paths $\rho \in S^*$. See [11] for details. For $\rho = s_0 s_1 \dots s_n$, we abbreviate $\text{Pr}_s^{\mathcal{M}}(\text{Cyl}(\rho))$ by $\text{Pr}_s^{\mathcal{M}}(\rho)$ and the measure yields $\text{Pr}_s^{\mathcal{M}}(\rho) = 0$ if $s_0 \neq s$ and $\text{Pr}_s^{\mathcal{M}}(\rho) = \prod_{j=0}^{n-1} P(s_j)(s_{j+1})$ otherwise. We write $\text{Pr}^{\mathcal{M}}$ for $\text{Pr}_{s_{init}}^{\mathcal{M}}$ and drop the superscript if $\mathcal{M}$ is clear from the context. Given a set of finite traces $T \subseteq (2^{AP})^{k+1}$ for some $k \in \mathbb{N}$, $\text{Pr}_s(T)$ denotes the probability to follow, when starting in s , a finite path $\pi = s s_1 \dots s_k$ with $\text{trace}(\pi) \in T$. $\mathbb{E}_s^{\mathcal{M}}(X)$ or simply $\mathbb{E}_s(X)$ denotes the *expected value* of a random variable X on $\text{Paths}(\mathcal{M})$ w.r.t. $\text{Pr}_s^{\mathcal{M}}$.

LTL. A popular logic for the specification of desired properties of LMCs is the *linear temporal logic* (LTL) which can be used to, e.g., specify properties such as reachability, safety or liveness [44, 11]. For $a \in AP$, LTL formulas are formed w.r.t. the grammar

$$\varphi ::= \text{true} \mid a \mid \neg\varphi \mid \varphi_1 \vee \varphi_2 \mid \bigcirc\varphi \mid \varphi_1 \mathbf{U} \varphi_2.$$

Here, $\bigcirc$ is the *next* operator, so $\pi \in \text{Paths}(\mathcal{M})$ satisfies $\bigcirc\varphi$ iff φ is true in $\pi[1]$. For the *until* operator $\mathbf{U}$, π satisfies $\varphi_1 \mathbf{U} \varphi_2$ iff, alongside π , φ_1 holds until φ_2 is true. As syntactic sugar we define the *reachability* operator $\Diamond\varphi \equiv \text{true} \mathbf{U} \varphi$ and the *always* operator $\Box\varphi \equiv \neg\Diamond\neg\varphi$.

For $B, C \subseteq S$ and $s \in S$, $\text{Pr}_s(\text{BUC})$ is the probability to reach a state in C via a (finite) path from s that only consists of states in B . Moreover, $\text{Pr}_s(\Diamond^{\leq n}\varphi)$ denotes the probability to reach a state satisfying φ from s in at most $n \in \mathbb{N}$ steps. For details on LTL, see [11].

Relations. Given a relation $R \subseteq S \times S$ and an $A \subseteq S$, $R(A) = \{t \in S \mid \exists s \in A: (s, t) \in R\}$ is the *image of A under R*. If R is reflexive then $A \subseteq R(A)$, and A is called *R-closed* if $R(A) \subseteq A$. When R is an equivalence, i.e., when it is reflexive, symmetric and transitive, the *equivalence class* of $s \in S$ is $[s]_R = R(\{s\}) = \{t \in S \mid (s, t) \in R\}$, and we set $S/R = \{[s]_R \mid s \in S\}$. For an equivalence R , the *R-closed sets* are precisely the (unions of) *R* equivalence classes.

Bisimulation. An equivalence $R \subseteq S \times S$ is a (*probabilistic*) *bisimulation* on $\mathcal{M}$ if for all $(s, t) \in R$ and all R -equivalence classes C it holds that $l(s) = l(t)$ and $P(s)(C) = P(t)(C)$. States $s, t \in S$ are (*probabilistically*) *bisimilar*, written $s \sim^{\mathcal{M}} t$ or simply $s \sim t$, if there is a bisimulation R on $\mathcal{M}$ with $(s, t) \in R$. We call two LMCs $\mathcal{M}, \mathcal{N}$ *bisimilar*, written $\mathcal{M} \sim \mathcal{N}$, if $s_{init}^{\mathcal{M}} \sim s_{init}^{\mathcal{N}}$ in $\mathcal{M} \oplus \mathcal{N}$. An alternative characterization of bisimulations can be found in, e.g., [23, 25, 27, 14]: an equivalence R is a bisimulation iff for all $(s, t) \in R$ and all R -closed sets $A \subseteq S$ it holds that $l(s) = l(t)$ and $P(s)(A) = P(t)(A)$. The (*probabilistic bisimulation*) *quotient* of $\mathcal{M}$ is the LMC $\mathcal{M}/\sim = (S/\sim, P_\sim, [s_{init}]_\sim, l_\sim)$ with $l_\sim([s]_\sim) = l(s)$, and $P_\sim([s]_\sim)([t]_\sim) = \sum_{q \in [t]_\sim} P(s)(q)$ for all $[s]_\sim, [t]_\sim \in S/\sim$. It holds that $\mathcal{M} \sim \mathcal{M}/\sim$. An important result is that *bisimilarity* $\sim$ preserves the satisfaction of PCTL* state formulas [8].

We also consider *weak* and *branching probabilistic bisimulations* [43, 51, 10, 35]. An equivalence R is a *weak probabilistic bisimulation* if, for all $(s, t) \in R$ and all R -equivalence classes $C \neq [s]_R = [t]_R$, it holds that $l(s) = l(t)$ and $\Pr_s(L(s)UC) = \Pr_t(L(t)UC)$. R is a *branching probabilistic bisimulation* if, instead of the second condition in the previous definition, $\Pr_s([s]_R UC) = \Pr_t([t]_R UC)$ holds. *Weak probabilistic bisimilarity* $\approx^w$ and *branching probabilistic bisimilarity* $\approx^b$ are defined like $\sim$, and lifted to LMCs in the same way.

3 ε -Bisimulation, ε -APB and Up-To- (n, ε) -Bisimulation

If not specified otherwise, we always assume $\varepsilon \in [0, 1]$ and $\mathcal{M} = (S, P, s_{init}, l)$ to be finitely branching. This section summarizes various notions of approximate probabilistic bisimulation from the literature. We first provide their formal definitions and discuss how the notions interrelate. Afterwards, in Section 3.2, we present some logical preservation results.

3.1 Definitions and Interrelation

We start with the seminal notion of ε -*bisimulations* of Desharnais *et al.* [25]. While originally introduced for *labeled Markov processes* [21, 22], ε -bisimulations were later adapted to other models like LMCs [14, 38] or Segala's *probabilistic automata* [45, 47].

► **Definition 3.1** ([25, 14]). A reflexive¹ and symmetric relation $R \subseteq S \times S$ is an ε -bisimulation if for all $(s, t) \in R$ and any $A \subseteq S$ it holds that


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

States s, t are ε -bisimilar, denoted $s \sim_\varepsilon t$, if there is an ε -bisimulation R with $(s, t) \in R$. LMCs $\mathcal{M}, \mathcal{N}$ are ε -bisimilar, denoted $\mathcal{M} \sim_\varepsilon \mathcal{N}$, if $s_{init}^{\mathcal{M}} \sim_\varepsilon s_{init}^{\mathcal{N}}$ in $\mathcal{M} \oplus \mathcal{N}$.


Intuitively, $s \sim_\varepsilon t$ if both states can mimic the other's transition probabilities to any $A \subseteq S$ by transitioning to the (potentially bigger) set $\sim_\varepsilon(A)$ with a probability that is smaller by at most ε than the original one. The parameter ε describes how much the behavior of related states may differ: for ε close to 1 more states can be related, while for $\varepsilon \approx 0$ related states behave almost equivalently. In the extreme case of $\varepsilon = 0$, we have $\sim_0 = \sim$ [25, 14].

Instead of being transitive, ε -bisimulations are *additive* in their tolerances: $s \sim_{\varepsilon_1} t$ and $t \sim_{\varepsilon_2} u$ implies $s \sim_{\varepsilon'} u$ for some $0 \leq \varepsilon' \leq \min\{1, \varepsilon_1 + \varepsilon_2\}$ [25]. As the next example suggests, transitivity is not always desirable for ε -bisimulations if $\varepsilon > 0$.

¹ In contrast to [25, 14] we require reflexivity of ε -bisimulations. This is a rather natural assumption (a state should always simulate itself) that does not affect $\sim_\varepsilon$.



■ **Figure 2** The LMC used in Example 3.2.



■ **Figure 3** The LMC used in Example 3.5, adapted from [14].

► **Example 3.2.** Let $\varepsilon = \frac{1}{n}$ for $n \geq 1$ and consider the LMC of Figure 2. There, the reflexive and symmetric closure of $R = \{(s_i, s_{i+1}) \mid 0 \leq i \leq n-1\}$ is an ε -bisimulation. Hence, s_0 and s_n are related by a chain of ε -bisimilar states, even though they behave completely different: s_0 transitions to the $\{a\}$ -labeled state x with probability 0, s_n does so with probability 1.

Desharnais *et al.* [25] describe how to check condition (ii) of Definition 3.1 in terms of the values of maximum flows in specific flow networks à la [46, 9]. Their result is well-suited for algorithmic purposes, but is restricted to finite models. Equivalently, one can characterize ε -bisimulations by the existence of *weight functions* $\Delta: S \rightarrow \text{Distr}(S)$ that describe how to split the successor probabilities of related states. This formulation is used in, e.g., [47, 38, 39, 31]. The following lemma provides one such characterization that is proved using ideas from [14, 15] and a technical measure-theoretic statement from [16].

► **Lemma 3.3.** *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 $\Delta: \text{Succ}(s) \rightarrow \text{Distr}(\text{Succ}(t))$ such that*

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

Intuitively, Lemma 3.3 tells us that, if $s \sim_\varepsilon t$, the successors s' of s can be mapped to distributions $\Delta(s')$, i.e., convex combinations, of successors of t . More precisely, it shows that (i) if we move from s to a successor s' with probability $P(s)(s')$ and, afterwards, from s' to a successor t' of t with probability $\Delta(s')(t')$, then we reach t' with probability $P(t)(t')$, and that (ii) the overall probability that the states s' and t' are ε -bisimilar is at least $1 - \varepsilon$.

A second notion of approximate probabilistic bisimulation are ε -APBs, which stands short for *approximate probabilistic bisimulations with precision ε* [27, 1, 2]. In contrast to ε -bisimulations, where the differences in transition probabilities of related states are bounded w.r.t. all subsets $A \subseteq S$, an ε -APB R only requires a difference of at most ε for the probabilities of related states to transition to R -closed subsets of S .

► **Definition 3.4** ([27]). *A reflexive and symmetric relation $R \subseteq S \times S$ is an ε -APB if for all $(s, t) \in R$ and any R -closed set $A \subseteq S$ it holds that*

$$(i) \ l(s) = l(t) \quad \text{and} \quad (ii) \ |P(s)(A) - P(t)(A)| \leq \varepsilon.$$

We write $s \equiv_\varepsilon t$ if s and t are related by any ε -APB, and $\mathcal{M} \equiv_\varepsilon \mathcal{N}$ if $s_{init}^\mathcal{M} \equiv_\varepsilon s_{init}^\mathcal{N}$ in $\mathcal{M} \oplus \mathcal{N}$.

Like $\sim_\varepsilon$, ε -APBs are additive in their tolerances, and we have $\equiv_0 = \sim = \sim_0$ [25, 27].

► **Example 3.5** ([14]). Let $\varepsilon \in (0, 1]$ and $n = \lceil \frac{1}{\varepsilon} \rceil \in \mathbb{N}$. Consider $\mathcal{M}$ as in Figure 3, and let R be the reflexive and symmetric closure of $\{(s, t), (x, x), (y, y)\} \cup \{(u_i, u_{i+1}) \mid 0 \leq i \leq n-1\}$. The R -closed sets in $\mathcal{M}$ are $\{s, t\}, \{x\}, \{y\}, \{u_i \mid 0 \leq i \leq n\}$ and their unions. For all $(p, q) \in R$ and R -closed sets A it holds that $|P(p)(A) - P(q)(A)| \leq \varepsilon$, so R is an ε -APB.

Example 3.5 illustrates that the use of ε -APBs as a notion that relates states with almost equivalent behavior is questionable: even though states s and t in Figure 3 are related by $\equiv_\varepsilon$, they behave completely different. This is caused by the set $\{u_0, \dots, u_n\}$ of (unreachable) states being R -closed, which in turn allows to relate s and t by the relation R from the example. Such an anomaly cannot occur for $\sim_\varepsilon$, and we in fact have $s \sim_\varepsilon t$ in Figure 3 for every $\varepsilon \in (0, 1)$. In particular, this shows that $\sim_\varepsilon$ can be strictly finer than $\equiv_\varepsilon$ for $\varepsilon \in (0, 1)$.

Lastly, we introduce *up-to- (n, ε) -bisimulations* [25, 13], which are relations that require the behaviors of related states to differ by at most ε for at least n steps.

► **Definition 3.6** ([25, 13]). The up-to- (n, ε) -bisimulation $\sim_\varepsilon^n \subseteq S \times S$ is inductively defined on n via $s \sim_\varepsilon^0 t$ for all $s, t \in S$ and, for $n \geq 0$, $s \sim_\varepsilon^{n+1} t$ iff for all $A \subseteq S$

$$(i) \ l(s) = l(t), \quad (ii) \ P(s)(A) \leq P(t)(\sim_\varepsilon^n(A)) + \varepsilon \quad \text{and} \quad (iii) \ P(t)(A) \leq P(s)(\sim_\varepsilon^n(A)) + \varepsilon.$$

States s, t are (n, ε) -bisimilar if $s \sim_\varepsilon^n t$, and the notion is lifted to LMCs as usual. Similar to $\sim_\varepsilon$ and $\equiv_\varepsilon$, $\sim_\varepsilon^n$ is reflexive and symmetric, but not transitive. Instead, it is additive in the tolerances and monotonic in n and ε , i.e., for $n \geq n'$ and $\varepsilon \leq \varepsilon'$, $s \sim_\varepsilon^n t$ implies $s \sim_{\varepsilon'}^{n'} t$ [13].

It is clear that $s \sim_\varepsilon^n t$ for a fixed n does not necessarily imply $s \sim_\varepsilon t$ or $s \equiv_\varepsilon t$, as (n, ε) -bisimilarity only restricts the behavior of related states for n steps. However, considering the limit $n \rightarrow \infty$ makes $\sim_\varepsilon$ and $\sim_\varepsilon^n$ coincide, i.e., $s \sim_\varepsilon t$ iff $s \sim_\varepsilon^n t$ for all $n \in \mathbb{N}$ [25].

We now make precise the relationship between ε -APBs and up-to- (n, ε) -bisimulations.

► **Proposition 3.7.** If $\varepsilon \in (0, 1)$, $s \equiv_\varepsilon t$ implies $s \sim_\varepsilon^n t$ if $n \leq 2$, but not necessarily if $n \geq 3$.

3.2 Preservation of Logical Properties


A key application of exact probabilistic bisimilarity $\sim$ is the use of quotients $\mathcal{Q} = \mathcal{M}/\sim$ to speed up PCTL* model checking [37, 36]. As abstractions built by grouping states related by approximate probabilistic bisimulations can be smaller than $\mathcal{Q}$ [27], these notions might prove useful to combat the *state space explosion problem* of model checking [37, 11, 36]. It is hence of interest to see which logical properties these relations preserve.

We start by considering $\sim_\varepsilon$. As shown by Bian and Abate [14], ε -bisimilarity induces bounds on the absolute difference of satisfaction probabilities of *finite horizon* properties, i.e., of properties that only depend on traces of finite length, in related states.


► **Theorem 3.8** ([14]). Let $s \sim_\varepsilon t$, $k \in \mathbb{N}$ and $T \subseteq (2^{AP})^{k+1}$ a set of traces of length $k+1$. Then $|\Pr_s(T) - \Pr_t(T)| \leq 1 - (1 - \varepsilon)^k$.

Since any finite horizon LTL formula coincides with a set of traces of finite length, Theorem 3.8 in particular bounds the satisfaction probabilities of such formulas in ε -bisimilar states. Furthermore, as argued in [14] and the next example, this bound is tight.

► **Example 3.9.** Consider Figure 4. For $i \in \{0, \dots, n\}$, let $l(s_i) = l(t_i) = a_i$ for pairwise distinct a_i , $l(G_1) = g = l(G_2)$ and $l(F) = f$ for some $f \neq g$. Then $s_0 \sim_\varepsilon t_0$ and the upper bound of Theorem 3.8 is met exactly: $|\Pr_{s_0}(\Diamond^{\leq n+1} g) - \Pr_{t_0}(\Diamond^{\leq n+1} g)| = 1 - (1 - \varepsilon)^{n+1}$.



■ **Figure 4** An LMC in which $s_0 \sim_\varepsilon t_0$ and $|\Pr_{s_0}(\diamond^{\leq n+1}g) - \Pr_{t_0}(\diamond^{\leq n+1}g)| = 1 - (1 - \varepsilon)^{n+1}$.



■ **Figure 5** The LMC used in Example 3.10.

A disadvantage of the bound provided in Theorem 3.8 is, however, that it rapidly converges to 1 for increasing k and is thus not suitable when reasoning about long (or infinite) time horizons. In fact, it is the case that – without further assumptions – even simple unbounded reachability probabilities in ε -bisimilar states can strongly deviate.

► **Example 3.10.** Let $\varepsilon \geq 0$. The states s_0 , s_1 , and s_2 in Figure 5 are pairwise ε -bisimilar. However, if $\varepsilon > 0$, we have $\Pr_{s_0}(\diamond g) = \frac{1}{2}$, $\Pr_{s_1}(\diamond g) = 0$, and $\Pr_{s_2}(\diamond g) = 1$.

The difference in reachability probabilities observed in the last example is caused by $\sim_\varepsilon$ relating states that are able to reach a goal state g with positive probability to those that can not reach g at all. One way to avoid this issue is to require that states from which g is not reachable are labeled with a distinct label f . The existence of such a label f is a rather natural assumption, as a typical preprocessing step when computing reachability probabilities is to identify the states from which no goal state is reachable, i.e., to identify the states we assume to be labeled with f [11]. A result in the spirit of Theorem 3.8 that deals with *unbounded* reachability properties can then be obtained as follows.

► **Theorem 3.11** ([31, 32]). *Let some states in $\mathcal{M}$ be labeled with g , and let exactly the states that cannot reach a g -labeled state be labeled with f . Further, let $s \sim_\varepsilon t$, and let N be the random variable that counts the number of steps until reaching a g - or f -labeled state. Then,*

$$|\Pr_s(\diamond g) - \Pr_t(\diamond g)| \leq \varepsilon \cdot \mathbb{E}_s(N).$$


► **Remark 3.12.** A result similar to Theorem 3.11 is derived by Haesaert *et al.* in [31, 32] in the context of policy synthesis in control theory. In fact, their result is more general, as it considers all properties that can be described as the language of a deterministic finite automaton. These properties include, among others, the *syntactically co-safe* LTL formulas [40], which form a fragment of LTL built according to the grammar

$$\varphi ::= \text{true} \mid a \mid \neg a \mid \varphi_1 \vee \varphi_2 \mid \varphi_1 \wedge \varphi_2 \mid \bigcirc \varphi \mid \varphi_1 \mathbf{U} \varphi_2,$$

where $a \in AP$. As unbounded reachability $\diamond g$ is a syntactically co-safe LTL formula, the results of [31, 32] extend the bound in Theorem 3.11 to a broader class of properties.

Next, we show that the bound described in Theorem 3.11 is actually tight.

► **Example 3.13.** Let $p \in (0, 1)$, $\varepsilon < \frac{p}{2}$ and consider Figure 6, where $s \sim_\varepsilon t$. There, $\Pr_s(\diamond g) = \frac{1}{2}$, $\Pr_t(\diamond g) = \frac{1}{2} - \frac{\varepsilon}{p}$ and $\mathbb{E}_s(N) = \mathbb{E}_t(N) = \frac{1}{p}$. Hence, the bound in Theorem 3.11 is met exactly: $|\Pr_s(\diamond g) - \Pr_t(\diamond g)| = \frac{\varepsilon}{p} = \varepsilon \cdot \mathbb{E}_s(N) = \varepsilon \cdot \mathbb{E}_t(N)$.



■ **Figure 6** The LMC used in Example 3.13. The states s and t are ε -bisimilar.

Regarding ε -APBs and up-to- (n, ε) -bisimilarity, some preservation results w.r.t. the (approximate or robust) satisfaction of PCTL state-formulas can be found in the literature. Since, as we have seen in Section 3.1, $s \sim_\varepsilon t$ implies both $s \equiv_\varepsilon t$ and $s \sim_\varepsilon^n t$ for any $n \in \mathbb{N}$ and any two states s, t , the following results also hold for ε -bisimilar states.

An important property of ε -APBs is that related states satisfy the same ε -robust PCTL state formulas Φ_{robust} [27], i.e., that $s \equiv_\varepsilon t$ implies that $s \models \Phi_{\text{robust}}$ iff $t \models \Phi_{\text{robust}}$, where $\models$ is the usual PCTL satisfaction relation [11]. Intuitively, Φ_{robust} is ε -robust if for all subformulas ϕ of Φ_{robust} and all $s \in S$ either a strengthened version of ϕ , obtained by making ϕ 's probability thresholds harder to meet, holds in s , or even relaxing ϕ 's probability thresholds is not sufficient to ensure that s satisfies ϕ . For details, see [27].

Furthermore, it was shown in [13] that (n, ε) -bisimilar states approximately satisfy the same bounded PCTL state formulas. The fragment of PCTL considered does not allow unbounded until, and requires all until operator appearing in a formula to have the same time bound $k \in \mathbb{N}$. Under these assumptions, the precision of the approximation of satisfaction probabilities between (n, ε) -bisimilar states is proved to depend linearly on the parameters n and ε , as well as the common step bound k of the until operators. For details, see [13].

4 ε -Perturbed Bisimulation

In this section we consider finite LMCs. In [38], Kiefer and Tang define the notion of ε -quotients for $\varepsilon \geq 0$. Their goal is to construct, from a given *perturbed* LMC $\mathcal{M}'$, an abstraction that is as close as possible to the exact bisimulation quotient of an unknown, unperturbed LMC $\mathcal{M}$ corresponding to $\mathcal{M}'$. This inspires us to introduce ε -perturbed bisimulations, which relate two LMCs iff they can be made probabilistically bisimilar by small perturbations of their transition probabilities. Since we require ε -perturbed bisimulations to be equivalences, these relations are well-suited for the construction of quotients of a given model.

Like the ε -quotients of [38], we base our definition on ε -perturbations of LMCs.


► **Definition 4.1** ([38]). $\mathcal{M}' = (S, P', s_{\text{init}}, l)$ is an ε -perturbation of $\mathcal{M} = (S, P, s_{\text{init}}, l)$ if $\|P(s) - P'(s)\|_1 \leq \varepsilon$ for all $s \in S$.

$\mathcal{M}$ and any of its ε -perturbations $\mathcal{M}'$ have the same state space and labeling, and we often write $S' = \{s' \mid s \in S\}$ for the state space of $\mathcal{M}'$. Hence, $\mathcal{M}$ and $\mathcal{M}'$ only differ in their transition distribution functions. However, $\mathcal{M}'$ does not need to preserve the structure of $\mathcal{M}$, i.e., there can be transitions in $\mathcal{M}$ that have probability 0 in $\mathcal{M}'$ and vice versa. As the next lemma shows, the total probability mass of these transitions cannot exceed $\frac{\varepsilon}{2}$.


► **Lemma 4.2.** For all $s \in S$ and $A \subseteq S$ it holds that $|P(s)(A) - P'(s')(A')| \leq \frac{\varepsilon}{2}$.

We now define the novel notion of ε -perturbed bisimulation.

► **Definition 4.3.** An equivalence $R \subseteq S \times S$ is called an ε -perturbed bisimulation on $\mathcal{M}$ if there is an ε -perturbation $\mathcal{M}'$ of $\mathcal{M}$ such that R is a bisimulation on $\mathcal{M}'$. Two states $s, t \in S$ are ε -perturbed bisimilar, denoted $s \simeq_\varepsilon t$, if they are related by some ε -perturbed bisimulation. Given LMCs $\mathcal{M}$ and $\mathcal{N}$, then $\mathcal{M} \simeq_\varepsilon \mathcal{N}$ if $s_{\text{init}}^\mathcal{M} \simeq_\varepsilon s_{\text{init}}^\mathcal{N}$ in $\mathcal{M} \oplus \mathcal{N}$.



■ **Figure 7** An LMC in which there is no unique maximal transitive ε -bisimulation.




■ **Figure 8** Two LMCs $\mathcal{M}_s$ and $\mathcal{M}_t$ (left) with initial states s and t , respectively, and $\mathcal{Q} = \mathcal{M}_s/\sim$ (right), demonstrating that $\simeq_\varepsilon$ and $\sim_\varepsilon^*$ can differentiate bisimilar models and are not additive.

In the terminology of [38], $\mathcal{M} \simeq_\varepsilon \mathcal{N}$ iff $\mathcal{M}$ and $\mathcal{N}$ have bisimilar ε -perturbations iff there are bisimilar ε -quotients of $\mathcal{M}$ and $\mathcal{N}$. If all states of $\mathcal{M}$ and $\mathcal{N}$ are reachable, even the stronger characterization $\mathcal{M} \simeq_\varepsilon \mathcal{N}$ iff $\mathcal{M}$ and $\mathcal{N}$ have *isomorphic* ε -perturbations iff there are *isomorphic* ε -quotients of $\mathcal{M}$ and $\mathcal{N}$ holds. Since the unique 0-perturbation of any LMC is the LMC itself, $\mathcal{M} \simeq_0 \mathcal{N}$ iff $\mathcal{M} \sim \mathcal{N}$. Moreover, $\simeq_\varepsilon$ is symmetric and reflexive, but not always transitive, which implies that $\simeq_\varepsilon$ is not necessarily an ε -perturbed bisimulation itself.


Let $s \sim_\varepsilon^* t$ denote that states s and t are related by a *transitive* ε -bisimulation. We remark that both $\simeq_\varepsilon$ and $\sim_\varepsilon^*$ are definitions in the spirit of a notion called *ε -lumpability* (or *quasi-lumpability*), which describes that a LMC can be made exactly lumpable w.r.t. a given equivalence by slight changes (up to ε in each value) of its transition probabilities [17, 29, 28]. In contrast to the non-transitive case, any transitive ε -APB is also an ε -bisimulation.

The requirement of transitivity comes with the downside that there is not always a *unique* largest transitive ε -bisimulation: in Figure 7, no transitive ε -bisimulation R can contain both (s, t) and (t, u) , as otherwise also $(s, u) \in R$ must hold. However, $s \sim_\varepsilon^* t$ and $t \sim_\varepsilon^* u$ as $R_1 = \{\{s, t\}, \{u\}, \{x\}, \{y\}\}$ and $R_2 = \{\{s\}, \{t, u\}, \{x\}, \{y\}\}$ are transitive ε -bisimulations. Hence, the union of all transitive ε -bisimulations in a given model is thus not always a transitive ε -bisimulation itself. This is different than in the non-transitive case, where $\sim_\varepsilon$ is always an ε -bisimulation [25]. Since $s \simeq_\varepsilon t$ and $t \simeq_\varepsilon u$ but $s \not\simeq_\varepsilon u$ in Figure 7, it follows that there is also not always a unique largest ε -perturbed bisimulation.

Now consider, for $\varepsilon < \frac{1}{4}$, the LMCs $\mathcal{M}_s$ and $\mathcal{M}_t$ on the left of Figure 8, with initial states s and t , respectively. In both models, $\sim$ is the finest equivalence that contains (u_2, u_3) . Let R_1 be the finest equivalence that contains $(s, t), (u_1, u_2), (u_3, u_4)$, and let R_2 be the one that contains $(s, t), (u_1, u_3), (u_2, u_4)$. Both R_1 and R_2 are transitive ε -bisimulations, and since $u_1 \not\sim_\varepsilon u_4$ no other transitive ε -bisimulation can contain (s, t) . Hence, no such relation contains (u_2, u_3) . Let $\mathcal{Q} = \mathcal{M}_s/\sim$ be as on the right of the figure. Then $\mathcal{M}_s \sim \mathcal{Q}$ and $\mathcal{M}_s \simeq_\varepsilon \mathcal{M}_t$ as, e.g., the ε -perturbations $\mathcal{M}'_s$ and $\mathcal{M}'_t$ that enforce $u'_1 \sim u'_2$ and $u'_3 \sim u'_4$ and are otherwise unchanged are bisimilar. However, there are no bisimilar ε -perturbations of $\mathcal{M}_t$ and $\mathcal{Q}$, i.e., $\mathcal{M}_t \not\simeq_\varepsilon \mathcal{Q}$. Since $\simeq_0 = \sim$ this observation additionally yields that $\simeq_\varepsilon$ cannot be additive, as otherwise $\mathcal{M}_s \sim \mathcal{Q}$ and $\mathcal{M}_s \simeq_\varepsilon \mathcal{M}_t$ would have to imply $\mathcal{M}_t \simeq_\varepsilon \mathcal{Q}$. All in all, this leads to the following result.



■ **Figure 9** An LMC that demonstrates that $\simeq_\varepsilon$ is strictly finer than $\sim_\varepsilon^*$.



■ **Figure 10** The LMCs $\mathcal{M}_2$ (left) and $\mathcal{N}_2$ (right), as used in the proof of Theorem 4.8.

► **Proposition 4.4.** *The relation $\simeq_\varepsilon$ is not additive in the tolerances and can distinguish bisimilar LMCs in the following sense: there are LMCs $\mathcal{M}_1, \mathcal{M}_2$ and $\mathcal{N}$ such that $\mathcal{M}_1 \sim \mathcal{M}_2$ and $\mathcal{M}_1 \simeq_\varepsilon \mathcal{N}$, but $\mathcal{M}_2 \not\simeq_\varepsilon \mathcal{N}$.*

This behavior of $\simeq_\varepsilon$ is in contrast to, e.g., $\sim_\varepsilon$, as $s_1 \sim s_2$ and $s_1 \sim_\varepsilon t$ always implies $s_2 \sim_\varepsilon t$. In particular, the non-additivity does not hinge on the existence of bisimilar states in the model. To see this consider, e.g., slight perturbations $\mathcal{M}'_s$ and $\mathcal{M}'_t$ of the LMCs on the left of Figure 8, where for some $\delta < \varepsilon$ we set $P(u_2)(v) = \frac{1}{2} - \varepsilon - \delta$ and $P(u_2)(w) = \frac{1}{2} + \varepsilon + \delta$, and leave the rest of the models unchanged. Then $u_2 \sim u_3$ in $\mathcal{M}'_s$ and $\mathcal{M}'_t$, but still $\mathcal{M}'_s \simeq_\delta \mathcal{Q}$ and $\mathcal{M}'_s \simeq_\varepsilon \mathcal{M}'_t$ while $\mathcal{M}'_t \not\simeq_{\varepsilon+\delta} \mathcal{Q}$, where $\mathcal{Q}$ is again the (unperturbed) LMC on the right of the figure. Similar results hold for $\sim_\varepsilon^*$, as $\mathcal{M}_s \sim_\varepsilon^* \mathcal{M}_t$ and $\mathcal{M}_s \sim \mathcal{Q}$, but $\mathcal{M}_t \not\sim_\varepsilon^* \mathcal{Q}$.

We now discuss how $\simeq_\varepsilon$ relates to $\sim_\varepsilon$ and $\sim_\varepsilon^*$, starting with the direction from left to right. From [39] it follows directly that $\mathcal{M} \simeq_\varepsilon \mathcal{N}$ implies $\mathcal{M} \sim_\varepsilon \mathcal{N}$. As we show next, the claim also holds when considering the stronger requirement of *transitive* ε -bisimilarity.

► **Lemma 4.5.** *$\mathcal{M} \simeq_\varepsilon \mathcal{N}$ implies $\mathcal{M} \sim_\varepsilon^* \mathcal{N}$.*

It is thus possible to transfer known results for $\sim_\varepsilon$ like, e.g., the preservation of approximate satisfaction of bounded PCTL state formulas [13], the exact preservation of ε -robust PCTL [27], or the bounds on finite horizon [14] and syntactically co-safe [31, 32] LTL satisfaction probabilities to ε -perturbed bisimilar LMCs.

Regarding the reverse implication, consider Figure 9. There, the finest equivalence that relates u_1, u_2 and u_3 and contains (s, t) is a transitive ε -bisimulation. However, there is no ε -perturbation of the LMC in which s and t are bisimilar. Hence, $s \sim_\varepsilon^* t$, but $s \not\simeq_\varepsilon t$.

► **Lemma 4.6.** *$\simeq_\varepsilon$ is strictly finer than $\sim_\varepsilon^*$.*

In fact, ε -bisimilarity is not even guaranteed to imply δ -perturbed bisimilarity if $\varepsilon \ll \delta$, or if the Markov chains in question are graph-isomorphic.

► **Theorem 4.7.** *Let $\varepsilon \in (0, \frac{1}{4}]$. There are LMCs $\mathcal{M}$ and $\mathcal{N}$ with $\mathcal{M} \sim_\varepsilon \mathcal{N}$ but $\mathcal{M} \not\sim_{\frac{1}{4}} \mathcal{N}$.*

► **Theorem 4.8.** *There is a family $\mathcal{F} = \{(\mathcal{M}_n, \mathcal{N}_n) \mid n \in \mathbb{N}_{\geq 1}\}$ of pairs of finite LMCs such that, for all $n \in \mathbb{N}_{\geq 1}$ and $\varepsilon \in \left(0, \frac{1}{n \cdot (n+1)^2}\right]$, $\mathcal{M}_n$ and $\mathcal{N}_n$ are graph-isomorphic and ε -bisimilar, but $\mathcal{M}_n \not\sim_{\delta} \mathcal{N}_n$ for any $\delta < n\varepsilon$.*

Proof sketch. We sketch the case $n = 2$, with $\mathcal{M}_2$ and $\mathcal{N}_2$ as in Figure 10, $\varepsilon \in (0, \frac{1}{18}]$ and $\sim_{\varepsilon}$ the symmetric and reflexive closure of $\{(s, t), (s_1, t_1), (s_1, t_2), (s_2, t_2), (s_2, t_3), (s_3, t_3), (x, y)\}$. Any bisimilar perturbations $\mathcal{M}'_2$ and $\mathcal{N}'_2$ must ensure $s' \sim t'$. The smallest (w.r.t. the required tolerances) perturbations that achieve this make s'_2, s'_3 and t'_3 , as well as s'_1, t'_1 and t'_2 , bisimilar, and set the total probability mass from s' (resp. t') to reach these (sets of) state(s) to $\frac{1}{2}$ each. But this requires a perturbation by at least $\delta = \frac{1}{9} \geq 2\varepsilon$. ◀

► **Remark 4.9.** Theorems 4.7 and 4.8 seem to resemble results of [38, 39]. There, an LMC is presented in which a specific order of merging ε -bisimilar states results in an approximate quotient that requires tolerance $\geq \frac{1}{4}$, and a family of LMCs is provided [39, Thm. 12] in which merging ε -bisimilar states yields an approximate quotient that requires tolerance $\geq n\varepsilon$. Our results differ in that we consider the existence of bisimilar ε -perturbations of two LMCs, and in that we show that no suitable *smaller* tolerance exists.

The observation that $\simeq_{\varepsilon}$ is strictly finer than $\sim_{\varepsilon}$ (and even $\sim_{\varepsilon}^*$) raises the question whether there are logical properties which are preserved under $\simeq_{\varepsilon}$, but not necessarily under $\sim_{\varepsilon}^{(*)}$. It is future work to make this precise. Here, we note that the bound for reachability probabilities from Theorem 3.11 remains tight under $\simeq_{\varepsilon}$: the LMCs $\mathcal{M}$ and $\mathcal{N}$ in Figure 6 satisfy $\mathcal{M} \simeq_{\varepsilon} \mathcal{N}$, but the bounds are tight by Example 3.13.

The following theorem characterizes $\simeq_{\varepsilon}$ in terms of transitive ε -bisimulations that satisfy an additional *centroid property* specified as in Equation (1) below.

► **Theorem 4.10.** *The following statements are equivalent:*

- (i) $\mathcal{M} \simeq_{\varepsilon} \mathcal{N}$.
- (ii) *There is an ε -perturbation of $\mathcal{M} \oplus \mathcal{N}$ in which $s_{init}^{\mathcal{M}} \sim s_{init}^{\mathcal{N}}$.*
- (iii) *There is a transitive ε -bisimulation R on $\mathcal{M} \oplus \mathcal{N}$ with $(s_{init}^{\mathcal{M}}, s_{init}^{\mathcal{N}}) \in R$ such that for each $A \in S/R$, where S is the disjoint union of $S^{\mathcal{M}}$ and $S^{\mathcal{N}}$, there is a $P_A^* \in \text{Distr}(S/R)$ with*

$$|P(s)(C) - P_A^*(C)| \leq \frac{\varepsilon}{2} \text{ for all } s \in A \text{ and all } R\text{-closed sets } C. \quad (1)$$

From the next lemma it follows immediately that, for a given equivalence $R \subseteq S \times S$, the centroid property in Equation (1) can be checked efficiently.

► **Lemma 4.11.** *For a finite set X and $\mu_1, \dots, \mu_k \in \text{Distr}(X)$, the following are equivalent:*

- (i) *There exists $\mu^* \in \text{Distr}(X)$ with $|\mu_l(B) - \mu^*(B)| \leq \frac{\varepsilon}{2}$ for all $l \in \{1, \dots, k\}$ and $B \subseteq X$.*
- (ii) *There exists $\mu \in \text{Distr}(X)$ with $\|\mu_l - \mu\|_1 \leq \varepsilon$ for all $l \in \{1, \dots, k\}$.*
- (iii) *The following linear constraint system over non-negative variables $\delta_{l,i}$ and x_i for $l \in \{1, \dots, k\}$ and $i \in X$ is solvable:*

$$\sum_{i \in X} x_i = 1 \quad \text{and} \quad x_i - \mu_l(i) \leq \delta_{l,i} \quad \text{and} \quad \mu_l(i) - x_i \leq \delta_{l,i} \quad \text{and} \quad \sum_{i \in X} \delta_{l,i} \leq \varepsilon.$$

The equivalence to (iii) further implies that $\mu^ = \mu$ can be computed in polynomial time.*

However, as we show next, for given $\mathcal{M}, \mathcal{N}$ and ε it is NP-complete to decide if $\mathcal{M} \simeq_{\varepsilon} \mathcal{N}$ and if $\mathcal{M} \sim_{\varepsilon}^* \mathcal{N}$. This stands in contrast to the polynomial time computability of $\sim_{\varepsilon}$ [25], which is possible in $\mathcal{O}(|S|^7)$ by iteratively solving maximum flow problems. Our proofs are inspired by [39, Thm. 1], which proves that deciding if a LMC has an ε -quotient with a fixed number of states is NP-complete.

► **Theorem 4.12.** *For given finite LMCs $\mathcal{M}$ and $\mathcal{N}$ and given $\varepsilon \in (0, 1]$, it is NP-complete to decide if (i) $\mathcal{M} \simeq_\varepsilon \mathcal{N}$ and to decide if (ii) $\mathcal{M} \sim_\varepsilon^* \mathcal{N}$.*

Nevertheless, one can check in polynomial time if a *given* equivalence R is a transitive ε -bisimulation or an ε -perturbed bisimulation. Since constructing quotients w.r.t. these relations by collapsing equivalence classes into single states can be done efficiently as well, the notions are therefore still suitable for constructing abstractions in practical applications.

► **Proposition 4.13.** *Given an equivalence R , one can decide in polynomial time if (i) R is a transitive ε -bisimulation and if (ii) R is an ε -perturbed bisimulation.*

5 Branching and Weak ε -Bisimulation

We now introduce approximate versions of branching and weak probabilistic bisimulation. A similar approach has been discussed sporadically in the context of noninterference under the term “weak bisimulation with precision ε ” [4, 7, 5, 6, 26, 3]. While our notion of branching ε -bisimilarity is a branching variant of transitive ε -bisimilarity $\sim_\varepsilon^*$, the weak ε -bisimilarity we propose is a weak variant of $\sim_\varepsilon$. Hence, the former is tailored to the construction of quotients of a given model, while the latter is closer to classic process relations.

► **Definition 5.1.** *An equivalence $R \subseteq S \times S$ is a branching ε -bisimulation if for all $(s, t) \in R$ and all R -closed sets $A \subseteq S$ it holds that*

$$(i) \ l(s) = l(t) \quad \text{and} \quad (ii) \ |\Pr_s([s]_R \cup A) - \Pr_t([t]_R \cup A)| \leq \varepsilon.$$

We call $s, t \in S$ branching ε -bisimilar, written $s \approx_\varepsilon^b t$, if they are related by a branching ε -bisimulation. LMCs $\mathcal{M}$ and $\mathcal{N}$ are branching ε -bisimilar, written $\mathcal{M} \approx_\varepsilon^b \mathcal{N}$, if $s_{init}^\mathcal{M} \approx_\varepsilon^b s_{init}^\mathcal{N}$ in $\mathcal{M} \oplus \mathcal{N}$.

We require branching ε -bisimulations to be equivalences, as their goal is to abstract from stutter steps inside a state’s equivalence class. Because of transitivity, Definition 5.1 can also be formulated in the style of Definition 3.1 and should thus not be understood as an explicit extension of Definition 3.4. With the same arguments as for $\sim_\varepsilon^*$ and $\simeq_\varepsilon$, transitivity causes that there may not be a unique maximal branching ε -bisimulation, that $\approx_\varepsilon^b$ is not additive in the tolerances, and that it can differentiate bisimilar models: the first claim follows from $s \approx_\varepsilon^b t$ and $t \approx_\varepsilon^b u$ but $s \not\approx_\varepsilon^b u$ in Figure 7, the others from $\sim_\varepsilon^* = \approx_\varepsilon^b$ in Figure 8.

► **Definition 5.2.** *A reflexive and symmetric relation $R \subseteq S \times S$ is a weak ε -bisimulation if for all $(s, t) \in R$ and all $A \subseteq S$ it holds that*


$$(i) \ l(s) = l(t) \quad \text{and} \quad (ii) \ \Pr_s(L(s) \cup A) \leq \Pr_t(L(t) \cup R(A)) + \varepsilon.$$

We call $s, t \in S$ weakly ε -bisimilar, written $s \approx_\varepsilon^w t$, if they are related by a weak ε -bisimulation. LMCs $\mathcal{M}$ and $\mathcal{N}$ are weakly ε -bisimilar, written $\mathcal{M} \approx_\varepsilon^w \mathcal{N}$, if $s_{init}^\mathcal{M} \approx_\varepsilon^w s_{init}^\mathcal{N}$ in $\mathcal{M} \oplus \mathcal{N}$.


In contrast to branching ε -bisimulations, we do not require transitivity for weak ε -bisimulations. As it turns out, $\approx_\varepsilon^w$ is instead additive in the tolerances.

► **Lemma 5.3.** *$s \approx_\varepsilon^w t$ and $t \approx_\delta^w u$ implies $s \approx_{\varepsilon+\delta}^w u$.*

Further, $\approx_0^w$ and $\approx_0^b$ coincide with $\approx^w$ and $\approx^b$, respectively, so our notions are conservative extensions of their exact counterparts. In particular, as $\approx^w = \approx^b$ for LMCs [10], it follows that $\approx_0^w = \approx_0^b$. For $\varepsilon > 0$ the notions can, however, become incomparable. This is different compared to the nonprobabilistic case, where $\approx^b$ is strictly finer than $\approx^w$ [51].



■ **Figure 11** The LMCs used in the proof of Proposition 5.4.



■ **Figure 12** The LMCs used in the proof of (i) of Proposition 5.5.

► **Proposition 5.4.** For $0 < \varepsilon < \frac{1}{4}$, $s \approx_\varepsilon^b t \not\approx s \approx_\varepsilon^w t$ and $s \approx_\varepsilon^w t \not\approx s \approx_\varepsilon^b t$.

Proof. Let $\varepsilon \in (0, \frac{1}{4})$ and consider Figure 11. In the left LMC, $s \approx_\varepsilon^b t$, as the largest branching ε -bisimulation is induced by the equivalence classes $\{\{s, t\}, \{s_1, t_1\}, \{x\}, \{y\}\}$ and, in particular, $s \not\approx_\varepsilon^b s_1$ and $t \not\approx_\varepsilon^b t_1$. However, $s \not\approx_\varepsilon^w t$ as $\Pr_t(L(t) \cup \{x\}) = \frac{5}{8} + \frac{5}{4}\varepsilon - \varepsilon^2 > \frac{5}{8} + \varepsilon = \Pr_s(L(s) \cup \{x\}) + \varepsilon$. Furthermore, in the right LMC, $s \approx_\varepsilon^w t$ while $s \not\approx_\varepsilon^b t$ since any branching ε -bisimulation R that contains (s, t) must also contain (u, w) due to transitivity, which is not possible as, e.g., $|\Pr_u([u]_R \cup [x]_R) - \Pr_w([w]_R \cup [x]_R)| > \varepsilon$. ◀

The major difference between $\approx_\varepsilon^w$, $\approx_\varepsilon^b$ and $\sim_\varepsilon, \equiv_\varepsilon$ is that the former can abstract from (some) stutter steps. Consequently, if no stuttering is possible, i.e., when $P(s)(L(s)) = 0$ for all $s \in S$, we have $\sim_\varepsilon^* = \approx_\varepsilon^b$ and $\sim_\varepsilon = \approx_\varepsilon^w$. Otherwise, the notions become incomparable.

► **Proposition 5.5.** Let $\approx_\varepsilon \in \{\approx_\varepsilon^b, \approx_\varepsilon^w\}$. Then there are LMCs with states $s, t \in S$ such that (i) $s \sim_\varepsilon t$ and $s \equiv_\varepsilon t$ but $s \not\approx_\varepsilon t$, and (ii) $s \approx_\varepsilon t$ but $s \not\sim_\varepsilon t$ and $s \not\equiv_\varepsilon t$. Hence, $\approx_\varepsilon$ and $\sim_\varepsilon, \equiv_\varepsilon$ are incomparable. Furthermore, (i) and (ii) also hold for $\sim_\varepsilon^*$ and $\equiv_\varepsilon^*$ instead of $\sim_\varepsilon$ and $\equiv_\varepsilon$.

Proof. To show (i) we do a case distinction on $\approx_\varepsilon$. If $\approx_\varepsilon = \approx_\varepsilon^b$, consider the LMC on the left of Figure 12 where $\varepsilon_1, \varepsilon_2 \in (0, 1)$, $\varepsilon_1 \neq \varepsilon_2$, $\varepsilon_1 + \varepsilon_2 < 1$, and $\varepsilon = |\varepsilon_1 - \varepsilon_2|$. In this model, both $s \sim_\varepsilon t$ and $s \equiv_\varepsilon t$. However, for any equivalence R that only relates states with the same label, $|\Pr_s([s]_R \cup \{x_1\}) - \Pr_t([t]_R \cup \{x_1\})| = \frac{|\varepsilon_1 - \varepsilon_2|}{\varepsilon_1 + \varepsilon_2} \stackrel{\varepsilon_1 + \varepsilon_2 < 1}{>} |\varepsilon_1 - \varepsilon_2| = \varepsilon$, so $s \not\approx_\varepsilon^b t$.


If $\approx_\varepsilon = \approx_\varepsilon^w$, consider the right of Figure 12 with $\varepsilon \in (0, 1)$. There, $s \sim_\varepsilon t$ and $s \equiv_\varepsilon t$. However, $\Pr_t(L(t) \cup \{x\}) = 1 > \frac{4(1-\varepsilon)^2}{(2-\varepsilon)^2} = \Pr_s(L(t) \cup \{x\})$ for all $\varepsilon \in (0, 1)$, so $s \not\approx_\varepsilon^w t$.

The second claim follows when considering an LMC with three states, say s, t and x , with initial state s and $l(s) = l(t) \neq l(x)$ as well as $P(s)(t) = P(t)(x) = P(x)(x) = 1$. There, $s \approx_\varepsilon^b t$ and $s \approx_\varepsilon^w t$ for any ε , but neither $s \equiv_\varepsilon t$ nor $s \sim_\varepsilon t$.

The claims are shown analogously when replacing $\sim_\varepsilon$ and $\equiv_\varepsilon$ with $\sim_\varepsilon^*$ resp. $\equiv_\varepsilon^*$. ◀

Note that the anomaly of $\equiv_\varepsilon$ described in Example 3.5 does not occur for branching ε -bisimilarity, as here transitivity would enforce $u_i \approx_\varepsilon^b u_j$ for all i, j in Figure 3 if $s \approx_\varepsilon^b t$.

The next lemma bounds the probabilities of states related by $\approx_\varepsilon^b$ or $\approx_\varepsilon^w$ to stutter forever.



■ **Figure 13** An LMC $\mathcal{M}$ (left) and its transformation $\mathcal{M}_R$ (right) w.r.t. the branching ε -bisimulation R with equivalence classes $\{s, t, s_1, t_1\}, \{p, q\}, \{x\}$ for $0 < \varepsilon < 1 - 2\delta$ and $\text{div}_R = \{\{s, t, s_1, t_1\}, \{x\}\}$.

► **Lemma 5.6.** *Let $\varepsilon \in [0, 1]$ and let R be a branching ε -bisimulation.*

1. *If $(s, t) \in R$ and $C = [s]_R = [t]_R$ then $|\Pr_s(\Box C) - \Pr_t(\Box C)| \leq \varepsilon$.*
2. *If $\mathcal{M}$ is finite and $(s, t) \in R$ then, for any $C \in S/R$, either (i) $\Pr_s(\Box C) = 0$ or (ii) $\Pr_s(\Box C) \geq 1 - \varepsilon$ for all $s \in C$.*
3. *If $s \approx_\varepsilon^w t$ and $b = l(s) = l(t)$ then $|\Pr_s(\Box b) - \Pr_t(\Box b)| \leq \varepsilon$.*

Since $\approx_\varepsilon^b$ and $\approx_\varepsilon^w$ cannot differentiate single steps from steps after an arbitrary (but finite) amount of stuttering, they do not preserve any next-step probabilities. Furthermore, in Figure 4 both $s \approx_\varepsilon^b t$ and $s \approx_\varepsilon^w t$, so by Example 3.9 we cannot expect a better bound for finite horizon satisfaction probabilities in related states than the one from [14] stated in Theorem 3.8. We can, however, extend Theorem 3.11 to states related by $\approx_\varepsilon^b$ and $\approx_\varepsilon^w$.

Given an equivalence R on a finite LMC $\mathcal{M}$, let $\text{div}_R \subseteq S/R$ be the set of *divergent* R -equivalence classes, i.e., $C \in \text{div}_R$ iff $\Pr_s(\Box C) \geq 1 - \varepsilon$ for all $s \in C$. We construct from $\mathcal{M}$ an LMC $\mathcal{M}_R$ and an equivalence R^b on $\mathcal{M}_R$ with $R \subseteq R^b$. Intuitively, $\mathcal{M}_R$ is obtained from $\mathcal{M}$ by redirecting the probabilities $\Pr_s(\Box C)$ for $C = [s]_R$ to fresh “divergence states” s_C .

► **Definition 5.7.** *Given a finite LMC $\mathcal{M}$ and an equivalence R that only relates states with the same label, let $\mathcal{M}_R = (S_R, P_R, s_{\text{init}}, l_R)$ with*

- $S_R = S \cup \{s_C \mid C \in \text{div}_R\}$ where the s_C are fresh, pairwise different states
- $l_R(s) = l(s)$ if $s \in S$ and $l(s_C) = l(s)$ for some $s \in C$
- for $s \in S$ and $C = [s]_R$, the values of the distribution $P_R(s)$ are defined by

$$P_R(s)(t) = \begin{cases} \Pr_s(C \cup t), & \text{if } t \in S \setminus C \\ \Pr_s(\Box C), & \text{if } s \neq s_C \text{ and } t = s_C \\ 1, & \text{if } s = t = s_C \\ 0, & \text{otherwise} \end{cases}.$$

An example for the transformation from $\mathcal{M}$ to $\mathcal{M}_R$ can be found in Figure 13. We now show the connection between branching ε -bisimulations R on finite LMCs $\mathcal{M}$ and transitive ε -bisimulations on their transformations $\mathcal{M}_R$.

► **Lemma 5.8.** *Let $\mathcal{M}$ be finite, R an equivalence relating only equally labeled states, and R^b the finest equivalence on S_R with $R \subseteq R^b$ and $(s, s_C) \in R^b$ for all $C \in \text{div}_R$ and $s \in C$. Then R is a branching ε -bisimulation on $\mathcal{M}$ iff R^b is a transitive ε -bisimulation on $\mathcal{M}_R$.*

It is clear from the definition of $\mathcal{M}_R$ that for every $C \in S/R$ and all $s \in S$ with $s \notin C$ we have $\Pr_s^{\mathcal{M}}([s]_R \cup C) = P_R(s)(C)$. Hence, Lemma 5.8 allows us to transfer Theorem 3.11 to states $s \approx_\varepsilon^b t$, since they are ε -bisimilar in $\mathcal{M}_R$. As in $\mathcal{M}_R$ any transition from s to a $u \in S$ represents an equivalence class change in $\mathcal{M}$, the random variable N^b now has to count the number of *equivalence class changes* on paths to a g - or f -labeled state.

► **Corollary 5.9.** *Let $\mathcal{M}$ be finite, let some states in $\mathcal{M}$ be labeled with g , and let exactly the states that cannot reach a g -labeled state be labeled with f . Further, let $s \approx_\varepsilon^b t$, and let N^b denote the random variable that counts the number of equivalence class changes until a g - or f -labeled state is reached. Then $|\Pr_s(\Diamond g) - \Pr_t(\Diamond g)| \leq \varepsilon \cdot \mathbb{E}_s(N^b)$.*

Furthermore, it is possible to extend Theorem 3.11 to weakly ε -bisimilar states.

► **Proposition 5.10.** *Let $\mathcal{M}$, f and g be as in Corollary 5.9, let $s \approx_\varepsilon^w t$ and let N^w denote the random variable that counts the number of label changes until a g - or f -labeled state is reached. Then $|\Pr_s(\Diamond g) - \Pr_t(\Diamond g)| \leq \varepsilon \cdot \mathbb{E}_s(N^w)$.*

Proof sketch. Let $\mathcal{L} = \{b \in 2^{AP} \mid \exists s \in S: \Pr_s(\Box b) > 0\}$. From $\mathcal{M}$ we construct an LMC $\mathcal{M}^w$, almost similar to $\mathcal{M}_R$ in Definition 5.7. The main differences are that we introduce fresh states s_b for all $b \in \mathcal{L}$, and that we set $P^w(s)(t) = \Pr_s(L(s) \cup t)$ for all $s, t \in S$ with $l(s) \neq l(t)$ as well as $P^w(s)(s_b) = \Pr_s(\Box b)$ if $l(s) = b \in \mathcal{L}$. Because for any weak ε -bisimulation R the finest reflexive and symmetric relation R^w on $\mathcal{M}^w$ with $R \subseteq R^w$ and $(s, s_b) \in R^w$ iff $b = l(s)$ and $\Pr_s(\Box b) \geq 1 - \varepsilon$ is an ε -bisimulation on $\mathcal{M}^w$, the result follows from Theorem 3.11. ◀

As the LMCs in Figure 6 are both branching $\frac{\varepsilon}{p}$ -bisimilar and weak $\frac{\varepsilon}{p}$ -bisimilar, and since in these models $\mathbb{E}_s(N^b) = \mathbb{E}_s(N^w) = 1$, the bounds are again tight by Example 3.13.

We finish this section by analyzing the complexity of deciding if two given states s, t are branching ε -bisimilar, i.e., if $s \approx_\varepsilon^b t$. The analogous problem for $\approx_\varepsilon^w$ is left open.

► **Theorem 5.11.** *Given a finite $\mathcal{M}$, $s, t \in S$, and $\varepsilon \in (0, 1]$, deciding if $s \approx_\varepsilon^b t$ is NP-complete.*

6 Conclusion and Future Work

We investigated several new types of approximate probabilistic bisimulation and showed how they interrelate, as well as how they are connected to notions from the literature like, e.g., $\sim_\varepsilon$ and $\equiv_\varepsilon$ (see Figure 1). These connections in turn allowed the transfer of known preservation results for logical formulas between the different notions, which we extended by tight bounds on the absolute difference of unbounded reachability probabilities in weak and branching ε -bisimilar states. Additionally, we established complexity results for most of our relations.

The results of Section 4 indicate that ε -perturbed bisimilarity $\simeq_\varepsilon$ and transitive ε -bisimilarity $\sim_\varepsilon^*$ show some anomalies (lack of additivity, the possibility to differentiate bisimilar models and the fact that they themselves are not necessarily an ε -perturbed resp. a transitive ε -bisimulation) when viewed as process relations. However, both relations can be interesting for algorithmic purposes as they permit efficient quotienting techniques: given a transitive ε -bisimulation R (with or without the centroid property) on an LMC $\mathcal{M}$, one can build in polynomial time a quotient LMC that arises from $\mathcal{M}$ by collapsing all R -equivalence classes into single states. The quotient under an ε -perturbed bisimulation R_1 enjoys the property that every state s and its R_1 -equivalence class $[s]_{R_1}$ are $\frac{\varepsilon}{2}$ -bisimilar [38], while for the quotients under a transitive ε -bisimulation R_2 that lacks the centroid property we can only guarantee $s \sim_\varepsilon [s]_{R_2}$. On the other hand, transitive ε -bisimulations can identify more states and hence can induce smaller quotients.

Similarly, the transitivity of branching ε -bisimulations causes the same anomalies as for $\simeq_\varepsilon$ and $\sim_\varepsilon^*$. However, checking if a given equivalence is a branching ε -bisimulation and constructing a corresponding quotient is again possible in polynomial time. Hence, investigating the potential of transitive (or branching) ε -bisimulations as abstraction techniques for an approximate analysis of LMCs in practice is an interesting future research direction.

Other open questions include the search for a characterization of logical formulas that distinguish $\sim_\varepsilon, \sim_\varepsilon^*, \approx_\varepsilon^w, \approx_\varepsilon^b$ and $\simeq_\varepsilon$, and how our results relate to bisimilarity distances [49].

References

- 1 Alessandro Abate. Approximation metrics based on probabilistic bisimulations for general state-space Markov processes: A survey. *Electronic Notes in Theoretical Computer Science*, 297:3–25, 2013. Proceedings of the first workshop on Hybrid Autonomous Systems.
- 2 Alessandro Abate, Marta Kwiatkowska, Gethin Norman, and David Parker. Probabilistic model checking of labelled Markov processes via finite approximate bisimulations. In Franck van Breugel, Elham Kashefi, Catuscia Palamidessi, and Jan Rutten, editors, *Horizons of the Mind. A Tribute to Prakash Panangaden: Essays Dedicated to Prakash Panangaden on the Occasion of His 60th Birthday*, volume 8464 of *Lecture Notes in Computer Science (LNCS)*, pages 40–58. Springer International Publishing, Cham, 2014.
- 3 Alessandro Aldini. A note on the approximation of weak probabilistic bisimulation, 2009.
- 4 Alessandro Aldini, Mario Bravetti, Alessandra Di Pierro, Roberto Gorrieri, Chris Hankin, and Herbert Wiklicky. Two formal approaches for approximating noninterference properties. In Riccardo Focardi and Roberto Gorrieri, editors, *Foundations of Security Analysis and Design II (FOSAD 2001)*, volume 2946 of *Lecture Notes in Computer Science (LNCS)*, pages 1–43. Springer, Berlin, Heidelberg, 2004.
- 5 Alessandro Aldini, Mario Bravetti, and Roberto Gorrieri. A process-algebraic approach for the analysis of probabilistic noninterference. *J. Comput. Secur.*, 12(2):191–245, April 2004.
- 6 Alessandro Aldini and Alessandra Di Pierro. A quantitative approach to noninterference for probabilistic systems. *Electronic Notes in Theoretical Computer Science*, 99:155–182, 2004. Proceedings of the MEFISTO Project 2003, Formal Methods for Security and Time.
- 7 Alessandro Aldini and Roberto Gorrieri. Security analysis of a probabilistic non-repudiation protocol. In Holger Hermanns and Roberto Segala, editors, *Process Algebra and Probabilistic Methods: Performance Modeling and Verification*, volume 2399 of *Lecture Notes in Computer Science (LNCS)*, pages 17–36. Springer, Berlin, Heidelberg, 2002.
- 8 Adnan Aziz, Vigyan Singhal, Felice Balarin, Robert K. Brayton, and Alberto L. Sangiovanni-Vincentelli. It usually works: The temporal logic of stochastic systems. In Pierre Wolper, editor, *Computer Aided Verification (CAV 1995)*, volume 939 of *Lecture Notes in Computer Science (LNCS)*, pages 155–165. Springer, Berlin, Heidelberg, 1995.
- 9 Christel Baier. Polynomial time algorithms for testing probabilistic bisimulation and simulation. In Rajeev Alur and Thomas A. Henzinger, editors, *Computer Aided Verification (CAV 1996)*, volume 1102 of *Lecture Notes in Computer Science (LNCS)*, pages 50–61. Springer, Berlin, Heidelberg, 1996.
- 10 Christel Baier and Holger Hermanns. Weak bisimulation for fully probabilistic processes. In Orna Grumberg, editor, *Computer Aided Verification (CAV 1997)*, volume 1254 of *Lecture Notes in Computer Science (LNCS)*, pages 119–130. Springer, Berlin, Heidelberg, 1997.
- 11 Christel Baier and Joost-Pieter Katoen. *Principles of Model Checking*. The MIT Press, 2008.
- 12 Christel Baier, Joost-Pieter Katoen, Holger Hermanns, and Verena Wolf. Comparative branching-time semantics for Markov chains. *Information and Computation*, 200(2):149–214, 2005.
- 13 Massimo Bartoletti, Maurizio Murgia, and Roberto Zunino. Sound approximate and asymptotic probabilistic bisimulations for PCTL. *Logical Methods in Computer Science*, 19(1), 2023.
- 14 Gaoang Bian and Alessandro Abate. On the relationship between bisimulation and trace equivalence in an approximate probabilistic context. In Javier Esparza and Andrzej S. Murawski, editors, *Foundations of Software Science and Computation Structures (FoSSaCS)*, volume 10203 of *Lecture Notes in Computer Science (LNCS)*, pages 321–337. Springer, Berlin, Heidelberg, 2017.
- 15 Gaoang Bian and Alessandro Abate. On the relationship between bisimulation and trace equivalence in an approximate probabilistic context (extended version). *CoRR*, abs/1701.04547, 2017. Full version of [14]. [arXiv:1701.04547](https://arxiv.org/abs/1701.04547).
- 16 Béla Bollobás and Nicolas Th. Varopoulos. Representation of systems of measurable sets. *Mathematical Proceedings of the Cambridge Philosophical Society*, 78(2):323–325, 1975.
- 17 Peter Buchholz. Exact and ordinary lumpability in finite Markov chains. *Journal of Applied Probability*, 31:59–75, January 1995.

- 18 David de Frutos-Escrig, Jeroen J. A. Keiren, and Tim A. C. Willemse. Games for bisimulations and abstraction. *CoRR*, abs/1611.00401, 2016. [arXiv:1611.00401](#).
- 19 Salem Derisavi, Holger Hermanns, and William H. Sanders. Optimal state-space lumping in Markov chains. *Information Processing Letters*, 87(6):309–315, 2003.
- 20 Josée Desharnais, Vineet Gupta, Radha Jagadeesan, and Prakash Panangaden. Metrics for labeled Markov systems. In Jos C. M. Baeten and Sjouke Mauw, editors, *CONCUR’99 Concurrency Theory (CONCUR 1999)*, volume 1664 of *Lecture Notes in Computer Science (LNCS)*, pages 258–273. Springer, Berlin, Heidelberg, 1999.
- 21 Josée Desharnais, Abbas Edalat, and Prakash Panangaden. A logical characterization of bisimulation for labeled Markov processes. In *Proceedings of the Thirteenth Annual IEEE Symposium on Logic in Computer Science (LiCS)*, pages 478–487, 1998.
- 22 Josée Desharnais, Abbas Edalat, and Prakash Panangaden. Bisimulation for labelled Markov processes. *Information and Computation*, 179(2):163–193, 2002.
- 23 Josée Desharnais, Vineet Gupta, Radha Jagadeesan, and Prakash Panangaden. Approximating labelled Markov processes. *Information and Computation*, 184(1):160–200, 2003.
- 24 Josée Desharnais, Vineet Gupta, Radha Jagadeesan, and Prakash Panangaden. Weak bisimulation is sound and complete for pCTL*. *Information and Computation*, 208(2):203–219, 2010.
- 25 Josée Desharnais, François Laviolette, and Mathieu Tracol. Approximate analysis of probabilistic processes: Logic, simulation and games. In *Fifth International Conference on Quantitative Evaluation of Systems (QEST 2008)*, pages 264–273, 2008.
- 26 Alessandra Di Pierro, Chris Hankin, and Herbert Wiklicky. Measuring the confinement of probabilistic systems. *Theoretical Computer Science*, 340(1):3–56, 2005. Theoretical Foundations of Security Analysis and Design II.
- 27 Alessandro D’Innocenzo, Alessandro Abate, and Joost-Pieter Katoen. Robust PCTL model checking. In *Proceedings of the 15th ACM International Conference on Hybrid Systems: Computation and Control (HSCC 2012)*, pages 275–286, New York, NY, USA, 2012. Association for Computing Machinery.
- 28 András Faragó. On the convergence rate of quasi lumpable Markov chains. In András Horváth and Miklós Telek, editors, *Formal Methods and Stochastic Models for Performance Evaluation*, volume 4054 of *Lecture Notes in Computer Science (LNCS)*, pages 138–147. Springer, Berlin, Heidelberg, 2006.
- 29 Giuliana Franceschinis and Richard R. Muntz. Bounds for quasi-lumpable markow chains. *Perform. Evaluation*, 20:223–243, 1994.
- 30 Alessandro Giacalone, Chi-Chang Jou, and Scott A. Smolka. Algebraic reasoning for probabilistic concurrent systems. In Manfred Broy and Cliff B. Jones, editors, *Proceedings of the IFIP TC2 Working Conference on Programming Concepts and Methods*, pages 443–458, 1990.
- 31 Sofie Haesaert, Petter Nilsson, and Sadegh Soudjani. Formal multi-objective synthesis of continuous-state mdps. In *2021 American Control Conference (ACC)*, pages 3428–3433, 2021.
- 32 Sofie Haesaert and Sadegh Soudjani. Robust dynamic programming for temporal logic control of stochastic systems. *IEEE Transactions on Automatic Control*, 66(6):2496–2511, 2021.
- 33 Hans Hansson and Bengt Jonsson. A logic for reasoning about time and reliability. *Formal Aspects of Computing*, 6:512–535, 1994.
- 34 Christian Hensel, Sebastian Junges, Joost-Pieter Katoen, Tim Quatmann, and Matthias Volk. The probabilistic model checker storm. *International Journal on Software Tools for Technology Transfer (STTT)*, 24:589–610, 2022.
- 35 David N. Jansen, Jan Friso Groote, Ferry Timmers, and Pengfei Yang. A Near-Linear-Time Algorithm for Weak Bisimilarity on Markov Chains. In Igor Konnov and Laura Kovács, editors, *31st International Conference on Concurrency Theory (CONCUR 2020)*, volume 171 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 8:1–8:20, Dagstuhl, Germany, 2020. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
- 36 Joost-Pieter Katoen. The probabilistic model checking landscape. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science (LiCS’16)*, pages 31–45, New York, NY, USA, 2016. Association for Computing Machinery.

- 37 Joost-Pieter Katoen, Tim Kemna, Ivan Zapreev, and David N. Jansen. Bisimulation minimisation mostly speeds up probabilistic model checking. In Orna Grumberg and Michael Huth, editors, *Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2007)*, volume 4424 of *Lecture Notes in Computer Science (LNCS)*, pages 87–101. Springer, Berlin, Heidelberg, 2007.
- 38 Stefan Kiefer and Qiyi Tang. Approximate Bisimulation Minimisation. In Mikołaj Bojańczyk and Chandra Chekuri, editors, *41st IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2021)*, volume 213 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 48:1–48:16, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
- 39 Stefan Kiefer and Qiyi Tang. Approximate bisimulation minimisation. *CoRR*, abs/2110.00326, 2021. Full version of [38]. [arXiv:2110.00326](https://arxiv.org/abs/2110.00326).
- 40 Orna Kupferman and Moshe Y. Vardi. Model checking of safety properties. In Nicolas Halbwachs and Doron Peled, editors, *Computer Aided Verification (CAV 1999)*, volume 1633 of *Lecture Notes in Computer Science (LNCS)*, pages 172–183, Berlin, Heidelberg, 1999. Springer, Berlin, Heidelberg.
- 41 Kim G. Larsen and Arne Skou. Bisimulation through probabilistic testing. *Information and Computation*, 94(1):1–28, 1991.
- 42 R. Milner. *Communication and Concurrency*. Prentice-Hall, Inc., USA, 1989.
- 43 Robin Milner. *A Calculus of Communicating Systems*, volume 92 of *Lecture Notes in Computer Science (LNCS)*. Springer-Verlag, Berlin, Heidelberg, 1982.
- 44 Amir Pnueli. The temporal logic of programs. In *18th Annual Symposium on Foundations of Computer Science (SFCS 1977)*, pages 46–57, 1977.
- 45 Roberto Segala. *Modeling and verification of randomized distributed real-time systems*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, USA, 1995.
- 46 Roberto Segala and Nancy Lynch. Probabilistic simulations for probabilistic processes. In Bengt Jonsson and Joachim Parrow, editors, *CONCUR '94: Concurrency Theory*, volume 836 of *Lecture Notes in Computer Science (LNCS)*, pages 481–496. Springer, Berlin, Heidelberg, 1994.
- 47 Mathieu Tracol, Josée Desharnais, and Abir Zhiova. Computing distances between probabilistic automata. In Mieke Massink and Gethin Norman, editors, *Proceedings Ninth Workshop on Quantitative Aspects of Programming Languages (QAPL 2011)*, Saarbrücken, Germany, April 1-3, 2011, volume 57 of *EPTCS*, pages 148–162, 2011.
- 48 Antti Valmari and Giuliana Franceschinis. Simple $O(m \log n)$ time Markov chain lumping. In Javier Esparza and Rupak Majumdar, editors, *Tools and Algorithms for the Construction and Analysis of Systems*, volume 6015 of *Lecture Notes in Computer Science (LNCS)*, pages 38–52, Berlin, Heidelberg, 2010. Springer, Berlin, Heidelberg.
- 49 Franck van Breugel. Probabilistic bisimilarity distances. *ACM SIGLOG News*, 4(4):33–51, 2017.
- 50 Rob J. van Glabbeek, Jan Friso Groote, and Erik P. de Vink. A complete axiomatization of branching bisimilarity for a simple process language with probabilistic choice. In Mário S. Alvim, Kostas Chatzikokolakis, Carlos Olarte, and Frank Valencia, editors, *The Art of Modelling Computational Systems: A Journey from Logic and Concurrency to Security and Privacy: Essays Dedicated to Catuscia Palamidessi on the Occasion of Her 60th Birthday*, volume 11760 of *Lecture Notes in Computer Science (LNCS)*, pages 139–162, Cham, 2019. Springer International Publishing.
- 51 Rob J. van Glabbeek and W. Peter Weijland. Branching time and abstraction in bisimulation semantics. *J. ACM*, 43(3):555–600, May 1996.
- 52 Ralf Wimmer and Bernd Becker. Correctness issues of symbolic bisimulation computation for Markov chains. In Bruno Müller-Clostermann, Klaus Ehtle, and Erwin P. Rathgeb, editors, *Measurement, Modelling, and Evaluation of Computing Systems and Dependability and Fault Tolerance (MMB&DFT 2010)*, volume 5987 of *Lecture Notes in Computer Science (LNCS)*, pages 287–301. Springer, Berlin, Heidelberg, 2010.