



A Sound Up-to- n , δ Bisimilarity for PCTL

Massimo Bartoletti^{1(✉)}, Maurizio Murgia^{2,3,4}, and Roberto Zunino²

¹ Università Degli Studi di Cagliari, Cagliari, Italy
bart@unica.it

² Università degli Studi di Trento, Trento, Italy

³ Università degli Studi di Bologna, Bologna, Italy

⁴ Gran Sasso Science Institute, L'Aquila, Italy

Abstract. We tackle the problem of establishing the soundness of approximate bisimilarity with respect to PCTL and its relaxed semantics. To this purpose, we consider a notion of bisimilarity similar to the one introduced by Desharnais, Laviolette, and Tracol, which is parametric with respect to an approximation error δ , and to the depth n of the observation along traces. Essentially, our soundness theorem establishes that, when a state q satisfies a given formula up-to error δ and steps n , and q is bisimilar to q' up-to error δ' and enough steps, we prove that q' also satisfies the formula up-to a suitable error δ'' and steps n . The new error δ'' is computed from δ, δ' and the formula, and only depends linearly on n . We provide a detailed overview of our soundness proof.

Keywords: PCTL · Probabilistic processes · Approximate bisimulation

1 Introduction

The behaviour of many real-world systems can be formally modelled as probabilistic processes, e.g. as discrete-time Markov chains. Specifying and verifying properties on these systems requires probabilistic versions of temporal logics, such as PCTL [25]. PCTL allows to express probability bounds using the formula $\text{Pr}_{\geq\pi}[\psi]$, which is satisfied by those states starting from which the path formula ψ holds with probability $\geq \pi$. A well-known issue is that real-world systems can have tiny deviations from their mathematical models, while logical properties, such as those written in PCTL, impose sharp constraints on the behaviour. To address this issue, one can use a *relaxed* semantics for PCTL, as in [21]. There, the semantics of formulae is parameterised over the error $\delta \geq 0$ one is willing to tolerate. While in the standard semantics of $\text{Pr}_{\geq\pi}[\psi]$ the bound $\geq \pi$ is *exact*, in relaxed PCTL this bound is weakened to $\geq \pi - \delta$. So, the relaxed semantics generalises the standard PCTL semantics of [25], which can be obtained by choosing $\delta = 0$. Instead, choosing an error $\delta > 0$ effectively provides a way to measure “how much” a state satisfies a given formula: some states might require only a very small error, while others a much larger one.

When dealing with temporal logics such as PCTL, one often wants to study some notion of state equivalence which preserves the semantics of formulae: that is, when two states are equivalent, they satisfy the same formulae. For instance, probabilistic bisimilarities like those in [16, 19, 29] preserve the semantics of formulae for PCTL and other temporal logics. Although *strict* probabilistic bisimilarity preserves the semantics of relaxed PCTL, it is not *robust* against small deviations in the probability of transitions in Markov chains [24]. A possible approach to deal with this issue is to also relax the notion of probabilistic bisimilarity, by making it parametric with respect to an error δ [21]. Relaxing bisimilarity in this way poses a choice regarding which properties of the strict probabilistic bisimilarity are to be kept. In particular, transitivity is enjoyed by the strict probabilistic bisimilarity, but it is *not* desirable for the relaxed notion. Indeed, we could have three states q, q' and q'' where the behaviour of q and q' is similar enough (within the error δ), the behaviour of q' and q'' is also similar enough (within δ), but the distance between q and q'' is larger than the allowed error δ . At best, we can have a sort of “triangular inequality”, where q and q'' can still be related but only with a larger error $2 \cdot \delta$.




Fig. 1. A Markov chain modelling repeated tosses of a fair coin.

Bisimilarity is usually defined by coinduction, essentially requiring that the relation is preserved along an arbitrarily long sequence of moves. Still, in some settings, observing the behaviour over a very long run is undesirable. For instance, consider the PCTL formula $\phi = \Pr_{\geq 0.5}[\text{true } U^{\leq n} \text{ a}]$, which is satisfied by those states from which, with probability ≥ 0.5 , a is satisfied within n steps. In this case, a behavioural equivalence relation that preserves the semantics of ϕ can neglect the long-run behaviour after n steps. More generally, if all the until operators are *bounded*, as in $\phi_1 U^{\leq k} \phi_2$, then each formula has an upper bound of steps n after which a behavioural equivalence relation can ignore what happens next. Observing the behaviour after this upper bound is unnecessarily strict, and indeed in some settings it is customary to neglect what happens in the very long run. For instance, a real-world player repeatedly tossing a coin is usually considered equivalent to a Markov chain with two states and four transitions with probability $1/2$ (see Fig. 1), even if in the long run the real-world system will diverge from the ideal one (e.g., when the player dies).

Another setting where observing the long-term behaviour is notoriously undesirable is that of cryptography. When studying the security of systems modelling

cryptographic protocols, two states are commonly considered equivalent when their behaviour is similar (up to a small error δ) in the short run, even when in the long run they diverge. For instance, a state q could represent an ideal system where no attacks can be performed by construction, while another state q' could represent a real system where an adversary can try to disrupt the cryptographic protocol. In such a scenario, if the protocol is secure, we would like to have q and q' equivalent, since the behaviour of the real system is close to the one of the ideal system. Note that in the real system an adversary can repeatedly try to guess the secret cryptographic keys, and break security in the very long run, with very high probability. Accordingly, standard security definitions require that the behaviour of the ideal and real system are within a small error, but only for a *bounded* number of steps, after which their behaviour could diverge.

Contributions. To overcome the above mentioned issues, in this work we introduce a bounded, approximate notion of bisimilarity \sim_{δ}^n , that only observes the first n steps, and allows for an error δ . Unlike standard bisimilarity, our relation is naturally defined by *induction* on n . We call this looser variant of bisimilarity an *up-to- n, δ* bisimilarity. We showcase up-to- n, δ bisimilarity on a running example (Examples 1, 2 and 4), comparing an ideal combination padlock against a real one which can be opened by an adversary guessing its combination. We show that the two systems are bisimilar up-to- n, δ , while they are not bisimilar according to the standard coinductive notion. We then discuss how the two systems satisfy a basic security property expressed in PCTL, with suitable errors. To make our theory amenable to reason about infinite-state systems, such as those usually found when modelling cryptographic protocols, all our results apply to Markov chains with countably many states. In this respect, our work departs from most literature on probabilistic bisimulations [21, 33] and bisimilarity distances [7, 9, 13, 22, 34–36], which usually assume finite-state Markov chains. In Example 3 we exploit infinite-state Markov chains to compare a biased random bit generator with an ideal one.

Our main contribution is a soundness theorem establishing that, when a state q satisfies a PCTL formula ϕ (up to a given error), any bisimilar state $q' \sim q$ must also satisfy ϕ , at the cost of a slight increase of the error. More precisely, if ϕ only involves until operators bounded by $\leq n$, state q satisfies ϕ up to some error, and bisimilarity holds for enough steps and error δ , then q' satisfies ϕ with an *additional* asymptotic error $O(n \cdot \delta)$.

This asymptotic behaviour is compatible with the usual assumptions of computational security in cryptography. There, models of security protocols include a security parameter η , which affects the length of the cryptographic keys and the running time of the protocol: more precisely, a protocol is assumed to run for $n(\eta)$ steps, which is polynomially bounded w.r.t. η . As already mentioned above, cryptographic notions of security do not observe the behaviour of the systems after this bound $n(\eta)$, since in the long run an adversary can surely guess the secret keys by brute force. Coherently, a protocol is considered to be secure if (roughly) its actual behaviour is *approximately* equivalent to the ideal one for $n(\eta)$ steps and up to an error $\delta(\eta)$, which has to be a negligible function,

asymptotically approaching zero faster than any rational function. Under these bounds on n and δ , the asymptotic error $O(n \cdot \delta)$ in our soundness theorem is negligible in η . Consequently, if two states q and q' represent the ideal and actual behaviour, respectively, and they are bisimilar up to a negligible error, they will satisfy the same PCTL formulae with a negligible error.

We provide a detailed overview of the proof of our soundness theorem, explaining the techniques that we used. Due to space limitations, we present the full proofs of our results in [5].

Related Work. There is a well-established line of research on establishing soundness and completeness of probabilistic bisimulations against various kinds of probabilistic logics [19, 23, 26, 29, 31, 32].

The work closest to ours is that of D’Innocenzo, Abate and Katoen [21], which addresses the model checking problem on relaxed PCTL. Their relaxed PCTL differs from ours in a few aspects. First, their syntax allows for an individual bound for each until operator $U^{\leq k}$, while we assume all such bounds are equal. This difference does not seem to affect our results, which could be extended to cover the general case. Second, their main result shows that bisimilar states up to a given error ϵ satisfy the same formulae ψ , provided that ψ ranges over the so-called ϵ -robust formulae. Instead, our soundness result applies to *all* PCTL formulae, and ensures that when moving from a state satisfying ϕ to a bisimilar one, ϕ is still satisfied, but at the cost of slightly increasing the error. Third, their relaxed semantics differs from ours. In ours, we relax all the probability bounds by the same amount δ . Instead, the relaxation in [21] affects the bounds by a different amount which depends on the error ϵ , the until bound k , and the underlying DTMC.

Desharnais, Laviolette and Tracol [20] use a coinductive approximate probabilistic bisimilarity, up-to an error δ . Using such coinductive bisimilarity, [20] establishes the soundness and completeness with respect to a Larsen-Skou logic [29] (instead of PCTL). In [20], a bounded, up-to n, δ version of bisimilarity is only briefly used to derive a decision algorithm for coinductive bisimilarity under the assumption that the state space is finite. In our work, instead, the bounded up-to n, δ bisimilarity is the main focus of study. In particular, our soundness result only assumes n, δ bisimilarity, which is strictly weaker than coinductive bisimilarity. Another minor difference is that [20] considers a labelled Markov process, i.e. the probabilistic variant of a labelled transition system, while we instead focus on DTMCs having labels on states.

Bian and Abate [6] study bisimulation and trace equivalence up-to an error ϵ , and show that ϵ -bisimilar states are also ϵ' -trace equivalent, for a suitable ϵ' which depends on ϵ . Furthermore, they show that ϵ -trace equivalent states satisfy the same formulae in a bounded LTL, up-to a certain error. In our work, we focus instead on the branching logic PCTL.

A related research line is that on *bisimulation metrics* [7, 8, 10]. Some of these metrics, like our up-to bisimilarity, take approximations into account [12, 17]. Similarly to our bisimilarity, bisimulation metrics allow to establish two states equivalent up-to a certain error (but usually do not take into account the bound

on the number of steps). Interestingly, Castiglioni, Gebler and Tini [12] introduce a notion of distance between Larsen-Skou formulae, and prove that the bisimulation distance between two processes corresponds to the distance between their mimicking formulae. De Alfaro, Majumdar, Raman and Stoelinga [2] elegantly characterise bisimulation metrics with a quantitative μ -calculus. Such logic allows to specify interesting properties such as maximal reachability and safety probability, and the maximal probability of satisfying a general ω -regular specification, but not full PCTL. Mio [30] characterises a bisimulation metric based on total variability with a more general quantitative μ -calculus, dubbed Łukasiewicz μ -calculus, able to encode PCTL. Both [2] and [30] do not take the number of steps into account, therefore their applicability to the analysis of security protocols is yet to be investigated.

Metrics with discount [1, 3, 9, 15, 18] are sometimes used to relate the behaviour of probabilistic processes, weighing less those events that happen in the far future compared to those happening in the first steps. Often, in these metrics each step causes the probability of the next events to be multiplied by a constant factor $c < 1$, in order to diminish their importance. Note that this discount makes it so that after η steps, this diminishing factor becomes c^η , which is a negligible function of η . As discussed before, in cryptographic security one needs to consider as important those events happening within polynomially many steps, while neglecting the ones after such a polynomial threshold. Using an exponential discount factor c^η after only η steps goes against this principle, since it would cause a secure system to be at a negligible distance from an insecure one which can be violated after just η steps. For this reason, instead of using a metric with discount, in this paper we resort to a bisimilarity that is parametrized over the number of steps n and error δ , allowing us to obtain a notion which distinguishes between the mentioned secure and insecure systems.

Several works develop algorithms to decide probabilistic bisimilarity, and to compute metrics [11, 13, 22, 34–36]. To this purpose, they restrict to finite-state systems, like e.g. probabilistic automata. Our results, instead, apply also to infinite-state systems.

2 PCTL

Assume a set \mathcal{L} of labels, ranged over by l , and let δ, π range over non-negative reals. A *discrete-time Markov chain* (DTMC) is a standard model of probabilistic systems. Throughout this paper, we consider a DTMC having a countable, possibly infinite, set of states q , each carrying a subset of labels $\ell(q) \subseteq \mathcal{L}$.

Definition 1 (Discrete-Time Markov Chain). *A (labelled) DTMC is a triple $(\mathcal{Q}, \text{Pr}, \ell)$ where:*

- \mathcal{Q} is a countable set of states;
- $\text{Pr} : \mathcal{Q}^2 \rightarrow [0, 1]$ is a function, named transition probability function;
- $\ell : \mathcal{Q} \rightarrow \mathcal{P}(\mathcal{L})$ is a labelling function

Given $q \in \mathcal{Q}$ and $Q \subseteq \mathcal{Q}$, we write $\Pr(q, Q)$ for $\sum_{q' \in Q} \Pr(q, q')$ and we require that $\Pr(q, \mathcal{Q}) = 1$ for all $q \in \mathcal{Q}$.

A *trace* is an infinite sequence of states $t = q_0 q_1 \dots$, where we write $t(i)$ for q_i , i.e. the i -th element of t . A *trace fragment* is a finite, non-empty sequence of states $\tilde{t} = q_0 \dots q_{n-1}$, where $|\tilde{t}| = n \geq 1$ is its length. Given a trace fragment \tilde{t} and a state q , we write $\tilde{t}q^\omega$ for the trace $\tilde{t}qqq\dots$.

It is well-known that, given an initial state q_0 , the DTMC induces a σ -algebra of measurable sets of traces T starting from q_0 , i.e. the σ -algebra generated by cylinders. More in detail, given a trace fragment $\tilde{t} = q_0 \dots q_{n-1}$, its cylinder $Cyl(\tilde{t}) = \{t \mid \tilde{t} \text{ is a prefix of } t\}$ is given probability $\Pr(Cyl(\tilde{t})) = \prod_{i=0}^{n-2} \Pr(q_i, q_{i+1})$. As usual, if $n = 1$ the product is empty and evaluates to 1. Closing the family of cylinders under countable unions and complement we obtain the family of measurable sets. The probability measure on cylinders then uniquely extends to all the measurable sets.

Given a set of trace fragments \tilde{T} , all starting from the same state q_0 and having the same length, we let $\Pr(\tilde{T}) = \Pr(\bigcup_{\tilde{t} \in \tilde{T}} Cyl(\tilde{t})) = \sum_{\tilde{t} \in \tilde{T}} \Pr(Cyl(\tilde{t}))$. Note that using same-length trace fragments ensures that their cylinders are disjoint, hence the second equality holds.

Below, we define PCTL formulae. Our syntax is mostly standard, except for the *until* operator. There, for the sake of simplicity, we do not bound the number of steps in the syntax $\phi_1 \mathbf{U} \phi_2$, but we do so in the semantics. Concretely, this amounts to imposing the same bound to *all* the occurrences of \mathbf{U} in the formula. Such bound is then provided as a parameter to the semantics. We foresee no problems in relaxing this restriction, which we make only to simplify our treatment.

Definition 2 (PCTL syntax). *The syntax of PCTL is given by the following grammar, defining state formulae ϕ and path formulae ψ :*

$$\begin{aligned} \phi &:: = l \mid \text{true} \mid \neg\phi \mid \phi \wedge \phi \mid \Pr_{\triangleright\pi}[\psi] & \text{where } \triangleright \in \{>, \geq\} \\ \psi &:: = \mathbf{X} \phi \mid \phi \mathbf{U} \phi \end{aligned}$$

As syntactic sugar, we write $\Pr_{<\pi}[\psi]$ for $\neg\Pr_{\geq\pi}[\psi]$, and $\Pr_{\leq\pi}[\psi]$ for $\neg\Pr_{>\pi}[\psi]$.

Given a PCTL formula ϕ , we define its maximum \mathbf{X} -nesting $\circ_{max}(\phi)$ and its maximum \mathbf{U} -nesting $\mathbf{U}_{max}(\phi)$ inductively as follows:

Definition 3 (Maximum nesting). *For $\circ \in \{\mathbf{X}, \mathbf{U}\}$, we define:*

$$\begin{aligned} \circ_{max}(l) &= 0 & \circ_{max}(\text{true}) &= 0 & \circ_{max}(\neg\phi) &= \circ_{max}(\phi) \\ \circ_{max}(\phi_1 \wedge \phi_2) &= \max(\circ_{max}(\phi_1), \circ_{max}(\phi_2)) & \circ_{max}(\Pr_{\triangleright\pi}[\psi]) &= \circ_{max}(\psi) \\ \circ_{max}(\mathbf{X}\phi) &= \circ_{max}(\phi) + \begin{cases} 1 & \text{if } \circ = \mathbf{X} \\ 0 & \text{otherwise} \end{cases} \\ \circ_{max}(\phi_1 \mathbf{U} \phi_2) &= \max(\circ_{max}(\phi_1), \circ_{max}(\phi_2)) + \begin{cases} 1 & \text{if } \circ = \mathbf{U} \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

We now define a semantics for PCTL where the probability bounds $\triangleright\pi$ in $\text{Pr}_{\triangleright\pi}[\psi]$ can be relaxed or strengthened by an error δ . Our semantics is parameterized over the *until* bound n , the error $\delta \in \mathbb{R}^{\geq 0}$, and a direction $r \in \{+1, -1\}$. Given the parameters, the semantics associates each PCTL state formula with the set of states satisfying it. Intuitively, when $r = +1$ we relax the semantics of the formula, so that increasing δ causes more states to satisfy it. More precisely, the probability bounds $\triangleright\pi$ in positive occurrences of $\text{Pr}_{\triangleright\pi}[\psi]$ are decreased by δ , while those in negative occurrences are increased by δ . Dually, when $r = -1$ we strengthen the semantics, modifying $\triangleright\pi$ in the opposite direction.

Definition 4 (PCTL semantics). *The semantics of PCTL formulas is given below. Let $n \in \mathbb{N}$, $\delta \in \mathbb{R}^{\geq 0}$ and $r \in \{+1, -1\}$.*

$$\begin{aligned}
\llbracket l \rrbracket_{\delta,r}^n &= \{q \in \mathcal{Q} \mid l \in \ell(q)\} \\
\llbracket \text{true} \rrbracket_{\delta,r}^n &= \mathcal{Q} \\
\llbracket \neg\phi \rrbracket_{\delta,r}^n &= \mathcal{Q} \setminus \llbracket \phi \rrbracket_{\delta,-r}^n \\
\llbracket \phi_1 \wedge \phi_2 \rrbracket_{\delta,r}^n &= \llbracket \phi_1 \rrbracket_{\delta,r}^n \cap \llbracket \phi_2 \rrbracket_{\delta,r}^n \\
\llbracket \text{Pr}_{\triangleright\pi}[\psi] \rrbracket_{\delta,r}^n &= \{q \in \mathcal{Q} \mid \text{Pr}(\text{Cyl}(q) \cap \llbracket \psi \rrbracket_{\delta,r}^n) + r \cdot \delta \triangleright \pi\} \\
\llbracket \text{X}\phi \rrbracket_{\delta,r}^n &= \{t \mid t(1) \in \llbracket \phi \rrbracket_{\delta,r}^n\} \\
\llbracket \phi_1 \text{U}\phi_2 \rrbracket_{\delta,r}^n &= \{t \mid \exists i \in 0..n. t(i) \in \llbracket \phi_2 \rrbracket_{\delta,r}^n \wedge \forall j \in 0..i-1. t(j) \in \llbracket \phi_1 \rrbracket_{\delta,r}^n\}
\end{aligned}$$

The semantics is mostly standard, except for $\text{Pr}_{\triangleright\pi}[\psi]$ and $\phi_1 \text{U}\phi_2$. The semantics of $\text{Pr}_{\triangleright\pi}[\psi]$ adds $r \cdot \delta$ to the probability of satisfying ψ , which relaxes or strengthens (depending on r) the probability bound as needed. The semantics of $\phi_1 \text{U}\phi_2$ uses the parameter n to bound the number of steps within which ϕ_2 must hold.

Our semantics enjoys monotonicity. The semantics of state and path formulae is increasing w.r.t. δ if $r = +1$, and decreasing otherwise. The semantics also increases when moving from $r = -1$ to $r = +1$.

Lemma 1 (Monotonicity). *Whenever $\delta \leq \delta'$, we have:*

$$\begin{aligned}
\llbracket \phi \rrbracket_{\delta,+1}^n &\subseteq \llbracket \phi \rrbracket_{\delta',+1}^n & \llbracket \phi \rrbracket_{\delta',-1}^n &\subseteq \llbracket \phi \rrbracket_{\delta,-1}^n & \llbracket \phi \rrbracket_{\delta,-1}^n &\subseteq \llbracket \phi \rrbracket_{\delta,+1}^n \\
\llbracket \psi \rrbracket_{\delta,+1}^n &\subseteq \llbracket \psi \rrbracket_{\delta',+1}^n & \llbracket \psi \rrbracket_{\delta',-1}^n &\subseteq \llbracket \psi \rrbracket_{\delta,-1}^n & \llbracket \psi \rrbracket_{\delta,-1}^n &\subseteq \llbracket \psi \rrbracket_{\delta,+1}^n
\end{aligned}$$

Note that monotonicity does *not* hold for the parameter n , i.e. even if $n \leq n'$, we can *not* conclude $\llbracket \phi \rrbracket_{\delta,+1}^n \subseteq \llbracket \phi \rrbracket_{\delta,+1}^{n'}$. As a counterexample, let $\mathcal{Q} = \{q_0, q_1\}$, $\ell(q_0) = \emptyset$, $\ell(q_1) = \{\mathbf{a}\}$, $\text{Pr}(q_0, q_1) = \text{Pr}(q_1, q_1) = 1$, and $\text{Pr}(q, q') = 0$ elsewhere. Given $\phi = \text{Pr}_{\leq 0}[\text{true U } \mathbf{a}]$, we have $q_0 \in \llbracket \phi \rrbracket_{0,+1}^0$ since in $n = 0$ steps it is impossible to reach a state satisfying \mathbf{a} . However, we do *not* have $q_0 \in \llbracket \phi \rrbracket_{0,+1}^1$ since in $n' = 1$ steps we always reach q_1 , which satisfies \mathbf{a} .

Example 1. We compare an ideal combination padlock to a real one from the point of view of an adversary. The ideal padlock has a single state q_{ok} , representing a closed padlock that can not be opened. Instead, the real padlock is under attack from the adversary who tries to open the padlock by repeatedly guessing its 5-digit PIN. At each step the adversary generates a (uniformly) random




Fig. 2. A Markov chain modelling an ideal (left) and a real (right) padlock.

PIN, different from all the ones which have been attempted so far, and tries to open the padlock with it. The states of the real padlock are q_0, \dots, q_{N-1} (with $N = 10^5$), where q_i represents the situation where i unsuccessful attempts have been made, and an additional state q_{err} that represents that the padlock was opened.

Since after i attempts the adversary needs to guess the correct PIN among the $N - i$ remaining combinations, the real padlock in state q_i moves to q_{err} with probability $1/(N - i)$, and to q_{i+1} with the complementary probability.

Summing up, we simultaneously model both the ideal and real padlock as a single DTMC with the following transition probability function (see Fig. 2):

$$\begin{aligned}
 \Pr(q_{ok}, q_{ok}) &= 1 \\
 \Pr(q_{err}, q_{err}) &= 1 \\
 \Pr(q_i, q_{err}) &= 1/(N - i) && 0 \leq i < N \\
 \Pr(q_i, q_{i+1}) &= 1 - 1/(N - i) && 0 \leq i < N - 1 \\
 \Pr(q, q') &= 0 && \text{otherwise}
 \end{aligned}$$

We label the states with labels $\mathcal{L} = \{err\}$ by letting $\ell(q_{err}) = \{err\}$ and $\ell(q) = \emptyset$ for all $q \neq q_{err}$.

The PCTL formula $\phi = \Pr_{\leq 0}[\text{true U } err]$ models the expected behaviour of an unbreakable padlock, requiring that the set of traces where the padlock is eventually opened has zero probability. Formally, ϕ is satisfied by state q when

$$\begin{aligned}
 q \in \llbracket \phi \rrbracket_{\delta, +1}^n &\iff q \in \llbracket \neg \Pr_{> 0}[\text{true U } err] \rrbracket_{\delta, +1}^n \\
 &\iff q \notin \llbracket \Pr_{> 0}[\text{true U } err] \rrbracket_{\delta, -1}^n \\
 &\iff \neg(\Pr(Cyl(q) \cap \llbracket \text{true U } err \rrbracket_{\delta, -1}^n) - \delta > 0) \\
 &\iff \Pr(Cyl(q) \cap \llbracket \text{true U } err \rrbracket_{\delta, -1}^n) \leq \delta
 \end{aligned} \tag{1}$$

When $q = q_{ok}$ we have that $Cyl(q_{ok}) \cap \llbracket \text{true U } err \rrbracket_{\delta, -1}^n = \emptyset$, hence the above probability is zero, which is surely $\leq \delta$. Consequently, ϕ is satisfied by the ideal padlock q_{ok} , for all $n \geq 0$ and $\delta \geq 0$.

By contrast, ϕ is not always satisfied by the real padlock $q = q_0$, since we have $q_0 \in \llbracket \phi \rrbracket_{\delta, +1}^n$ only for some values of n and δ . To show why, we start by considering some trivial cases. Choosing $\delta = 1$ makes Eq. (1) trivially true for all n . Furthermore, if we choose $n = 1$, then $Cyl(q_0) \cap \llbracket \text{true U } err \rrbracket_{\delta, -1}^n = \{q_0 q_{err}^\omega\}$

is a set of traces with probability $1/N$. Therefore, equation (1) holds only when $\delta \geq 1/N$. More in general, when $n \geq 1$, we have

$$\text{Cyl}(q_0) \cap [\text{true} \cup \text{err}]_{\delta, -1}^n = \{q_0 q_{\text{err}}^\omega, q_0 q_1 q_{\text{err}}^\omega, q_0 q_1 q_2 q_{\text{err}}^\omega, \dots, q_0 \dots q_{n-1} q_{\text{err}}^\omega\}$$

The probability of the above set is the probability of guessing the PIN within n steps. The complementary event, i.e. not guessing the PIN for n times, has probability

$$\frac{N-1}{N} \cdot \frac{N-2}{N-1} \dots \frac{N-n}{N-(n-1)} = \frac{N-n}{N}$$

Consequently, (1) simplifies to $n/N \leq \delta$, suggesting the least value of δ (depending on n) for which q_0 satisfies ϕ . For instance, when $n = 10^3$, this amounts to claiming that the real padlock is secure, up to an error of $\delta = n/N = 10^{-2}$.

3 Up-to- n, δ Bisimilarity

We now define a relation on states $q \sim_\delta^n q'$ that intuitively holds whenever q and q' exhibit similar behaviour for a bounded number of steps. The parameter n controls the number of steps, while δ controls the error allowed in each step. Note that since we only observe the first n steps, our notion is *inductive*, unlike unbounded bisimilarity which is co-inductive, similarly to [12].

Definition 5 (Up-to- n, δ Bisimilarity). We define the relation $q \sim_\delta^n q'$ as follows by induction on n :

1. $q \sim_\delta^0 q'$ always holds
2. $q \sim_\delta^{n+1} q'$ holds if and only if, for all $Q \subseteq \Omega$:
 - (a) $\ell(q) = \ell(q')$
 - (b) $\Pr(q, Q) \leq \Pr(q', \sim_\delta^n(Q)) + \delta$
 - (c) $\Pr(q', Q) \leq \Pr(q, \sim_\delta^n(Q)) + \delta$

where $\sim_\delta^n(Q) = \{q' \mid \exists q \in Q. q \sim_\delta^n q'\}$ is the image of the set Q according to the bisimilarity relation.

We now establish two basic properties of the bisimilarity. Our notion is reflexive and symmetric, and enjoys a triangular property. Furthermore, it is monotonic on both n and δ .

Lemma 2. The relation \sim satisfies:

$$q \sim_\delta^n q \quad q \sim_\delta^n q' \implies q' \sim_\delta^n q \quad q \sim_\delta^n q' \wedge q' \sim_{\delta'}^n q'' \implies q \sim_{\delta+\delta'}^n q''$$

Proof. Straightforward induction on n . □

Lemma 3 (Monotonicity).

$$n \geq n' \wedge \delta \leq \delta' \wedge p \sim_\delta^n q \implies p \sim_{\delta'}^{n'} q$$

Example 2. We use up-to- n, δ bisimilarity to compare the behaviour of the ideal padlock q_{ok} and the real one, in any of its states, when observed for n steps. When $n = 0$ bisimilarity trivially holds, so below we only consider $n > 0$.

We start from the simplest case: bisimilarity does not hold between q_{ok} and q_{err} . Indeed, q_{ok} and q_{err} have distinct labels ($\ell(q_{\text{ok}}) = \emptyset \neq \{\text{err}\} = \ell(q_{\text{err}})$), hence we do not have $q_{\text{ok}} \sim_{\delta}^n q_{\text{err}}$, no matter what $n > 0$ and δ are.

We now compare q_{ok} with any q_i . When $n = 1$, both states have an empty label set, i.e. $\ell(q_{\text{ok}}) = \ell(q_i) = \emptyset$, hence they are bisimilar for any error δ . We therefore can write $q_{\text{ok}} \sim_{\delta}^1 q_i$ for any $\delta \geq 0$.

When $n = 2$, we need a larger error δ to make q_{ok} and q_i bisimilar. Indeed, if we perform a move from q_i , the padlock can be broken with probability $1/(N - i)$, in which case we reach q_{err} , thus violating bisimilarity. Accounting for such probability, we only obtain $q_{\text{ok}} \sim_{\delta}^2 q_i$ for any $\delta \geq 1/(N - i)$.

When $n = 3$, we need an even larger error δ to make q_{ok} and q_i bisimilar. Indeed, while the first PIN guessing attempt has probability $1/(N - i)$, in the second move the guessing probability increases to $1/(N - i - 1)$. Choosing δ equal to the largest probability is enough to account for both moves, hence we obtain $q_{\text{ok}} \sim_{\delta}^3 q_i$ for any $\delta \geq 1/(N - i - 1)$. Technically, note that the denominator $N - i - 1$ might be zero, since when $i = n - 1$ the first move always guesses the PIN, and the second guess never actually happens. In such case, we instead take $\delta = 1$.

More in general, for an arbitrary $n \geq 2$, we obtain through a similar argument that $q_{\text{ok}} \sim_{\delta}^n q_i$ for any $\delta \geq 1/(N - i - n + 2)$. Intuitively, $\delta = 1/(N - i - n + 2)$ is the probability of guessing the PIN in the last attempt (the n -th), which is the attempt having the highest success probability. Again, when the denominator $N - i - n + 2$ becomes zero (or negative), we instead take $\delta = 1$.

Note that the DTMC of the ideal and real padlocks (Example 1) has finitely many states. Our bisimilarity notion and results, however, can also deal with DTMCs with a countably infinite set of states, as we show in the next example.

Example 3. We consider an ideal system which randomly generates bit streams in a fair way. We model such a system as having two states $\{q_a, q_b\}$, with transition probabilities $\Pr(x, y) = 1/2$ for any $x, y \in \{q_a, q_b\}$, as in Fig. 1. We label state q_a with label a denoting bit 0, and state q_b with label b denoting bit 1.




Fig. 3. A Markov chain modelling an unfair random generator of bit streams.

We compare this ideal system with a real system which generates bit streams in an unfair way. At each step, the real system draws a ball from an urn, initially having g_0 **a**-labelled balls and g_0 **b**-labelled balls. After each drawing, the ball is placed back in the urn. However, every time an **a**-labelled ball is drawn, an additional **a**-labelled ball is put in the urn, making the next drawings more biased towards **a**.

We model the real system using the infinite¹ set of states $\mathbb{N} \times \{\mathbf{a}, \mathbf{b}\}$, whose first component counts the number of **a**-labelled balls in the urn, and the second component is the label of the last-drawn ball. The transition probabilities are as follows, where $g_0 \in \mathbb{N}^+$ (see Fig. 3):

$$\begin{aligned} \Pr((g, x), (g + 1, \mathbf{a})) &= g / (g + g_0) \\ \Pr((g, x), (g, \mathbf{b})) &= g_0 / (g + g_0) \\ \Pr((g, x), (g', x')) &= 0 \quad \text{otherwise} \end{aligned}$$

We label each such state with its second component.

We now compare the ideal system to the real one. Intuitively, the ideal system, when started from state q_a , produces a sequence of states whose labels are uniform independent random values in $\{\mathbf{a}, \mathbf{b}\}$. Instead, the real system slowly becomes more and more biased towards label **a**. More precisely, when started from state (g_0, \mathbf{a}) , in the first drawing the next label is uniformly distributed between **a** and **b**, as in the ideal system. When the sampled state has label **a**, this causes the component g to be incremented, increasing the probability $g / (g + g_0)$ of sampling another **a** in the next steps. Indeed, the value g is always equal to g_0 plus the number of sampled **a**-labelled states so far.

Therefore, unlike the ideal system, on the long run the real system will visit **a**-labelled states with very high probability, since the g component slowly but steadily increases. While this fact makes the two systems *not* bisimilar according to the standard probabilistic bisimilarity [28], if we restrict the number of steps to $n \ll g_0$ and tolerate a small error δ , we can obtain $q_a \sim_{\delta}^n (g_0, \mathbf{a})$.

For instance, if we let $g_0 = 1000$, $n = 100$ and $\delta = 0.05$ we have $q_a \sim_{\delta}^n (g_0, \mathbf{a})$. This is because, in n steps, the first component g of a real system (g, x) will at most reach 1100, making the probability of the next step to be $(g + 1, \mathbf{a})$ to be at most $1100/2100 \simeq 0.523$. This differs from the ideal probability 0.5 by less than δ , hence bisimilarity holds.

4 Soundness

Our soundness theorem shows that, if we consider any state q satisfying ϕ (with steps n and error δ'), and any state q' which is bisimilar to q (with enough steps and error δ), then q' must satisfy ϕ , with the same number n of steps, at the cost of suitably increasing the error. For a fixed ϕ , the “large enough” number of steps and the increase in the error depend linearly on n .

¹ Modelling this behaviour inherently requires an *infinite* set of states, since each number of **a**-labelled balls in the urn leads to a unique transition probability function.

Theorem 1 (Soundness). *Given a formula ϕ , let $k_X = X_{\max}(\phi)$ be the maximum X -nesting of ϕ , and let $k_U = U_{\max}(\phi)$ be the maximum U -nesting of ϕ . Then, for all n, δ, δ' we have:*

$$\sim_{\delta}^{\bar{n}} (\llbracket \phi \rrbracket_{\delta', +1}^n) \subseteq \llbracket \phi \rrbracket_{\bar{n} \cdot \delta + \delta', +1}^n \quad \text{where } \bar{n} = n \cdot k_U + k_X + 1$$

Example 4. We now apply Theorem 1 to our padlock system in the running example. To do so, we take the same formula $\phi = \text{Pr}_{\leq 0}[\text{true } U \text{ err}]$ of Example 1 and choose $n = 10^3$ and $\delta' = 0$. Since ϕ has only one until operator and no next operators, the value \bar{n} in the statement of the theorem is $\bar{n} = 10^3 \cdot 1 + 0 + 1 = 1001$. Therefore, from Theorem 1 we obtain, for all δ :

$$\sim_{\delta}^{1001} (\llbracket \phi \rrbracket_{0, +1}^{1000}) \subseteq \llbracket \phi \rrbracket_{1001 \cdot \delta, +1}^{1000}$$

In Example 1 we discussed how the ideal padlock q_{ok} satisfies the formula ϕ for any number of steps and any error value. In particular, choosing 1000 steps and zero error, we get $q_{\text{ok}} \in \llbracket \phi \rrbracket_{0, +1}^{1000}$.

Moreover, in Example 2 we observed that states q_{ok} and q_0 are bisimilar with $\bar{n} = 1001$ and $\delta = 1/(N - 0 - \bar{n} + 2) = 1/99001$, i.e. $q_{\text{ok}} \sim_{\delta}^{\bar{n}} q_0$.

In such case, the theorem ensures that $q_0 \in \llbracket \phi \rrbracket_{1001/99001, +1}^{1000}$, hence the real padlock can be considered unbreakable if we limit our attention to the first $n = 1000$ steps, up to an error of $1001/99001 \approx 0.010111$. Finally, we note that such error is remarkably close to the least value that would still make q_0 satisfy ϕ , which we computed in Example 1 as $n/N = 10^3/10^5 = 0.01$.

In the rest of this section, we describe the general structure of the proof in a top-down fashion, leaving the detailed proof for [5].

We prove the soundness theorem by induction on the state formula ϕ , hence we also need to deal with path formulae ψ . Note that the statement of the theorem considers the image of the semantics of the state formula ϕ w.r.t. bisimilarity (i.e., $\sim_{\delta}^{\bar{n}} (\llbracket \phi \rrbracket_{\delta', +1}^n)$). Analogously, to deal with path formulae we also need an analogous notion on sets of traces. To this purpose, we consider the set of traces in the definition of the semantics: $T = \text{Cyl}(p) \cap \llbracket \psi \rrbracket_{\delta, r}^n$. Then, given a state q bisimilar to p , we define the set of *pointwise bisimilar traces* starting from q , which we denote with $\tilde{R}_{\delta, q}^n(T)$. Technically, since ψ can only observe a finite portion of a trace, it is enough to define $\tilde{R}_{\delta, q}^n(\tilde{T})$ on sets of *trace fragments* \tilde{T} .

Definition 6. *Write $F_{q_0}^n$ for the set of all trace fragments of length n starting from q_0 . Assuming $p \sim_{\delta}^{\bar{n}} q$, we define $\tilde{R}_{\delta, q}^n(\cdot) : \mathcal{P}(F_p^n) \rightarrow \mathcal{P}(F_q^n)$ as follows:*

$$\tilde{R}_{\delta, q}^n(\tilde{T}) = \{\tilde{u} \in F_q^n \mid \exists \tilde{t} \in \tilde{T}. \forall 0 \leq i < n. \tilde{t}(i) \sim_{\delta}^{\bar{n}-i} \tilde{u}(i)\}$$

The key inequality we exploit in the proof (Lemma 4) compares the probability of a set of trace fragments \tilde{T} starting from p to the one of the related set of trace fragments $\tilde{R}_{\delta, q}^n(\tilde{T})$ starting from a q bisimilar to p . We remark that the component $\bar{n}\delta$ in the error that appears in Theorem 1 results from the component $m\delta$ appearing in the following lemma.

Lemma 4. *If $p \sim_{\delta}^n q$ and \tilde{T} is a set of trace fragments of length m , with $m \leq n$, starting from p , then:*

$$\Pr(\tilde{T}) \leq \Pr(\tilde{R}_{\delta, q}^m(\tilde{T})) + m\delta$$

Lemma 4 allows \tilde{T} to be an infinite set (because the set of states Ω can be infinite). We reduce this case to that where \tilde{T} is finite. We first recall a basic calculus property: any inequality $a \leq b$ can be proved by establishing instead $a \leq b + \epsilon$ for all $\epsilon > 0$. Then, since the probability distribution of trace fragments of length m is discrete, for any $\epsilon > 0$ we can always take a finite subset of the infinite set \tilde{T} whose probability differs from that of \tilde{T} less than ϵ . It is therefore enough to consider the case where \tilde{T} is finite, as done in the following lemma.

Lemma 5. *If $p \sim_{\delta}^n q$ and \tilde{T} is a finite set of trace fragments of length $n > 0$ starting from p , then:*

$$\Pr(\tilde{T}) \leq \Pr(\tilde{R}_{\delta, q}^n(\tilde{T})) + n\delta$$

We prove Lemma 5 by induction on n . In the inductive step, we partition the traces according to their first move, i.e., on their next state after p (for the trace fragments in T) or q (for the bisimilar counterparts). A main challenge here is caused by the probabilities of such moves being weakly connected. Indeed, when p moves to p' , we might have several states q' , bisimilar to p' , such that q moves to q' . Worse, when p moves to another state p'' , we might find that some of the states q' we met before are also bisimilar to p'' . Such overlaps make it hard to connect the probability of p moves to that of q moves.

To overcome these issues, we exploit the technical lemma below. Let set A represent the p moves, and set B represent the q moves. Then, associate to each set element $a \in A, b \in B$ a value $(f_A(a), f_B(b))$ in the lemma) representing the move probability. The lemma ensures that each $f_A(a)$ can be expressed as a weighted sum of $f_B(b)$ for the elements b bisimilar to a . Here, the weights $h(a, b)$ make it possible to relate a p move to a “weighted set” of q moves. Furthermore, the lemma ensures that no $b \in B$ has been cumulatively used for more than a unit weight ($\sum_{a \in A} h(a, b) \leq 1$).

Lemma 6. *Let A be a finite set and B be a countable set, equipped with functions $f_A : A \rightarrow \mathbb{R}_0^+$ and $f_B : B \rightarrow \mathbb{R}_0^+$. Let $g : A \rightarrow 2^B$ be such that $\sum_{b \in g(a)} f_B(b)$ converges for all $a \in A$. If, for all $A' \subseteq A$:*

$$\sum_{a \in A'} f_A(a) \leq \sum_{b \in \bigcup_{a \in A'} g(a)} f_B(b) \quad (2)$$

then there exists $h : A \times B \rightarrow [0, 1]$ such that:

$$\forall b \in B : \sum_{a \in A} h(a, b) \leq 1 \quad (3)$$

$$\forall A' \subseteq A : \sum_{a \in A'} f_A(a) = \sum_{a \in A'} \sum_{b \in g(a)} h(a, b) f_B(b) \quad (4)$$




Fig. 4. Graphical representation of Lemma 6 (left) and its proof (right).

We visualize Lemma 6 in Fig. 4 through an example. The leftmost graph shows a finite set $A = \{a_1, a_2, a_3\}$ where each a_i is equipped with its associated value $f_A(a_i)$ and, similarly, a finite set $B = \{b_1, \dots, b_4\}$ where each b_i has its own value $f_B(b_i)$. The function g is rendered as the edges of the graph, connecting each a_i with all $b_j \in g(a_i)$.

The graph satisfies the hypotheses, as one can easily verify. For instance, when $A' = \{a_1, a_2\}$ inequality (2) simplifies to $0.3 + 0.5 \leq 0.5 + 0.6$. The thesis ensures the existence of a weight function $h(-, -)$ whose values are shown in the graph on the left over each edge.

These values indeed satisfy (3): for instance, if we pick $b = b_2$ the inequality reduces to $0.5 + 0.16 \leq 1$. Furthermore, (4) is also satisfied: for instance, taking $A' = \{a_2\}$ the equation reduces to $0.5 = 0.4 \cdot 0.5 + 0.5 \cdot 0.6$, while taking $A' = \{a_3\}$ the equation reduces to $0.2 = 0.16 \cdot 0.6 + 1.0 \cdot 0.05 + 1.0 \cdot 0.05$.

The rightmost graph in Fig. 4 instead sketches how our proof devises the desired weight function h , by constructing a network flow problem, and exploiting the well-known min-cut/max-flow theorem [14], following the approach of [4]. We start by adding a source node to the right (white bullet in the figure), connected to nodes in B , and a sink node to the left, connected to nodes in A . We write the capacity over each edge: we use $f_B(b_i)$ for the edges connected to the source, $f_A(a_i)$ for the edges connected to the sink, and $+\infty$ for the other edges in the middle.

Then, we argue that the leftmost cut C shown in the figure is a min-cut. Intuitively, if we take another cut C' not including some edge in C , then C' has to include other edges making C' not any better than C . Indeed, C' can surely not include any edge in the middle, since they have $+\infty$ capacity. Therefore, if C' does not include an edge from some a_i to the sink, it has to include all the edges from the source to each $b_j \in g(a_i)$. In this case, hypothesis (2) ensures that doing so does not lead to a better cut. Hence, C is indeed a min-cut.

From the max-flow corresponding to the min-cut, we derive the values for $h(-, -)$. Thesis (3) follows from the flow conservation law on each b_i , and the fact that the incoming flow of each b_j from the source is bounded by the capacity of the related edge. Finally, thesis (4) follows from the flow conservation law on

each a_i , and the fact that the outgoing flow of each a_i to the sink is exactly the capacity of the related edge, since the edge is on a min-cut.

5 Conclusions

In this paper we studied how the (relaxed) semantics of PCTL formulae interacts with (approximate) probabilistic bisimulation. In the regular, non relaxed case, it is well-known that when a state q satisfies a PCTL formula ϕ , then all the states that are probabilistic-bisimilar to q also satisfy ϕ ([19]). Theorem 1 extends this to the relaxed semantics, establishing that when a state q satisfies a PCTL formula ϕ up-to n steps and error δ , then all the states that are approximately probabilistic bisimilar to q with error δ' (and enough steps) also satisfy ϕ up-to n steps and suitably increased error. We provide a way to compute the new error in terms of n, δ, δ' .

Our results are a first step towards a novel approach to the security analysis of cryptographic protocols using probabilistic bisimulations. Ideally, when one is able to prove that a real-world specification of a cryptographic protocol is bisimilar to an ideal one, with an error that is asymptotically negligible as long as the number of steps is kept polynomial, then one can invoke Theorem 1 and claim that the two models satisfy the same PCTL formulae (with negligible error and a polynomial number of steps), essentially reducing the security proof of the cryptographic protocol to verifying the ideal model. A relevant line for future work is to study the applicability of our theory in this setting. Related to this, one would like to investigate proof techniques for establishing approximate bisimilarity and refinement [27].

Another possible line of research would be devising algorithms for approximate bisimilarity, along the lines of [11, 13, 22, 34–36]. This direction, however, would require restricting our theory to finite-state systems, which contrasts with our general motivation coming from cryptographic security. Indeed, in the analysis of cryptographic protocols, security is usually to be proven against an arbitrary adversary, hence also against infinite-state ones. Hence, model-checking of finite-state systems would not directly be applicable in this setting.

Acknowledgements. Massimo Bartoletti is partially supported by Conv. Fondazione di Sardegna & Atenei Sardi projects F74I19000900007 *ADAM* and F75F21001220007 *ASTRID*. Maurizio Murgia and Roberto Zunino are partially supported by MIUR PON *Distributed Ledgers for Secure Open Communities*.

References

1. de Alfaro, L., Henzinger, T.A., Majumdar, R.: Discounting the future in systems theory. In: Baeten, J.C.M., Lenstra, J.K., Parrow, J., Woeginger, G.J. (eds.) ICALP 2003. LNCS, vol. 2719, pp. 1022–1037. Springer, Heidelberg (2003). https://doi.org/10.1007/3-540-45061-0_79
2. de Alfaro, L., Majumdar, R., Raman, V., Stoelinga, M.: Game refinement relations and metrics. *Log. Methods Comput. Sci.* 4(3), 1–28 (2008)

3. Bacci, G., Bacci, G., Larsen, K.G., Mardare, R., Tang, Q., van Breugel, F.: Computing probabilistic bisimilarity distances for probabilistic automata. *Log. Methods Comput. Sci.* **17**(1), 9:1-9:36 (2021)
4. Baier, C.: Algorithmic verification methods for probabilistic systems. Habilitation Thesis, Universität Mannheim (1998). https://www.inf.tu-dresden.de/content/institutes/thi/algi/publikationen/texte/15_98.pdf
5. Bartoletti, M., Murgia, M., Zunino, R.: A sound up-to- n, δ bisimilarity for PCTL. *CoRR abs/2111.03117* (2021)
6. Bian, G., Abate, A.: On the relationship between bisimulation and trace equivalence in an approximate probabilistic context. In: Esparza, J., Murawski, A.S. (eds.) *FoSSaCS 2017*. LNCS, vol. 10203, pp. 321–337. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-662-54458-7_19
7. van Breugel, F.: Probabilistic bisimilarity distances. *ACM SIGLOG News* **4**(4), 33–51 (2017)
8. van Breugel, F., Hermida, C., Makkai, M., Worrell, J.: An accessible approach to behavioural pseudometrics. In: Caires, L., Italiano, G.F., Monteiro, L., Palamidessi, C., Yung, M. (eds.) *ICALP 2005*. LNCS, vol. 3580, pp. 1018–1030. Springer, Heidelberg (2005). https://doi.org/10.1007/11523468_82
9. van Breugel, F., Sharma, B., Worrell, J.: Approximating a behavioural pseudometric without discount for probabilistic systems. *Log. Methods Comput. Sci.* **4**(2) (2008). [https://doi.org/10.2168/LMCS-4\(2:2\)2008](https://doi.org/10.2168/LMCS-4(2:2)2008)
10. van Breugel, F., Worrell, J.: A behavioural pseudometric for probabilistic transition systems. *Theor. Comput. Sci.* **331**(1), 115–142 (2005). <https://doi.org/10.1016/j.tcs.2004.09.035>
11. van Breugel, F., Worrell, J.: The complexity of computing a bisimilarity pseudometric on probabilistic automata. In: van Breugel, F., Kashefi, E., Palamidessi, C., Rutten, J. (eds.) *Horizons of the Mind. A Tribute to Prakash Panangaden*. LNCS, vol. 8464, pp. 191–213. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-06880-0_10
12. Castiglioni, V., Gebler, D., Tini, S.: Logical characterization of bisimulation metrics. In: *Workshop on Quantitative Aspects of Programming Languages and Systems (QAPL), EPTCS*, vol. 227, pp. 44–62 (2016). <https://doi.org/10.4204/EPTCS.227.4>
13. Chen, D., van Breugel, F., Worrell, J.: On the complexity of computing probabilistic bisimilarity. In: Birkedal, L. (ed.) *FoSSaCS 2012*. LNCS, vol. 7213, pp. 437–451. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-28729-9_29
14. Dantzig, G.B., Fulkerson, D.R.: *On the Max Flow Min Cut Theorem of Networks*. RAND Corporation, Santa Monica, CA (1955)
15. Deng, Y., Chothia, T., Palamidessi, C., Pang, J.: Metrics for action-labelled quantitative transition systems. *Electron. Notes Theor. Comput. Sci.* **153**(2), 79–96 (2006). <https://doi.org/10.1016/j.entcs.2005.10.033>
16. 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>
17. Desharnais, J., Gupta, V., Jagadeesan, R., Panangaden, P.: Metrics for labeled Markov systems. In: Baeten, J.C.M., Mauw, S. (eds.) *CONCUR 1999*. LNCS, vol. 1664, pp. 258–273. Springer, Heidelberg (1999). https://doi.org/10.1007/3-540-48320-9_19
18. Desharnais, J., Gupta, V., Jagadeesan, R., Panangaden, P.: Metrics for labelled Markov processes. *Theor. Comput. Sci.* **318**(3), 323–354 (2004). <https://doi.org/10.1016/j.tcs.2003.09.013>

19. Desharnais, J., Gupta, V., Jagadeesan, R., Panangaden, P.: Weak bisimulation is sound and complete for pCTL*. *Inf. Comput.* **208**(2), 203–219 (2010). <https://doi.org/10.1016/j.ic.2009.11.002>
20. Desharnais, J., Laviolette, F., Tracol, M.: Approximate analysis of probabilistic processes: Logic, simulation and games. In: Quantitative Evaluation of Systems (QEST), pp. 264–273. IEEE Computer Society (2008). <https://doi.org/10.1109/QEST.2008.42>
21. D’Innocenzo, A., Abate, A., Katoen, J.: Robust PCTL model checking. In: Dang, T., Mitchell, I.M. (eds.) Hybrid Systems: Computation and Control (HSCC), pp. 275–286. ACM (2012). <https://doi.org/10.1145/2185632.2185673>
22. Fu, H.: Computing game metrics on Markov decision processes. In: Czumaj, A., Mehlhorn, K., Pitts, A., Wattenhofer, R. (eds.) ICALP 2012. LNCS, vol. 7392, pp. 227–238. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-31585-5_23
23. Furber, R., Mardare, R., Mio, M.: Probabilistic logics based on riesz spaces. *Log. Methods Comput. Sci.* **16**(1) (2020). [https://doi.org/10.23638/LMCS-16\(1:6\)2020](https://doi.org/10.23638/LMCS-16(1:6)2020)
24. Giacalone, A., Jou, C., Smolka, S.A.: Algebraic reasoning for probabilistic concurrent systems. In: Broy, M., Jones, C.B. (eds.) Programming concepts and methods: Proceedings of the IFIP Working Group 2.2, 2.3 Working Conference on Programming Concepts and Methods, pp. 443–458. North-Holland (1990)
25. Hansson, H., Jonsson, B.: A logic for reasoning about time and reliability. *Formal Aspects Comput.* **6**(5), 512–535 (1994). <https://doi.org/10.1007/BF01211866>
26. Hermanns, H., Parma, A., Segala, R., Wachter, B., Zhang, L.: Probabilistic logical characterization. *Inf. Comput.* **209**(2), 154–172 (2011). <https://doi.org/10.1016/j.ic.2010.11.024>
27. Jonsson, B., Larsen, K.G.: Specification and refinement of probabilistic processes. In: IEEE Symp. on Logic in Computer Science (LICS), pp. 266–277. IEEE Computer Society Press (1991)
28. Larsen, K.G., Skou, A.: Bisimulation through probabilistic testing. In: ACM Symposium on Principles of Programming Languages (POPL), pp. 344–352. ACM Press (1989). <https://doi.org/10.1145/75277.75307>
29. 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](https://doi.org/10.1016/0890-5401(91)90030-6)
30. Mio, M.: Upper-expectation bisimilarity and Lukasiewicz μ -calculus. In: Muscholl, A. (ed.) FoSSaCS 2014. LNCS, vol. 8412, pp. 335–350. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-54830-7_22
31. Mio, M.: Riesz modal logic with threshold operators. In: Dawar, A., Grädel, E. (eds.) ACM/IEEE Symposium on Logic in Computer Science (LICS), pp. 710–719. ACM (2018). <https://doi.org/10.1145/3209108.3209118>
32. Mio, M., Simpson, A.: Lukasiewicz μ -calculus. *Fundam. Informaticae* **150**(3–4), 317–346 (2017). <https://doi.org/10.3233/FI-2017-1472>
33. Song, L., Zhang, L., Godskesen, J.C., Nielson, F.: Bisimulations meet PCTL equivalences for probabilistic automata. *Logical Methods Comput. Sci.* **9**(2) (2013). [https://doi.org/10.2168/LMCS-9\(2:7\)2013](https://doi.org/10.2168/LMCS-9(2:7)2013)
34. Tang, Q., van Breugel, F.: Computing probabilistic bisimilarity distances via policy iteration. In: International Conference on Concurrency Theory (CONCUR). LIPIcs, vol. 59, pp. 22:1–22:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2016). <https://doi.org/10.4230/LIPIcs.CONCUR.2016.22>

35. Tang, Q., van Breugel, F.: Algorithms to compute probabilistic bisimilarity distances for labelled markov chains. In: International Conference on Concurrency Theory (CONCUR). LIPIcs, vol. 85, pp. 27:1–27:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2017). <https://doi.org/10.4230/LIPIcs.CONCUR.2017.27>
36. Tang, Q., van Breugel, F.: Deciding probabilistic bisimilarity distance one for labelled Markov chains. In: Chockler, H., Weissenbacher, G. (eds.) CAV 2018. LNCS, vol. 10981, pp. 681–699. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-96145-3_39