Robust PCTL Model Checking *

Alessandro D'Innocenzo
Department of Electrical and Information Engineering, Center of Excellence DEWS University of L'Aquila, Italy
alessandro.dinnocenzo@univaq.it

Alessandro Abate
Delft Center for Systems and Control TU Delft – Delft University of Technology, The Netherlands
a.abate@tudelft.nl

Joost-Pieter Katoen
Software Modeling and Verification Group RWTH Aachen University, Germany
katoen@cs.rwth-aachen.de

ABSTRACT

This paper deals with the notion of approximate probabilistic bisimulation (APB) relation for discrete-time labeled Markov Chains (LMC). In order to provide a quantified upper bound on a metric over probabilistic realizations for LMC, we exploit the structure and properties of the APB and leverage the mathematical framework of Markov set-Chains. Based on this bound, the article proves that the existence of an APB implies the preservation of robust PCTL formulae, which are formulae that allow being properly relaxed or strengthened, according to the underlying APB. This leads to a notion of robustness for probabilistic model checking.

Categories and Subject Descriptors

G.3, G.4 [**Mathematics of Computing**]: Probability and Statistics; Mathematical Software.

General Terms

Markov processes; Stochastic processes; Verification.

1. INTRODUCTION

For complex probabilistic systems with large state space dimension, automatic verification of probabilistic properties can be computationally prohibitive. An approach that is successfully used to cope with the issue of computational complexity and scalability that arise in formal verification of complex models is that of abstraction: a system – equivalent in some sense to the original system – with smaller state space or simpler dynamics is sought. The abstraction is usually an aggregated or lumped version of the concrete model. Equivalence is usually introduced with reference to the notion of bisimulation [26, 27], which is an equivalence relation induced on the state space of the original system. System equivalence implies that certain properties of the original (complex) system are preserved by the (simpler) abstraction. Thus a specific property of interest can be checked more efficiently (that is, with a lower computational complexity) on the abstraction [21].

Often though, the exact notion of equivalence appears to be quite conservative, because it requires an exact match between trajectories of the concrete and of the abstract systems. In practice, severe problems may occur with the computation of exact bisimulation due to numerical errors, in particular for models with large state spaces [29]. This issue is even more taxing for probabilistic models: if an abstract system is verified on a model obtained with quantization/discretization errors in the transition probabilities, the effect of a small perturbation can invalidate the outcome of the procedure.

These issues have led to the introduction of an *approximate* notion of equivalence, which for deterministic systems was introduced in [15] via the notion of approximate bisimulation – a notion based on metrics over the distance between trajectories of concrete and abstract models.

In the context of non-deterministic models, the notion of approximate bisimulation has been used to perform model reduction while preserving properties expressed by temporal logics, e.g. TCTL [19]. Such properties are subsequently verified by the use of model checking procedures. In [8], a quantitative version of transition systems is considered, and different versions of trace and bisimulation distance are defined.

For probabilistic systems, the concept of (strong) probabilistic bisimulation has been introduced in [25] over discrete-time, finite-state Markov Chains. Bisimulation corresponds to what is also known as lumping. The use of approximate notions is advocated in [14] and motivated by the robustness issues mentioned above. Of course approximate notions are also likely to result in coarser bisimulations than those obtained with the exact notion. The work in [12] discusses approximate notions of bisimulations for discrete-time labeled Markov processes.

The reference model framework in this paper is that of discrete time labeled Markov Chains (LMC). We tailor the definition of approximate probabilistic bisimulation (APB) in [12] to LMC. As a first result of this paper, we provide a quantitative upper bound for a probabilistic realization

*This work is supported by the European Commission MoVeS project FP7-ICT-2009-5 257005, by the European Commission Marie Curie grant MANTRAS 249295, by the European Commission NoE HYCON2 FP7-ICT-2009-5 257462, and by the NWO VENI grant 016.103.020.

HSCC 2012, April 17–19, 2012, Beijing, China.

metric in time (both over a finite and an infinite horizon) with an expression that depends on the approximation level of the APB. We provide these bounds first by exploiting the structure of the APB, then by employing the theory of Markov set-Chains (MSC) [18]. Since the APB notion induces a coarser bisimulation than the corresponding exact notion, it cannot be used to prove that all PCTL properties of the original system are preserved by an approximately bisimilar abstraction. As the main result of the paper, we prove that an APB with precision ε implies the preservation of ε-robust PCTL properties over the original system. PCTL [2, 22] is a discrete-time probabilistic temporal logic that allows modeling probabilistic specifications, to be then verified by probabilistic model checking using tools such as PRISM [23, 24] or MRMC [20].

The results obtained in this paper for finite LMC represent the first step towards robustness analysis for model checking of infinite state space systems as general as Stochastic Hybrid Systems (SHS). Indeed we aim to leverage the abstraction procedure we developed in [1], where we show how to derive a finite LMC abstraction of a SHS characterized by a finite precision, to subsequently verify properties of SHS by using classical PCTL model checking algorithms over the abstraction. In order to guarantee that the model checking result obtained over the abstraction also holds for the SHS, the accuracy of the abstraction needs to satisfy the constraints discussed in this paper.

The paper is structured as follows. In Section 2, we define LMC and introduce the notions of APB. In Section 3 we state quantified relations between the APB and a probabilistic realization metric both by exploiting the structure and properties of the APB (Section 3.1) and leveraging the mathematical framework of Markov set-Chains (Section 3.3). In Section 4 we recall basic notions of PCTL model checking (Section 4.1), and prove that the presence of an APB with precision ε implies the preservation of ε-robust PCTL formulae (Section 4.2). A case study is developed throughout the paper, to apply and clarify the presented notions and results.

2. APPROXIMATE BISIMULATION

Let AP be a finite, fixed set of atomic propositions.

Definition 1. *[Discrete-time labeled Markov Chain] We define a discrete-time labeled Markov Chain (LMC) as a tuple (Q, P, L) consisting of:*

- *Q, a non-empty set of states of finite cardinality $n \in \mathbb{N}$;*
- *$P : Q \times Q \to [0, 1]$, a stochastic matrix that associates to each pair $(q_1, q_2) \in Q \times Q$ the transition probability from state q_1 to state q_2;*
- *$L : Q \to 2^{AP}$, a labeling function that associates to each state $q \in Q$ the set $L(q)$ of atomic propositions that are valid in q.*

A LMC can be related to a DTMC, as in [3], and is a subclass of the labeled Markov process (LMP) model as in [12].

Consider a LMC $M = (Q, P, L)$ and $k \in \mathbb{N}$. We define by $P^k(q, q')$ the probability that state q' is reached in k steps by an execution of M starting from state q. Given a set $A \in 2^Q$, we define by $P^k(q, A) = \sum_{q' \in A} P^k(q, q')$ the probability that the set of states A is reached in k steps starting from state q.

Given a LMC M, a relation $\Gamma \subseteq Q \times Q$ and a set $A \in 2^Q$, we introduce the set

$$\Gamma(A) = \{q \in Q \mid \exists q' \in A, (q, q') \in \Gamma\},$$

and say that A is Γ-closed if $\Gamma(A) \subseteq A$. The following definition is inspired by [10, 25].

Definition 2. *[Probabilistic bisimulation] Given a LMC $M = (Q, P, L)$, a probabilistic bisimulation is an equivalence relation $\Gamma \subseteq Q \times Q$ such that for any $(q_1, q_2) \in \Gamma$ then $L(q_1) = L(q_2)$, and for any Γ-closed set $A \in 2^Q$,*

$$P(q_1, A) = P(q_2, A).$$

States $q_1, q_2 \in Q$ are probabilistic bisimilar, which is denoted by $q_1 \equiv q_2$, if there exists a probabilistic bisimulation Γ with $(q_1, q_2) \in \Gamma$.

The condition $P(q_1, A) = P(q_2, A)$ in the above definition is equivalent to the condition $P(q_1, \Gamma(A)) = P(q_2, \Gamma(A))$, since for an equivalence Γ we have $A = \Gamma(A)$. Any probabilistic bisimulation relation induces a partition of the state space, where the equivalence classes are made up of bisimilar states [3]. In particular, the equivalence classes are given by the Γ-closed sets $\{A_1, \ldots, A_m\}$. In the following, we will denote by m the number of Γ-closed sets that form a partition of the state space. A recursive algorithm to compute the maximal (coarsest) bisimulation of a LMC with time complexity $\mathcal{O}(|Q|^2 \log |Q|)$ has been proposed in [9]. Note that Definition 2 hinges on rather strong conditions on the transition probabilities, marginalized over the equivalence classes.

Example 1. *Craps is a dice game where the players bet on the outcome of dice rolls. An LMC (Q, P, L) characterizes the dynamics (the possible outcomes) of the game – its transition probability matrix P is reported in Figure 1 [3, Section 10]. Given a set $AP = \{\texttt{start}, \texttt{mid}, \texttt{won}, \texttt{lost}\}$, the Markov Chain states*

$$Q = \{\texttt{start}, 4, 10, 5, 9, 6, 8, \texttt{won}, \texttt{lost}\}$$

are associated to the following labels:

$$L(\texttt{start}) = \texttt{start}, \quad L(i) = \texttt{mid}, i \in \{4, 10, 5, 9, 6, 8\},$$
$$L(\texttt{won}) = \texttt{won}, \quad L(\texttt{lost}) = \texttt{lost}.$$

The Markov Chain admits an (exact) probabilistic bisimulation, depicted in Figure 2, with the following collection of

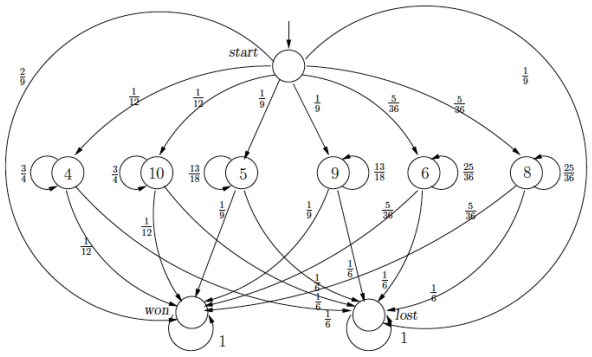

Figure 1: Concrete Markov Chain of the craps game.

6 *equivalence classes:*

$$\{\texttt{start}\}, \{4, 10\}, \{5, 9\}, \{6, 8\}, \{\texttt{won}\}, \{\texttt{lost}\}.$$

Each pair of vertices within an equivalence class denotes probabilistically bisimilar states.

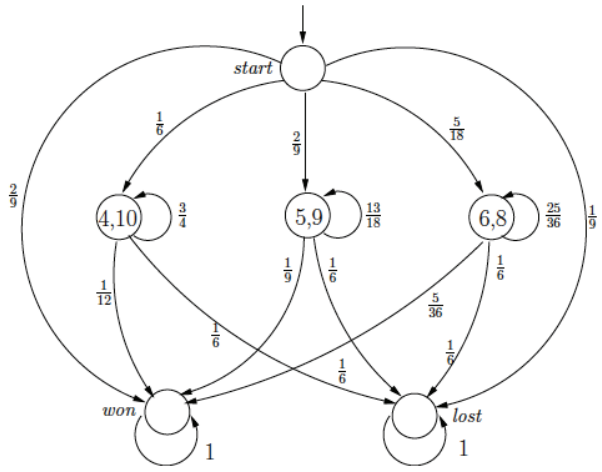

Figure 2: Bisimilar LMC for craps game.

Notice how even small perturbations on the transition probabilities or possible numerical approximations in matrix P would invalidate the exact probabilistic bisimulation relation of the above example. The concept of approximate probabilistic bisimulation has been introduced in [12] to overcome the above limitations. The definition in [12, Definition 10] has been introduced for LMP, and can be shown to directly tailor to the one given here for LMC.

DEFINITION 3. *[Approximate probabilistic bisimulation] Given a LMC $M = (Q, P, L)$, an approximate probabilistic bisimulation with precision $0 \leq \varepsilon \leq 1$ (APB with precision ε) is a reflexive and symmetric relation $\Gamma_\varepsilon \subseteq Q \times Q$ such that for any $(q_1, q_2) \in \Gamma_\varepsilon$, then $L(q_1) = L(q_2)$ and for any Γ_ε-closed set $A \in 2^Q$,*

$$|P(q_1, \Gamma_\varepsilon(A)) - P(q_2, \Gamma_\varepsilon(A))| \leq \varepsilon,$$

States $q_1, q_2 \in Q$ are probabilistic approximately bisimilar with precision ε, which is denoted by $q_1 \equiv_\varepsilon q_2$, if there exists an APB Γ_ε with precision ε, such that $(q_1, q_2) \in \Gamma_\varepsilon$.

It is easy to see that an APB with precision $\varepsilon = 0$ is an exact probabilistic bisimulation in the sense of Definition 2. For any given LMC with state space Q and approximation parameter $\varepsilon \in [0, 1]$, techniques to compute the largest APB with precision ε in time $\mathcal{O}(|Q|^7)$ are introduced in [12].

EXAMPLE 2. *Let us consider a perturbed LMC $(Q, \tilde{P}, L)$ of (Q, P, L) in Example 1, which is defined as follows:*

$$\tilde{P} = P + E,$$

where $E(i, j) = \varepsilon_{i,j}, |\varepsilon_{i,j}| \leq \min\{P(i, j), 1 - P(i, j)\}$, and $\sum_{j \in Q} \varepsilon_{i,j} = 0$. As an example, consider

$$\tilde{P}(\texttt{start}, i) = P(\texttt{start}, i) + \varepsilon_{\texttt{start},i}, \quad \sum_{i \in Q} \varepsilon_{\texttt{start},i} = 0.$$

Notice that $\tilde{P}$ is again a stochastic matrix. The LMC $(Q, \tilde{P}, L)$ admits an APB with precision $\varepsilon = 4\|E\|_\infty$. (The presence of the multiplicative factor is elucidated in the following set of inequalities.) In this example, the APB Γ_ε induces a collection of classes of ε-bisimilar pairs of states, which is equivalent to the partition of Q as described in Figure 2. This partition also corresponds to the collection of Γ_ε-closed sets. Let us consider the Γ_ε-closed set $A = \{\texttt{start}, \texttt{won}, \texttt{lost}\}$: since $\{5\} \equiv_\varepsilon \{9\}$, we obtain

$$\begin{aligned}
&|P(5, A) - P(9, A)| \\
&= \left| \left(0 + \frac{1}{9} + \varepsilon_{5,won} + \frac{1}{6} + \varepsilon_{5,lost}\right) \right. \\
&\quad \left. - \left(0 + \frac{1}{9} + \varepsilon_{9,won} + \frac{1}{6} + \varepsilon_{9,lost}\right) \right| \\
&\leq |\varepsilon_{5,won} - \varepsilon_{9,won}| + |\varepsilon_{5,lost} - \varepsilon_{9,lost}| \leq 2 \cdot 2\|E\|_\infty.
\end{aligned}$$

The precision bound $\varepsilon = 4\|E\|_\infty$ can be shown to hold regardless of the choice of ε-bisimilar pairs $q_1 \equiv_\varepsilon q_2$ and Γ_ε-closed set A.

Let us now introduce the quantity $\tilde{\varepsilon} := \frac{5}{36} - \frac{1}{9} = \frac{1}{9} - \frac{1}{12}$. If $\|E\|_\infty < \frac{\tilde{\varepsilon}}{2}$, then there exists a coarser APB that consists of 5 classes (instead of 6 as before) of ε-bisimilar states

$$\{\texttt{start}\}, \{4, 10, 5, 9\}, \{5, 9, 6, 8\}, \{\texttt{won}\}, \{\texttt{lost}\}.$$

The new APB is associated to a precision $\varepsilon = 4\tilde{\varepsilon}$, which can be checked as done above and imposing the bound on $\|E\|_\infty$. Note that the obtained APB with precision ε does not generate a partition of Q: this fact will be further discussed shortly. The collection of Γ_ε-closed sets is then

$$\{\texttt{start}\}, \{4, 10, 5, 9, 6, 8\}, \{\texttt{won}\}, \{\texttt{lost}\}.$$

As it is also evident in the above example, let us remark that an APB with precision ε does not usually induce a partition of the state space by equivalence classes that consist of ε-bisimilar states. This is due to the fact that APB with precision ε is not an equivalence relation, since in general it does not satisfy the transitive property. This entails that any two states belonging to the same Γ_ε-closed set are not necessarily ε-bisimilar: we show instead that they are $l\varepsilon$-bisimilar, with l a finite positive integer smaller than the diameter of the Γ_ε-closed set. To clarify and formalize this fact, we provide some connections between the concept of APB and graph theory [17].

DEFINITION 4. *[APB graph] Given a LMC $M = (Q, P, L)$ and an APB Γ_ε, we define the associated APB graph $\mathcal{G} = (V, E)$ as $V = Q$ and $E = \Gamma_\varepsilon$. Since Γ_ε is symmetric, then $\mathcal{G}$ is an undirected graph.*

The set $\{A_1, \ldots, A_m\}$ of all non-intersecting Γ_ε-closed sets forms a partition of the set Q. Moreover, the following straightforward result holds.

PROPOSITION 1. *Given a LMC $M = (Q, P, L)$ and an APB Γ_ε, then the associated APB graph $\mathcal{G}$ has m connected components $\{\mathcal{G}_i = (V_i, E_i)\}_{i=1}^m$, and $\forall i \in \{1, \ldots, m\}$, $V_i = A_i$.*

DEFINITION 5. *[Central vertex, radius, and diameter of APB graph] Given a LMC $M = (Q, P, L)$, an APB Γ_ε and the associated APB graph $\mathcal{G}$, we select for each Γ_ε-closed set A_i any element $\bar{a}_i \in A_i$ as a central vertex of $\mathcal{G}_i$, and define*

the radius r_i and the diameter d_i associated to A_i as the radius and diameter of $\mathcal{G}_i$.

The radius of a graph is the minimum eccentricity of any of its vertices. The eccentricity of a vertex is its greatest possible distance from any other vertex in the graph. Thus a radius of a graph (and its central vertex) can be computed by first running the all-pairs-shortest-path algorithm, then maximizing the computed distance for any vertex, and finally minimizing the obtained value over the vertices of the graph. This can be done over un-weighted, un-directed graphs in polynomial time, at worst in $\mathcal{O}(n^3)$ [4]. Analogous considerations hold for the concept of graph diameter.

Figure 3 provides an example of APB and associated APB graph. The dashed ellipses denote Γ_ε-bisimilar states in Q. The Γ_ε-closed sets are $A_1 = \{1, 2, 4, 5, 8\}$, $A_2 = \{3, 6\}$, $A_3 = \{7, 10\}$, $A_4 = \{9, 11, 12\}$. The central vertex a_1 of $\mathcal{G}_1$ can be indifferently state 4 or state 5. The radius and the diameter of $\mathcal{G}_1$ are respectively $r_1 = 2$ and $d_1 = 3$.

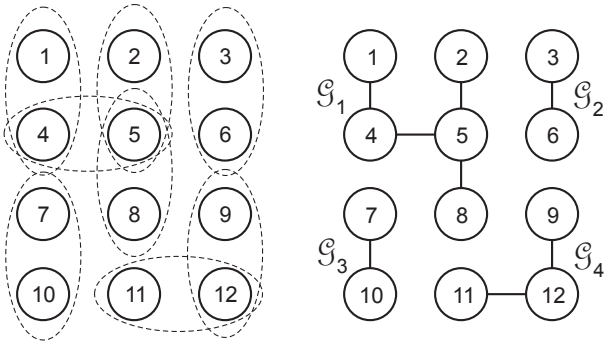

Figure 3: Left: APB - dashed ellipses define the Γ_ε relation. Right: associated APB graph $\mathcal{G} = \mathcal{G}_1 \cup \mathcal{G}_2 \cup \mathcal{G}_3 \cup \mathcal{G}_4$.

EXAMPLE 3. *With reference to the last instance of APB Γ_ε in Example 2, which is characterized by the following sets of Γ_ε bisimilar states*

$$\{\texttt{start}\}, \{4, 10, 5, 9\}, \{5, 9, 6, 8\}, \{\texttt{won}, \texttt{lost}\},$$

we obtain the following Γ_ε-closed sets:

$$A_1 = \{\texttt{start}\},\ A_2 = \{4, 10, 5, 9, 6, 8\},\ A_3 = \{\texttt{won}, \texttt{lost}\},$$

with radii $r_1 = 0, r_2 = 1, r_3 = 1$ and diameters $d_1 = 0, d_2 = 5, d_3 = 1$ The central vertices of the Γ_ε-closed sets are respectively `start` *for $\mathcal{G}_1$, either 5 or 9 for $\mathcal{G}_2$, and either* `won` *or* `lost` *for $\mathcal{G}_3$. The APB graph is represented in Figure 4.*

3. BISIMULATION BOUNDS ON REALIZATION DISTANCE

In this section we draw a connection between the notion of probabilistic realization distance over a LMC and that of APB. More precisely, we are interested in defining a metric on the distance in time between realizations of trajectories starting from two different initial conditions. We focus on initial conditions that are related by the notion of APB, and compute metrics over sets that are closed over this relation. We show explicit upper bounds in time for this distance.

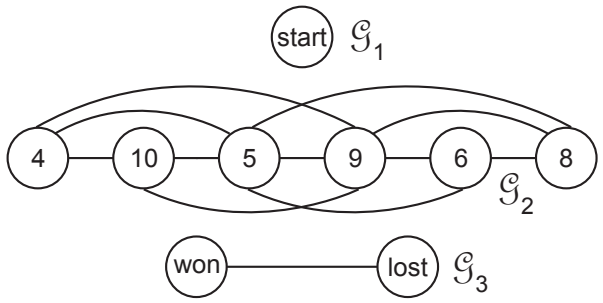

Figure 4: APB graph of Example 3.

DEFINITION 6. *[Probabilistic realization distance induced by Γ_ε] Given $M = (Q, P, L)$ and an APB $\Gamma_\varepsilon \subseteq Q \times Q$, we define a distance $d^k_{\Gamma_\varepsilon}(q_1, q_2)$ induced by Γ_ε at time $k \geq 1$ as follows:*

$$d^k_{\Gamma_\varepsilon}(q_1, q_2) = \max_{\forall A\ \Gamma_\varepsilon - closed} \left| P^k(q_1, \Gamma_\varepsilon(A)) - P^k(q_2, \Gamma_\varepsilon(A)) \right|.$$

We define $d^\infty_{\Gamma_\varepsilon}(q_1, q_2) = \lim_{k\to\infty} d^k_{\Gamma_\varepsilon}(q_1, q_2)$, if such limit exists.

Note that the probabilistic realization distance at time k induced by an APB Γ_ε between two LMC is the 1-norm distance at time k between the probability distributions, marginalized over the Γ_ε-closed sets.

Given $M = (Q, P, L)$ and an APB Γ_ε, let us introduce the quantity

$$d^k_\varepsilon(M) = \max_{(q_1,q_2)\in\Gamma_\varepsilon} d^k_{\Gamma_\varepsilon}(q_1, q_2). \tag{1}$$

3.1 Bound on finite-time realization distance

The next result draws an explicit bound on the finite-horizon realization distance. The bound depends on the existence of an APB Γ_ε, in particular on its precision ε and on the properties of the associated APB graph.

THEOREM 1. *Given a LMC $M = (Q, P, L)$ and an APB Γ_ε, then for any $(q_1, q_2) \in \Gamma_\varepsilon$ and for any $k \in \mathbb{N}$ the following holds:*

$$d^k_{\Gamma_\varepsilon}(q_1, q_2) \leq \varepsilon \left(m + \varepsilon \sum_{i=1}^{m} r_i \right)^{k-1}.$$

PROOF. Recall that $|Q| = n$ and that the m-dimensional $(m \leq n)$ collection $\{A_i\}_{i=1}^m$ of Γ_ε-closed sets results in a partition of Q. Each set A_i, $i \in \{1, \ldots, m\}$ is associated to a central vertex $\bar{a}_i$ and a radius r_i (cfr. Def. 5).
Consider any two states $q_1, q_2 \in Q : q_1 \equiv_\varepsilon q_2$. By Definition 3, for any Γ_ε-closed set A:

$$|P(q_1, A) - P(q_2, A)| \leq \varepsilon,$$

thus $d^1_{\Gamma_\varepsilon}(q_1, q_2) \leq \varepsilon$. As a next step, consider:

$$\left|P^2(q_1, A) - P^2(q_2, A)\right| = \left| \sum_{j=1}^{n} (P(q_1, q_j) - P(q_2, q_j)) P(q_j, A) \right|.$$

Notice that $\cup_{j=1}^n \{q_j\} = \cup_{i=1}^m A_i$. Consider a set $A_i \in \{A_i\}_{i=1}^m$, a generic state $q \in A_i$ and its central vertex $\bar{a}_i$, and the connected component $\mathcal{G}_i$ of the APB graph $\mathcal{G}$. By construction,

there exists a finite-discrete path

$$\{q = s_0, s_1, \ldots, s_k, s_{k+1}, \ldots, \bar{a}_i\}$$

of cardinality at most r_i, which connects q with $\bar{a}_i$. Therefore, $\forall q \in A_i$, it holds that

$$\begin{aligned} &P(q,A) \leq |P(q,A) - P(s_1,A)| \\ &+ |P(s_1,A) - \ldots - P(s_k,A)| + |P(s_k,A) - P(s_{k+1},A)| \\ &+ |P(s_{k+1},A) - \ldots - P(\bar{a}_i,A)| + P(\bar{a}_i,A). \end{aligned}$$

Since $(s_k, s_{k+1}) \in \Gamma_\varepsilon$ for all $0 \leq k \leq r_i - 1$ by definition of $\mathcal{G}_i$, then $\forall A, A\,\Gamma_\varepsilon$-closed, it holds that

$$\forall k, |P(s_k,A) - P(s_{k+1},A)| \leq \varepsilon.$$

We then obtain that

$$\forall q \in A_i, P(q,A) \leq P(\bar{a}_i,A) + r_i\varepsilon,$$

which leads to

$$\begin{aligned} &\left|P^2(q_1,A) - P^2(q_2,A)\right| \leq \\ &\left|\sum_{i=1}^{m} (P(q_1,A_i) - P(q_2,A_i))\,(P(\bar{a}_i,A) + r_i\varepsilon)\right| \leq \\ &\sum_{i=1}^{m} |P(q_1,A_i) - P(q_2,A_i)|\,(P(\bar{a}_i,A) + r_i\varepsilon) \leq \\ &\left|\sum_{i=1}^{m} \varepsilon\,(P(\bar{a}_i,A) + r_i\varepsilon)\right| \leq \varepsilon\left(m + \varepsilon\sum_{i=1}^{m} r_i\right). \end{aligned} \quad (2)$$

Inductively, we obtain

$$\left|P^k(q_1,A) - P^k(q_2,A)\right| \leq \varepsilon\left(m + \varepsilon\sum_{i=1}^{m} r_i\right)^{k-1},$$

which leads to the bound in the statement. □

Notice that the derived bound can be practically conservative, in particular due to the inequality in (2), and can thus be substituted by a tighter lower-approximant, as the following example displays.

Example 4. *Let us consider the first instance of APB discussed in Example 2. Note that the radii related to the Γ_ε-closed sets are equal either to 0 (for classes* `start`, `won`, `lost`*) or to 1 (for classes $(4,10),(5,9),(6,8)$).*

At step k, the derived bound results in the quantity $\varepsilon(6+3\varepsilon)^{k-1}$, however this worst-case bound can be actually refined. For instance, consider the two states 4 and 10, which are ε-bisimilar, and the Γ_ε-closed set $A = \{$`won`$\}$. Then, the inequality in (2) can be refined as follows (refer to the incoming edges into set $\{$`won`$\}$ in Figure 2):

$$\begin{aligned} &\left|P^k(4,A) - P^k(10,A)\right| \\ &\leq \varepsilon\left(\frac{2}{9} + 0\varepsilon + \frac{1}{12} + 1\varepsilon + \frac{1}{9} + 1\varepsilon + \frac{5}{36} + 1\varepsilon + 0 + 0\varepsilon\right)^{k-1} \\ &= \varepsilon\left(\frac{5}{9} + 3\varepsilon\right)^{k-1}. \end{aligned}$$

Notice that the new bound is not only lower than $\varepsilon(6+3\varepsilon)^{k-1}$, but is also decreasing in time if $\varepsilon < \frac{4}{27}$.

Theorem 1 provides a bound that in general increases quickly with k if the argument of the exponent is larger than the unity. In order to try to generalize this bound to hold over an infinite horizon (cfr. Section 3.3), we introduce next the formalism and theory of Markov set-Chains.

3.2 Markov Set-Chains

The results illustrated in this section are from [18] and are also summarized in [1].

Definition 7. *[18, Definition 2.5, Transition Set] Let $P, Q \in \mathbb{R}^{n\times n}$ be nonnegative matrices (not necessarily stochastic) with $P \leq Q$, where $\leq$ is element-wise. We define a transition set as $[\Pi] = [P,Q] = \{T \in \mathbb{R}^{n\times n} : T$ is a stochastic matrix and $P \leq T \leq Q\}$.*

In this paper, we assume that the set $[P,Q] \neq \emptyset$. A (discrete-time) Markov set-Chain can be formally introduced as follows.

Definition 8. *[18, Definition 2.5, Markov set-Chain] Let $[\Pi]$ be a transition set, i.e. a compact set of $n \times n$ stochastic matrices. Consider the set of all inhomogeneous Markov Chains having their transition matrices in $[\Pi]$. We call the sequence $[\Pi], [\Pi]^2, \ldots$ a Markov set-Chain (MSC), where $[\Pi]^k$ is defined by induction as the set of all possible products $T_1T_2\cdots T_k$, such that $\forall i = 1,\ldots,k, T_i \in [\Pi]$.*

Definition 9. *[18, Definition 1.2, Coefficient of Ergodicity of a Stochastic Matrix] For a stochastic matrix T, its coefficient of ergodicity is defined as follows:*

$$\mathcal{T}(T) = \frac{1}{2}\max_{i,j} ||a_i - a_j||,$$

where a_i, a_j are the i–th, j–th rows of T, and $\|\cdot\|$ is the standard 1-norm over row vectors: $\|x\| = \sum_k |x_k|$.

It can be shown that the condition $\mathcal{T}(T) < 1$, along with the condition of irreducibility of the chain, implies the existence of a unique limiting and invariant distribution for the associated Markov Chain [18, 13]. The previous definition can be directly extended to MSC.

Definition 10. *[18, Definition 3.1, Coefficient of Ergodicity of a Transition Set] For any transition set $[\Pi]$, its coefficient of ergodicity is defined over the stochastic matrices that define $[\Pi]$ as follows: $\mathcal{T}([\Pi]) = \max_{T\in[\Pi]} \mathcal{T}(T)$.*

Since $\mathcal{T}(\cdot)$ is a continuous function and $[\Pi]$ a compact set, the corresponding maximum argument exists.

Theorem 2. *[18, Theorem 3.1] Let $[\Pi]$ be the interval $[P,Q]$ and $T \in [\Pi]$, then $|\mathcal{T}([\Pi]) - \mathcal{T}(T)| \leq ||Q - P||$.*

The used matrix norm is the induced 1-norm over row-vectors: $\|T\| = \max_{x\neq 0} \frac{\|xT\|}{\|x\|}$. Let us define the diameter of a transition set:

$$\Delta([\Pi]) = \max_{T,T'\in[\Pi]} ||T - T'||.$$

The following result provides an upper bound for the diameter of the transition set $[\Pi]^k, k > 0$.

Theorem 3. *[18, Theorems 3.4, 3.11] Given a MSC with transition set $[\Pi] = [P,Q]$, then*

$$\Delta([\Pi]^k) \leq \mathcal{T}([\Pi])^k + (\mathcal{T}([\Pi])^{k-1} + \ldots + 1)\Delta([\Pi]).$$

In particular, if $\mathcal{T}([\Pi]) < 1$, given any initial distribution set $[\pi_0]$, there exists a unique set $[\pi_\infty]$ that is invariant, i.e. such that $[\pi_\infty][\Pi] = [\pi_\infty]$, and moreover such that

$$\lim_{k\to\infty}[\pi_k] = \lim_{k\to\infty}[\pi_0][\Pi]^k = [\pi_\infty].$$

Furthermore, the following holds:

$$\Delta([\pi_\infty]) \leq \frac{\Delta([\Pi])}{1-\mathcal{T}([\Pi])} \leq \frac{||Q-P||}{1-\mathcal{T}([\Pi])}.$$

The notion of limit of a vector interval hinges on the Hausdorff distance [18], which is a distance between sets. The derived bounds are not necessarily tight, however they are sufficient for the objectives of the study (finiteness of bounds), and as such they will be used below. Tighter results can be obtained with more sophistication (cfr. the Hi-Lo method and the notion of scrambling coefficient in [18]).

3.3 Bound on infinite-time realization distance

On the basis of the MSC theory illustrated in the previous section, we provide a new bound for the probabilistic realization distance induced by an APB Γ_ε.

THEOREM 4. *Given a LMC $M = (Q, P, L)$ and an APB Γ_ε, then for any $(q_1, q_2) \in \Gamma_\varepsilon$ and $k \in \mathbb{N}\cup\{\infty\}$ the following holds:*

$$d^k_{\Gamma_\varepsilon}(q_1, q_2) \leq \tau^k + \varepsilon\lambda m \sum_{l=0}^{k-1}\tau^l,$$

where $\tau = \mathcal{T}(M) + \varepsilon\lambda m$, and $\lambda = \max_{i=1,\ldots,m} d_i$ is the maximum diameter among the connected components of the APB graph.

PROOF. Consider the APB Γ_ε on M. We define a MSC over the state space $Q_{[M]} = \{A_1, \ldots, A_m\}$ of Γ_ε-closed sets, and characterized by the following transition set $P_{[M]}$: for each $A_1, A_2 \in Q_{[M]}$, define

$$P_{[M]}(A_1, A_2) = [\min_{q\in A_1}\{P(q, A_2)\},\ \max_{q\in A_1}\{P(q, A_2)\}].$$

Since the cardinality of the MSC $[M]$ is given by $m = |Q_{[M]}| \leq |Q| = n$, and since by Definition 3

$$\left|\min_{q\in A_1}\{P(q, A_2)\} - \max_{q\in A_1}\{P(q, A_2)\}\right| \leq \varepsilon d_1,$$

then $\Delta([M]) = \varepsilon\lambda m$: in fact (cfr. Section 3.2), $\Delta([M]) = \max_{i=1,\ldots,m}\sum_{j=1}^m \varepsilon d_i = \max_{i=1,\ldots,m}\varepsilon d_i m$. The coefficient of ergodicity $\mathcal{T}([M])$, according to Theorem 2, satisfies the inequality $\mathcal{T}(M) \leq \mathcal{T}([M]) \leq \mathcal{T}(M) + \varepsilon\lambda m = \tau$.

Given any $A_1, A_2 \in Q_{[M]}$ and any $q, q' \in A_1$, it follows that

$$P(q, A_2) \in P_{[M]}(A_1, A_2), \quad P(q', A_2) \in P_{[M]}(A_1, A_2).$$

Using the above set of inclusions, we can derive the following bound for any Γ_ε-closed set $A_2 \in Q_{[M]}$:

$$\left|P(q, A_2) - P(q', A_2)\right| \leq \Delta([M]),$$

Directly leveraging Theorem 3 this yields, $\forall k > 0$,

$$\left|P^k(q, A_2) - P^k(q', A_2)\right| \leq \Delta([M]^k) \leq \tau^k + \varepsilon\lambda m\sum_{l=0}^{k-1}\tau^l.$$

Thus, the following holds for all $(q_1, q_2) \in \Gamma_\varepsilon$:

$$d^k_{\Gamma_\varepsilon}(q_1, q_2) = \max_{\forall A\ \Gamma_\varepsilon - closed}\left|P^k(q_1, A) - P^k(q_2, A)\right| \leq \tau^k + \varepsilon\lambda m\sum_{l=0}^{k-1}\tau^l,$$

which proves the statement. □

The bound provided above is finite for any $k > 0$ if $\tau < 1$. Therefore, if the MSC $[M]$ is sufficiently ergodic (namely its coefficient of ergodicity is less then $1 - \varepsilon\lambda m$), then Theorem 4 also holds for $k \to \infty$ since the bound $d^k_\varepsilon(q_1, q_2)$ is finite in time for all $(q_1, q_2) \in \Gamma_\varepsilon$. In this instance, the result is stronger than the bound obtained in Theorem 1. Furthermore, a positive aspect of the above bound is that, under the ergodicity assumption, it does not accumulate in time, unlike other metrics [11, 28] that have to rely on discounting over time [12].

4. ROBUSTNESS OF PCTL FORMULAE

Robustness issue have been the main driver that has lead to an approximate concept of probabilistic bisimulation. In this section we go further along this research path by "robustifying" PCTL model checking.

4.1 Probabilistic Computation Tree Logic

Let $M = (Q, P, L)$ be a LMC, where $L = 2^{AP}$ and AP is a finite set of atomic propositions. We recall now syntax, semantics and model checking for PCTL.

DEFINITION 11. *[PCTL syntax, [16]] The syntax of PCTL is as follows:*

$$\Phi = \mathtt{true} \mid a \mid \neg\Phi \mid \Phi\wedge\Phi \mid \mathbb{P}_{\sim p}[\phi]$$
$$\phi = X\Phi \mid \Phi U^{\leq k}\Phi$$

where a is an atomic proposition, $\sim \in \{<, \leq, \geq, >\}$, $p \in [0, 1]$ and $k \in \mathbb{N}\cup\{\infty\}$.

PCTL formulae are interpreted over the states of a LMC. For the presentation of the syntax we distinguish between state formulae Φ and path formulae ϕ, which are evaluated over states and paths, respectively. A path $\mathbf{q}$ is an infinite sequence of states in the LMC $q_1\, q_2\, q_3 \ldots$, such that $P(q_i, q_{i+1}) > 0$. Let $\mathbf{q}(i)$ denote the i-th state in $\mathbf{q}$, i.e., $\mathbf{q}(i) = q_i$. For a state q and a PCTL formula Φ, we write $q \models \Phi$ to indicate that q satisfies Φ. Similarly, for a path $\mathbf{q}$ satisfying the path formula ϕ, we write $\mathbf{q} \models \phi$.

DEFINITION 12. *[PCTL semantics, [16]] The semantics of PCTL over LMC is defined as follows:*

$$\begin{array}{rcl} q \models \mathtt{true} & & \forall q \in Q \\ q \models a & \Leftrightarrow & a \in L(q) \\ q \models \neg\Phi & \Leftrightarrow & q \not\models \Phi \\ q \models \Phi\wedge\Psi & \Leftrightarrow & q \models \Phi \wedge q \models \Psi \\ q \models \mathbb{P}_{\sim p}[\phi] & \Leftrightarrow & Prob(\{\boldsymbol{q} : \boldsymbol{q}(0) = q : \boldsymbol{q} \models \phi\}) \sim p \\ \boldsymbol{q} \models X\Phi & \Leftrightarrow & \boldsymbol{q}(1) \models \Phi \\ \boldsymbol{q} \models \Phi U^{\leq k}\Psi & \Leftrightarrow & \exists i \in \mathbb{N} : (i \leq k\ \wedge\ \boldsymbol{q}(i) \models \Psi\ \wedge \\ & & \forall j < i, (\boldsymbol{q}(j) \models \Phi)). \end{array}$$

where $Prob(\{\mathbf{q} : \mathbf{q}(0) = q : \mathbf{q} \models \phi\})$ is the probability that q generates a path $\mathbf{q}$ satisfying formula ϕ. We now summarize

a model checking algorithm for PCTL over LMC [6, 7, 16]. The inputs to the algorithm are a LMC $M = (Q, P, L)$ and a PCTL formula Φ. The output is the set of states $Sat(\Phi) = \{q \in Q : q \models \Phi\}$, i.e. the set containing all the states of the model which satisfy Φ. The overall structure of the algorithm is identical to the model checking algorithm for CTL [5] (the non-probabilistic temporal logic which PCTL is based on) and can be summarized as follows:

$$\begin{aligned} Sat(\texttt{true}) &= Q \\ Sat(a) &= \{q \in Q : a \in L(q)\} \\ Sat(\neg\Phi) &= Q \setminus Sat(\Phi) \\ Sat(\Phi \wedge \Psi) &= Sat(\Phi) \cap Sat(\Psi) \\ Sat(\mathbb{P}_{\sim p}[\phi]) &= \{q \in Q : Prob^M(q, \phi) \sim p\}, \end{aligned}$$

where $Prob^M(q, \phi)$ is the probability that q generates a path satisfying formula ϕ. Model checking for the majority of these formulae is trivial to implement and is, in fact, the same as for the non-probabilistic logic CTL. The exception is made up of formulae with the form $\mathbb{P}_{\sim p}[\phi]$. For these formulae, we have to calculate for all states q of the LMC, the probability $Prob^M(q, \phi)$, then compare these values to the bound $\sim p$ in the formula. We now describe how to compute these values for the two cases: $\mathbb{P}_{\sim p}[X\Phi]$ and $\mathbb{P}_{\sim p}[\Phi U^{\leq k}\Psi]$. Because of the recursive nature of the PCTL model checking algorithm, we can assume that the relevant sets $Sat(\Phi)$ and $Sat(\Psi)$ are already known.

$\mathbb{P}_{\sim p}[X\Phi]$ **formulae.** In this case, we need to compute the probability $Prob^M(q, X\Phi)$ for each $q \in Q$. This requires the probabilities of the immediate transitions from q:

$$Prob^M(q, X\Phi) = P(q, Sat(\Phi)). \tag{3}$$

$\mathbb{P}_{\sim p}[\Phi U^{\leq k}\Psi], k \in \mathbb{N} \cup \{\infty\}$ **formulae.** For such formulae we need to determine the probabilities $Prob^M(q, \Phi U^{\leq k}\Psi)$ for all states q, where $k \in \mathbb{N} \cup \{\infty\}$. To this aim, we need the following definition:

Definition 13. *[Formula-dependent LMC, [16]] Given a LMC $M = (Q, P, L)$ and a PCTL formula Φ, let $M[\Phi] = (Q, P[\Phi], L)$, where if $q \not\models \Phi$, then $P[\Phi](q, q') = P(q, q')$ for all $q' \in Q$, and if $q \models \Phi$, then $P[\Phi](q, q) = 1$ and $P[\Phi](q, q') = 0$ for all $q' \neq q$.*

Using this transformation we characterize $Prob^M(q, \Phi U^{\leq k}\Psi)$ as follows.

Proposition 2. *[16] For any LMC $M = (Q, P, L)$, state $q \in Q$, PCTL formulae Φ and Ψ, and $k \in \mathbb{N}$:*

$$Prob^M(q, \Phi U^{\leq k}\Psi) = P^k_{M[\neg\Phi\vee\Psi]}(q, Sat(\Psi)). \tag{4}$$

Model checking can be performed directly on $M[\neg\Phi \vee \Psi]$. For the unbounded until case, we obtain:

$$Prob^M(q, \Phi U^{\leq \infty}\Psi) = \lim_{k\to\infty} P^k_{M[\neg\Phi\vee\Psi]}(q, Sat(\Psi)).$$

We remark that $\lim_{k\to\infty} P^k_{M[\neg\Phi\vee\Psi]}(q, Sat(\Psi))$ always exists and is unique. This can be shown using the theory of labeled Markov Chains, and is also implied by the existence and uniqueness of a solution to unbounded until formulae [16, 22].

4.2 Robust PCTL

Given a PCTL formula Φ, a LMC M and a APB Γ_ε over M, we propose in the following definition an iterative recipe to construct a strengthened formula $S_\varepsilon(\Phi)$ and a relaxed formula $R_\varepsilon(\Phi)$. This definition will be used in the following to define the set of formulae that are preserved by an APB with precision ε.

Definition 14. *[Strengthened and Relaxed PCTL formulae] Given a LMC M, a PCTL formula Φ and a non-negative real $\varepsilon > 0$, we define the ε–*strengthened *PCTL formula $S_\varepsilon(\Phi)$ and the ε–*relaxed *PCTL formula $R_\varepsilon(\Phi)$ by structural induction as follows:*

1. $S_\varepsilon(\texttt{true}) = \texttt{true}$, $R_\varepsilon(\texttt{true}) = \texttt{true}$.
2. $S_\varepsilon(a) = a$, $R_\varepsilon(a) = a$.
3. $S_\varepsilon(\neg\Phi) = \neg R_\varepsilon(\Phi)$, $R_\varepsilon(\neg\Phi) = \neg S_\varepsilon(\Phi)$.
4. $S_\varepsilon(\Phi \wedge \Psi) = S_\varepsilon(\Phi) \wedge S_\varepsilon(\Psi)$, $R_\varepsilon(\Phi \wedge \Psi) = R_\varepsilon(\Phi) \wedge R_\varepsilon(\Psi)$.
5. $S_\varepsilon(\mathbb{P}_{\sim p}[X\Phi]) = \mathbb{P}_{\sim p'}[XS_\varepsilon(\Phi)]$, *where:*
$$p' = \begin{cases} p - \varepsilon & \text{if } \sim \in \{<, \leq\} \\ p + \varepsilon & \text{if } \sim \in \{>, \geq\}, \end{cases}$$
$R_\varepsilon(\mathbb{P}_{\sim p}[X\Phi]) = \mathbb{P}_{\sim p'}[XR_\varepsilon(\Phi)]$, *where:*
$$p' = \begin{cases} p + \varepsilon & \text{if } \sim \in \{<, \leq\} \\ p - \varepsilon & \text{if } \sim \in \{>, \geq\}. \end{cases}$$
6. $S_\varepsilon(\mathbb{P}_{\sim p}[\Phi U^{\leq k}\Psi]) = \mathbb{P}_{\sim p'}[S_\varepsilon(\Phi)U^{\leq k}S_\varepsilon(\Psi)]$, *where $k \in \mathbb{N} \cup \{\infty\}$ and:*
$$p' = \begin{cases} p - d^k_\varepsilon(M) & \text{if } \sim \in \{<, \leq\} \\ p + d^k_\varepsilon(M) & \text{if } \sim \in \{>, \geq\}, \end{cases}$$
$R_\varepsilon(\mathbb{P}_{\sim p}[\Phi U^{\leq k}\Psi]) = \mathbb{P}_{\sim p'}[R_\varepsilon(\Phi)U^{\leq k}R_\varepsilon(\Psi)]$, *where $k \in \mathbb{N} \cup \{\infty\}$ and:*
$$p' = \begin{cases} p + d^k_\varepsilon(M) & \text{if } \sim \in \{<, \leq\} \\ p - d^k_\varepsilon(M) & \text{if } \sim \in \{>, \geq\}. \end{cases}$$

We say that a strengthened or relaxed formula $S_\varepsilon(\Phi)$ or $R_\varepsilon(\Phi)$ is *consistent* if in each step of the recursive substitution above $p' \in [0, 1]$. Note that *inconsistent* strengthened or relaxed formulae are either identically true (e.g. $\mathbb{P}_{\leq 1.1}[\cdot]$, $\mathbb{P}_{\geq -0.1}[\cdot]$) or identically false (e.g. $\mathbb{P}_{\geq 1.1}[\cdot]$, $\mathbb{P}_{\leq -0.1}[\cdot]$).

Example 5. *Consider the first perturbed LMC $(Q, \tilde{P}, L)$ introduced in Example 2 and the formula*

$$\Phi = \mathbb{P}_{\geq\gamma}[\texttt{start}\, U^{\leq k}\, \texttt{won}].$$

Let us first compute the distance $d^k_\varepsilon(Q, \tilde{P}, L) := d^k_\varepsilon$ in (1), employing the finite-time bounds derived in Theorem 1, which we tailor to the case study at hand as discussed in Example 4. A few cases need to be considered. Select $A = \{\texttt{won}\}$. Notice that if $(i, j) \in \{(4, 10), (5, 9), (6, 8)\}$, then

$$\left|P^k(\{i\}, A) - P^k(\{j\}, A)\right| \leq \varepsilon \left(\frac{5}{9} + 3\varepsilon\right)^{k-1}.$$

If $A = \{\texttt{lost}\}$*, then*

$$\left|P^k(\{i\}, A) - P^k(\{j\}, A)\right| \leq \varepsilon \left(\frac{11}{18} + 3\varepsilon\right)^{k-1}.$$

If $A = \{(4, 10)\}$ *or* $A = \{(5, 9)\}$ *or* $A = \{(6, 8)\}$*, then*

$$\left|P^k(\{i\}, A) - P^k(\{j\}, A)\right| \leq \varepsilon \left(\frac{5}{6} + 3\varepsilon\right)^{k-1}.$$

The instance $A = \{\texttt{start}\}$ *yields a trivial bound. In conclusion, we derive that* $d_\varepsilon^k = \varepsilon \left(\frac{5}{6} + 3\varepsilon\right)^{k-1}$*. It follows that*

$$S_\varepsilon(\Phi) = \mathbb{P}_{\geq \gamma + d_\varepsilon^k}[\texttt{start}\, U^{\leq k}\, \texttt{won}],$$

whereas

$$R_\varepsilon(\Phi) = \mathbb{P}_{\geq \gamma - d_\varepsilon^k}[\texttt{start}\, U^{\leq k}\, \texttt{won}].$$

Notice that assuming $\varepsilon \leq 1/18$*, the distance bound is decreasing with time* k*. Furthermore,* $S_0(\Phi) = R_0(\Phi) = \Phi$*.*

Definition 15. *[*ε*-robustness of a PCTL formula] Given a LMC* M*, a PCTL formula* Φ *and a non-negative real* $\varepsilon > 0$*, we say that* Φ *is* ε*-robust with respect to* M *if, for any* $q \in Q$ *and for any sub-formula* Ψ *of* Φ*, either* $q \in Sat(S_\varepsilon(\Psi))$ *or* $q \notin Sat(R_\varepsilon(\Psi))$*.*

The above definition requires that the case $q \models R_\varepsilon(\Psi) \wedge \neg S_\varepsilon(\Psi)$ can not occur. The following Theorem establishes the main result of this paper.

Theorem 5. *Given a LMC* M*, an APB* Γ_ε *and a PCTL formula* Φ*, let* Φ *be* ε*-robust with respect to* M*. Then for each* $q_1, q_2 \in Q$ *such that* $q_1 \equiv_\varepsilon q_2$ *the following holds:*

$$q_1 \in Sat(\Phi) \Leftrightarrow q_2 \in Sat(\Phi).$$

Proof. (By induction on formula depth).
According to [6, 7, 16], model checking is performed over the parse tree of Φ where the root node is labeled with Φ itself, and leaves of the tree are labeled with either `true` or an atomic proposition a.

We prove the theorem by structural induction on Φ. For the base case, $\Phi = a$ for an atomic proposition a. Since $q_1 \equiv_\varepsilon q_2$, it follows that $L(q_1) = L(q_2)$, hence $q_1 \in Sat(a)$ iff $q_2 \in Sat(a)$. Then, we prove the induction step for the negation, next and until (bounded and unbounded) operators. The induction step for the other formulae is trivial.

$\neg\Phi$ formulae (negation):
Let $q_1 \in Sat(\neg\Phi)$, then $q_1 \notin Sat(\Phi)$. The induction hypothesis implies that $\forall q_1, q_2$ such that $q_1 \equiv_\varepsilon q_2$ then

$$q_1 \in Sat(\Phi) \Leftrightarrow q_2 \in Sat(\Phi),$$

which implies that

$$q_1 \notin Sat(\Phi) \Leftrightarrow q_2 \notin Sat(\Phi).$$

The above statement directly implies that

$$q_1 \in Sat(\neg\Phi) \Leftrightarrow q_2 \in Sat(\neg\Phi).$$

$\mathbb{P}_{\sim p}[X\Phi]$ formulae (next operator):
Let $q_1 \in Sat(\mathbb{P}_{\sim p}[X\Phi])$, then equation (3) implies that:

$$P(q_1, Sat(\Phi)) \sim p.$$

The induction hypothesis implies that $\forall q_1', q_2'$ such that $q_1' \equiv_\varepsilon q_2'$ the following holds:

$$q_1' \in Sat(\Phi) \Leftrightarrow q_2' \in Sat(\Phi). \tag{5}$$

Therefore $Sat(\Phi)$ is a Γ_ε-closed set. Given any $q_2 \in Q$ such that $q_1 \equiv_\varepsilon q_2$, Definition 3 implies that:

$$|P(q_1, Sat(\Phi)) - P(q_2, Sat(\Phi))| \leq \varepsilon. \tag{6}$$

It follows that, if $\sim \in \{<, \leq\}$, then:

$$P(q_2, Sat(\Phi)) \sim P(q_1, Sat(\Phi)) + \varepsilon.$$

Since $q_1 \in Sat(\mathbb{P}_{\sim p}[X\Phi])$, the robustness assumption implies that $q_1 \in Sat(S_\varepsilon(\mathbb{P}_{\sim p}[X\Phi]))$, thus:

$$P(q_1, Sat(\Phi)) + \varepsilon \sim p - \varepsilon + \varepsilon = p.$$

Analogously, if $\sim \in \{>, \geq\}$, then:

$$P(q_2, Sat(\Phi)) \sim p + \varepsilon - \varepsilon = p.$$

The above reasoning implies that

$$q_1 \in Sat(\mathbb{P}_{\sim p}[X\Phi]) \Rightarrow q_2 \in Sat(\mathbb{P}_{\sim p}[X\Phi]).$$

To complete the induction step, we need to prove that

$$q_1 \notin Sat(\mathbb{P}_{\sim p}[X\Phi]) \Rightarrow q_2 \notin Sat(\mathbb{P}_{\sim p}[X\Phi]).$$

The proof is analogous and is thus omitted.

$\mathbb{P}_{\sim p}[\Phi U^{\leq k}\Psi], k \in \mathbb{N} \cup \{\infty\}$ formulae (bounded and unbounded until operators):
Let $q \in Sat(\mathbb{P}_{\sim p}[\Phi U^{\leq k}\Psi])$, $k \in \mathbb{N} \cup \{\infty\}$, then equation (4) implies that:

$$P^k_{M[\neg\Phi \vee \Psi]}(q, Sat(\Psi)) \sim p.$$

The induction hypothesis implies that $\forall q_1', q_2'$ such that $q_1' \equiv_\varepsilon q_2'$ the following hold:

$$q_1' \in Sat(\neg\Phi) \Leftrightarrow q_2' \in Sat(\neg\Phi),$$
$$q_1' \in Sat(\Psi) \Leftrightarrow q_2' \in Sat(\Psi).$$

Therefore $Sat(\neg\Phi)$ and $Sat(\Psi)$ are Γ_ε-closed sets, and the following holds:

$$\begin{aligned}\forall (q_1, q_2) \in \Gamma_\varepsilon,\ q_1' \in Sat(\neg\Phi \vee \Psi) \\ \Leftrightarrow q_1' \in Sat(\neg\Phi) \vee q_1' \in Sat(\Psi) \\ \Leftrightarrow q_2' \in Sat(\neg\Phi) \vee q_2' \in Sat(\Psi) \\ \Leftrightarrow q_2' \in Sat(\neg\Phi \vee \Psi).\end{aligned}$$

Therefore $Sat(\neg\Phi \vee \Psi)$ is a Γ_ε-closed set, and the following property holds for the LMC $M[\neg\Phi \vee \Psi]$:

$$\begin{aligned}\forall (q_1, q_2) \in \Gamma_\varepsilon,\ \text{either } q_1, q_2 \in Sat(\neg\Phi \vee \Psi) \\ \text{or } q_1, q_2 \notin Sat(\neg\Phi \vee \Psi).\end{aligned}$$

We prove now that Γ_ε is an APB over the LMC $M[\neg\Phi \vee \Psi]$. Pick any $(q_1, q_2) \in \Gamma_\varepsilon$. If $q_1, q_2 \in Sat(\neg\Phi \vee \Psi)$, then by Definition 13:

$$\begin{aligned}&P[\neg\Phi \vee \Psi](q_1, q_1) = P[\neg\Phi \vee \Psi](q_2, q_2) = 1,\\ &\forall q' \neq q_1,\ P[\neg\Phi \vee \Psi](q_1, q') = 0,\\ &\forall q' \neq q_2,\ P[\neg\Phi \vee \Psi](q_2, q') = 0.\end{aligned}$$

Therefore, for any Γ_ε-closed set $A \neq Sat(\neg\Phi \vee \Psi)$,

$$|P[\neg\Phi \vee \Psi](q_1, A) - P[\neg\Phi \vee \Psi](q_2, A)| = 0 - 0 \leq \varepsilon,$$

and

$$|P[\neg\Phi \vee \Psi](q_1, Sat(\neg\Phi \vee \Psi)) - P[\neg\Phi \vee \Psi](q_2, Sat(\neg\Phi \vee \Psi))| = 1 - 1 \leq \varepsilon.$$

If $q_1, q_2 \notin Sat(\neg\Phi \vee \Psi)$, then by Definition 13:

$$\forall q' \in Q,\ P[\neg\Phi \vee \Psi](q_1, q') = P(q_1, q') \text{ and}$$
$$P[\neg\Phi \vee \Psi](q_2, q') = P(q_2, q').$$

Therefore, for any Γ_ε-closed set A

$$|P[\neg\Phi \vee \Psi](q_1, A) - P[\neg\Phi \vee \Psi](q_2, A)|$$
$$= |P(q_1, A) - P(q_2, A)| \leq \varepsilon.$$

Since the condition of Definition 3 is satisfied, then Γ_ε is an APB over the LMC $M[\neg\Phi \vee \Psi]$. Therefore, Theorem 4 implies that for any Γ_ε-closed set A the following holds, for each $k \in \mathbb{N} \cup \{\infty\}$:

$$\left| P^k_{M[\neg\Phi\vee\Psi]}(q_1, A) - P^k_{M[\neg\Phi\vee\Psi]}(q_2, A) \right| \leq d^k_\varepsilon(M[\neg\Phi \vee \Psi]), \tag{7}$$

where $P^k_{M[\neg\Phi\vee\Psi]}(q, A)$ is the probability that the set A is reached in k steps by an execution of $M[\neg\Phi \vee \Psi]$ starting from state q. If $\sim \in \{<, \leq\}$ then using the robustness assumption and by applying Equation (7) with $A = Sat(\Psi)$, it follows that:

$$P^k_{M[\neg\Phi\vee\Psi]}(q_2, Sat(\Psi)) \sim P^k_{M[\neg\Phi\vee\Psi]}(q_1, Sat(\Psi)) + d^k_\varepsilon(M)$$
$$\sim p - d^k_\varepsilon(M) + d^k_\varepsilon(M) = p.$$

Analogously, if $\sim \in \{>, \geq\}$, then:

$$P^k_{M[\neg\Phi\vee\Psi]}(q_2, Sat(\Psi)) \sim p + d^k_\varepsilon(M) - d^k_\varepsilon(M) = p.$$

The above reasoning implies that

$$q_1 \in Sat(\mathbb{P}_{\sim p}[\Phi U^{\leq k} \Psi]) \Rightarrow q_2 \in Sat(\mathbb{P}_{\sim p}[\Phi U^{\leq k} \Psi]).$$

To complete the induction step, we need to prove that

$$q_1 \notin Sat(\mathbb{P}_{\sim p}[\Phi U^{\leq k} \Psi]) \Rightarrow q_2 \notin Sat(\mathbb{P}_{\sim p}[\Phi U^{\leq k} \Psi]).$$

The proof is analogous and is thus omitted.

The above steps imply that $q_1 \in Sat(\Phi) \Leftrightarrow q_2 \in Sat(\Phi)$, and this completes the proof. □

The above result allows the verification of a PCTL formula Φ over a numerical model M_ε obtained from a concrete model M, where the transition probabilities of M_ε are obtained by approximation with precision ε. If Φ is ε robust over M_ε, then the PCTL model checking results over M_ε can be exported over M. Otherwise, it is necessary to use a more refined numerical model. Characterizing the existence of a strictly positive precision $\varepsilon' > 0$ that allows PCTL model checking of Φ over $M_{\varepsilon'}$ is an interesting question for future work.

EXAMPLE 6. *Consider the LMC $(Q, \tilde{P}, L)$ discussed in Example 5, which is a perturbed version of (Q, P, L) in Example 1. As suggested in [10], the definitions of (probabilistic bisimulation and of) APB can be extended to relate two different LMC: this is achieved by considering an APB over the cross product of the two LMC. We are here interested in verifying properties of (Q, P, L) over $(Q, \tilde{P}, L)$. Suppose that $\tilde{P}$ is obtained by quantization of P. This example argues that an increase in the quantization precision allows verifying on $(Q, \tilde{P}, L)$ a larger set of PCTL properties for (Q, P, L).*

Assume that the quantization is obtained by truncating the elements of P within its third decimal digit: for instance $5/36 = 0.13\bar{8}$ is approximated by 0.139, whereas $1/9 = 0.\bar{1}$ is approximated by 0.111. It is easy to realize that $\tilde{P}$ can be related to an error bound $\|E\|_\infty = 5 \cdot 10^{-4}$ and thus to an APB with approximation precision $\varepsilon = 2 \cdot 10^{-3}$. Similarly, if $\tilde{P}$ is obtained by truncating the elements of P within its i^{th} decimal digit, then $\varepsilon = 2 \cdot 10^{-i}$. Recall that in Example 5 we have established that $d^k_\varepsilon = \varepsilon\left(\frac{5}{6} + 3\varepsilon\right)^{k-1}$. The bounded until formula

$$\Phi = \mathbb{P}_{\geq\gamma}[\mathtt{start}\, U^{\leq k}\, \mathtt{won}]$$

is robust if $d^k_\varepsilon \leq \gamma \leq 1 - d^k_\varepsilon$: this can be established by application of Definition 15 over any $q \in Q$. Figure 5 plots this upper bound on γ, over different approximation digits and time horizons. Select $\gamma = 0.7$, and assume that the approximation digit is strictly greater than one. If Φ is true on $(Q, \tilde{P}, L)$ for any $k > 1$, being Φ robust then Φ is also true on (Q, P, L) for any $k > 1$. If instead the approximation digit is one then regardless of the value of Φ on $(Q, \tilde{P}, L)$ for any $k > 1$, since Φ not robust, we cannot draw any conclusion on the validity of Φ over (Q, P, L).

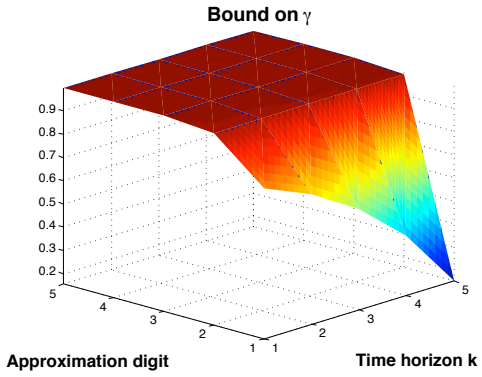

Figure 5: Robustness bounds on γ for formula $\Phi = \mathbb{P}_{\geq\gamma}[\mathtt{start}\, U^{\leq k}\, \mathtt{won}]$.

Consider now the formula $\Phi = \mathbb{P}_{\leq\gamma}[X\, \mathtt{lost}]$. Then $S_\varepsilon(\Phi) = \mathbb{P}_{\leq\gamma-\varepsilon}[X\, \mathtt{lost}]$ and $R_\varepsilon(\Phi) = \mathbb{P}_{\leq\gamma+\varepsilon}[X\, \mathtt{lost}]$. Robustness can again be studied via Definition 15, applied on any $q \in Q$. The allowed interval for ε is $[0, 5/12]$. In particular, if $0 \leq \varepsilon < 1/36$, then

$$\gamma \in [\varepsilon, 1/9] \cup [1/9 + \varepsilon, 1/6 - \varepsilon] \cup [1/6 + \varepsilon, 1 - \varepsilon],$$

whereas if $1/36 \leq \varepsilon < 1/18$, then

$$\gamma \in [\varepsilon, 1/9 - \varepsilon] \cup [1/6 + \varepsilon, 1 - \varepsilon],$$

and if $1/18 \leq \varepsilon \leq 5/12$, then

$$\gamma \in [1/6 + \varepsilon, 1 - \varepsilon].$$

With focus on the approximated LMC $(Q, \tilde{P}, L)$, again as the approximation increases to i digits, Φ grows in robustness. Given a particular approximation digit, if Φ is robust then we can safely claim that checking $j \models \Phi$ and $k \models \Phi$ is equivalent for any $j, k \in \{4, 10, 5, 9, 6, 8\}$.

5. CONCLUSIONS

With focus on labeled Markov Chains (LMC) in discrete-time, this work has utilized the notion of approximate probabilistic bisimulation (APB) to introduce bounds on probabilistic realization metrics in time between approximately

bisimilar states. As the main contribution of the paper, we have shown that the presence of an APB implies the preservation of (properly defined) robust PCTL formulae. This result allows the verification of a PCTL formula executed over an abstract model obtained as an approximation of a concrete model, e.g. over a numerical model where the transition probabilities are obtained by finite-precision approximation.

As for future work, it is of interest to understand what model properties (ergodicity, presence of absorbing classes) yield APB resulting in finite bounds for the probabilistic realization metrics. Furthermore, as discussed in Section 1, we aim to leverage on the abstraction procedure we developed in [1] and on the results developed in this paper to verify properties of continuous models as general as Stochastic Hybrid Systems by using classical PCTL model checking algorithms over a finite LMC abstraction which satisfies some accuracy constraints.

6. REFERENCES

[1] A. Abate, A. D'Innocenzo, and M.D. Di Benedetto. Approximate abstractions of stochastic hybrid systems. *IEEE Transactions on Automatic Control*, 56(11):2688–2694, November 2011.

[2] A. Aziz, K. Sanwal, V. Singhal, and R. Brayton. Model-checking continuous time Markov chains. *ACM Transactions on Computational Logic*, 1(1):162–170, 2000.

[3] C. Baier and J.-P. Katoen. *Principles of Model Checking*. The MIT Press, Cambridge, MA, 2008.

[4] T.M. Chan. All-pairs shortest paths for unweighted undirected graphs in o(mn) time. In *Proceedings of the seventeenth annual ACM-SIAM Symposium on Discrete Algorithm (SODA '06)*, pages 514–523, 2006.

[5] E. Clarke, E. Emerson, and A. Sistla. Automatic verification of finite-state concurrent systems using temporal logics. *ACM Transactions on Programming Languages and Systems*, 8(2):244–263, 1968.

[6] C. Courcoubetis and M. Yannakakis. Verifying temporal properties of finite state probabilistic programs. In *Proceedings of the* 29^{th} *Annual Symposium on Foundations of Computer Science (FOCS '88)*, pages 338–345, December 1988.

[7] C. Courcoubetis and M. Yannakakis. The complexity of probabilistic verification. *Journal of the ACM*, 42(4):857–907, 1995.

[8] L. de Alfaro, M. Faella, and M. Stoelinga. Linear and branching system metrics. *IEEE Transactions on Software Engineering*, 35(2):258–273, 2009.

[9] S. Derisavi, H. Hermanns, and W.H. Sanders. Optimal state-space lumping in Markov chains. *Information Processing Letters*, 87(6):309–315, September 2003.

[10] J. Desharnais, A. Edalat, and P. Panangaden. Bisimulation for labeled Markov processes. *Information and Computation*, 179(2):163–193, December 2002.

[11] J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden. Metrics for labeled Markov processes. *Theoretical Computer Science*, 318(3):323–354, 2004.

[12] J. Desharnais, F. Laviolette, and M. Tracol. Approximate analysis of probabilistic processes: logic, simulation and games. In *Proceedings of the International Conference on Quantitative Evaluation of SysTems (QEST '08)*, pages 264–273, Sept. 2008.

[13] R. Durrett. *Probability: Theory and Examples - Third Edition*. Duxbury Press, 2004.

[14] A. Giacalone, C.-C. Jou, and S.A. Smolka. Algebraic reasoning for probabilistic concurrent systems. In *Proceedings of the IFIP TC2 Working Conference on Programming Concepts and Methods*, pages 443–458, 1990.

[15] A. Girard and G.J. Pappas. Approximation metrics for discrete and continuous systems. *IEEE Transactions on Automatic Control*, 52(5):782–798, 2007.

[16] H. Hansson and B. Jonsson. A logic for reasoning about time and reliability. *Formal Aspects of Computing*, 6(5):512–535, 1994.

[17] F. Harary. *Graph Theory*. Addison-Wesley, 1994.

[18] H.J. Hartfiel. *Markov set-Chains*, volume 1695 of *Lecture Notes in Mathematics*. Springer-Verlag Berlin Heidelberg, 1998.

[19] T. A. Henzinger, R. Majumdar, and V. Prabhu. Quantifying similarities between timed systems. In *Proceedings of the Third International Conference on Formal Modeling and Analysis of Timed Systems (FORMATS '05)*, volume 3829 of *Lecture Notes in Computer Science*, pages 226–241. Springer, 2005.

[20] J.-P. Katoen, E.M. Hahn, H. Hermanns, D.N. Jansen, and I. Zapreev. The ins and outs of the probabilistic model checker MRMC. In *Proceedings of the International Conference on Quantitative Evaluation of SysTems (QEST '09)*, pages 167–176, 2009.

[21] J.-P. Katoen, T. Kemna, I.S. Zapreev, and D.N. Jansen. Bisimulation minimisation mostly speeds up probabilistic model checking. In Orna Grumberg and Michael Huth, editors, *Proceedings of Tools and Algorithms for the Construction and Analysis of Systems (TACAS '07)*, volume 4424 of *LNCS*, pages 87–101. Springer, 2007.

[22] M. Kwiatkowska, G. Norman, and D. Parker. Stochastic model checking. In M. Bernardo and J. Hillston, editors, *Formal Methods for the Design of Computer, Communication and Software Systems: Performance Evaluation (SFM '07)*, volume 4486 of *Lecture Notes in Computer Science*, pages 220–270. Springer Verlag, 2007.

[23] M. Kwiatkowska, G. Norman, and D. Parker. PRISM: Probabilistic model checking for performance and reliability analysis. *ACM SIGMETRICS Performance Evaluation Review*, 2009.

[24] M. Kwiatkowska, G. Norman, and D. Parker. PRISM 4.0: Verification of probabilistic real-time systems. In *Proc. 23rd International Conference on Computer Aided Verification (CAV'11)*, volume 6806 of *LNCS*, pages 585–591. Springer, 2011.

[25] K.G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Information and Computation*, 94:1–28, 1991.

[26] R. Milner. *A Calculus of Communicating Systems*. Springer-Verlag New York, Inc., Secaucus, NJ, 1982.

[27] D. Park. Concurrency and automata on infinite sequences. In *Proceedings of the 5th GI-Conference on Theoretical Computer Science*, pages 167–183, 1981.

[28] F. van Breugel and J. Worrell. Approximating and computing behavioural distances in probabilistic transition systems. *Theoretical Computer Science*, 360(1-3):373–385, 2006.

[29] R. Wimmer and B. Becker. Correctness issues of symbolic bisimulation computation for Markov chains. In *Proceedings of MMB-DFT, LNCS vol. 5987*, pages 287–301, 2010.