

On the Metric-based Approximate Minimization of Markov Chains^{☆☆☆}

Giovanni Bacci^{a,*}, Giorgio Bacci^a, Kim G. Larsen^a, Radu Mardare^a

^a*Department of Computer Science, Aalborg University, Denmark*

Abstract

In this paper we address the approximate minimization problem of Markov Chains (MCs) from a behavioral metric-based perspective. Specifically, given a finite MC and a positive integer k , we are looking for an MC with at most k states having *minimal* distance to the original. The metric considered in this work is the bisimilarity distance of Desharnais et al.. For this metric we show that (1) optimal approximations always exist; (2) the problem has a bilinear program characterization; and (3) prove that its threshold problem is in PSPACE and NP-hard.

In addition to the bilinear program solution, we present an approach inspired by expectation maximization techniques for computing suboptimal solutions to the problem. Experiments suggest that our method gives a practical approach that outperforms the bilinear program implementation run on state-of-the-art bilinear solvers.

Keywords: Behavioral Distances, Probabilistic Models, Automata Minimization

1. Introduction

Minimization of finite automata, i.e., the process of transforming a given finite automaton into an equivalent one with minimum number of states, has been a major subject since the 1950s due to its fundamental importance for any implementation of finite automata tools.

The first algorithm for the minimization of deterministic finite automata (DFAs) is due to Moore [2], with time complexity $O(n^2s)$, later improved by the now classical Hopcroft's algorithm [3] to $O(ns \log n)$, where n is the number of states and s the size

[☆]This article is an extended version of a paper accepted for publication at ICALP 2017 [1]. The current version provides proofs omitted in the original paper, additional examples, and a revised section on experimental results. Remarkably, we revise the proof of [1, Theorem 14] which contains a flaw.

^{☆☆}Work supported by the EU 7th Framework Programme (FP7/2007-13) under Grants Agreement nr.318490 (SENSATION), nr.601148 (CASSTING), the Sino-Danish Basic Research Center IDEA4CPS funded by Danish National Research Foundation and National Science Foundation China, the ASAP Project (4181-00360) funded by the Danish Council for Independent Research, the ERC Advanced Grant LASSO, and the Innovation Fund Denmark center DiCyPS.

*Corresponding author

Email addresses: `giovbacci@cs.aau.dk` (Giovanni Bacci), `grbacci@cs.aau.dk` (Giorgio Bacci), `kg1@cs.aau.dk` (Kim G. Larsen), `mardare@cs.aau.dk` (Radu Mardare)

Preprint submitted to Journal of Logical and Algebraic Methods in Programming

June 1, 2018

of the alphabet. Their algorithms are based on a partition refinement of the states into equivalence classes of the *Myhill-Nerode equivalence relation*. Partition refinement has been employed in the definition of efficient minimization procedures for a wide variety of automata: by Kanellakis and Smolka [4, 5] for the minimization of labelled transition systems (LTSs) w.r.t. Milner’s strong bisimulation [6]; by Baier [7] for the reduction of Markov Chains (MCs) w.r.t. Larsen and Skou’s probabilistic bisimulation [8]; by Alur et al. [9] and by Yannakakis and Lee [10], respectively, for the minimization of timed transition systems and timed-automata. This technique was used also in parallel and distributed implementations of the above algorithms [11, 12], and in the online reachability analysis of transition systems [13].

In [14], Jou and Smolka observed that for reasoning about the behavior of probabilistic systems (and more in general, all type of quantitative systems), rather than equivalences, a notion of distance is more reasonable in practice, since it permits “*a shift in attention from equivalent processes to probabilistically similar processes*”. This observation motivated the development of metric-based semantics for quantitative systems, that consists in proposing 1-bounded pseudometrics capturing the similarities of the behaviors in the presence of small variations of the quantitative data. These pseudometrics generalize behavioral equivalences in the sense that, two processes are at distance 0 iff they are equivalent, and at distance 1 if no significant similarities can be observed between them.

The first proposal of a behavioral pseudometric is due to Desharnais et al. [15] on labelled MCs, a.k.a. *probabilistic bisimilarity distance*, with the property that two MCs are at distance 0 iff they are probabilistic bisimilar. Its definition is parametric on a discount factor $\lambda \in (0, 1]$ that controls the significance of the future steps in the measurement. This pseudometric has been greatly studied by van Breugel and Worrell [16, 17] who noticed, among other notable results, its relation with the Kantorovich distance on probability distributions and provided a polynomial-time algorithm for its computation [18].

The introduction of metric-based semantics motivated the interest in the approximate minimization of quantitative systems. The goal of approximate minimization is to start from a minimal automaton and produce a smaller automaton that is close to the given one in a certain sense. The desired size of the approximating automaton is given as input. Inspired by the aggregation of equivalent states typical of partition refinement techniques, in [19], the approximate minimization problem has been approached by aggregating states having relative smaller distance. An example of this approach on MCs using the bisimilarity distance of Desharnais et al. is shown in Figure 1. Let \mathcal{M} be the MC on the left and assume we want to approximate it by an MC with at most 5 states. Since m_1, m_2 are the only two states at distance less than 1, the most natural choice for an aggregation shall collapse m_1 and m_2 , obtaining an instance $\mathcal{N}_{x,y}$ of the MC in the middle for some $x, y \geq 0$ such that $x + y \leq 1$. Any MC constructed in this way will not be closer than $\frac{1}{4}$ from \mathcal{M} . However, the MC on the right is a closer approximant of \mathcal{M} , at distance $\frac{1}{6}$ from it, showing that the approximate aggregation of states does not necessarily yield the closest optimal solution.

In this paper we address the issue of finding *optimal* solutions to the approximate minimization problem. Specifically we aim to answer to the following problem, left open in [19]: “*given a finite MC and a positive integer k , what is its ‘best’ k -state approximant? Here by ‘best’ we mean a k -state MC at minimal distance to the original*”. We refer to this problem as *Closest Bounded Approximant* (CBA) and we present the following results




Figure 1: An MC \mathcal{M} (left) with initial state m_0 ; generic 5-states approximant of \mathcal{M} aggregating m_1 and m_2 (for $x, y \in [0, 1]$ such that $x + y \leq 1$) (middle); an optimal 5-states approximant of \mathcal{M} (right). Labels are represented by different colors.

related to it.

1. We characterize CBA as a bilinear optimization problem, proving the existence of *optimal* solutions. As a consequence of this result, approximations of optimal solutions can be obtained by checking the feasibility of bilinear matrix inequalities (BMIs) [20, 21].

2. We provide upper- and lower-bound complexity results for the threshold problem of CBA, called *Bounded Approximant* problem (BA), that asks whether there exists a k -state approximant with distance from the original MC bounded by a given rational threshold. We show that BA is in PSPACE and NP-hard.

3. We introduce the *Minimum Significant Approximant Bound* (MSAB) problem, that asks what is the minimum size k for an approximant to have some significant similarity to the original MC (i.e., at distance strictly less than 1). We show that this problem is NP-complete when one considers the undiscounted bisimilarity distance.

4. Finally, we present an algorithm for finding suboptimal solutions of CBA that is inspired by Expectation Maximization (EM) techniques [22, 23]. Experiments suggest that our method gives a practical approach that outperforms the bilinear program implementation —state-of-the-art bilinear solvers [21] fail to handle MCs with more than 5 states!

Related Work. In [24], the approximate minimization of MCs is addressed via the notion of *quasi-lumpability*. An MC is quasi-lumpable if the given aggregations of the states can be turned into actual bisimulation-classes by a small perturbation of the transition probabilities. This approach differs from ours since there is no relation to a proper notion of behavioral distance (the approximation is w.r.t. the supremum norm of the difference of the stochastic matrices) and we do not consider any approximate aggregation of states. In [25], Balle et al. consider the approximate minimization of weighted finite automata (WFAs). Their method is via a truncation of a canonical normal form for WFAs that they introduced for the SVD decomposition of infinite Hankel matrices. Both [24] and [25] do not consider the issue of finding the *closest* approximant, which is the main focus of this paper, instead they give upper bounds on the distance from the given model.

Synopsis. Section 2 introduces some notation and definitions used in the paper, concluding with a useful characterization of the bisimilarity distance. Section 3 formalizes the CBA- λ problem and characterizes it as a bilinear optimization problem. Sections 4 and 5 introduce two decision problems closely related with CBA- λ , namely BA- λ and

MSAB-1, and study their computational complexity. Sections 6 and 7 present two efficient heuristics for computing sub-optimal solutions for CBA- λ and discuss about their performances and limitations. We conclude with Section 8 discussing possible applications of our results and interesting ideas for future work.

2. Markov Chains and Bisimilarity Pseudometric

In this section we introduce the notation and recall the definitions of (discrete-time) *Markov chains* (MCs), *probabilistic bisimilarity* of Larsen and Skou [8], and the *probabilistic bisimilarity pseudometric* of Desharnais et al. [26].

For $R \subseteq X \times X$ an equivalence relation, X/R denotes its quotient set and $[x]_R$ denotes the R -equivalence class of $x \in X$.

We denote by $\mathcal{D}(X)$ the set of discrete probability distributions on X , i.e., functions $\mu: X \rightarrow [0, 1]$, such that $\mu(X) = 1$, where $\mu(E) = \sum_{x \in E} \mu(x)$ for $E \subseteq X$. For $x \in X$, the Dirac distribution concentrated at x is the function $1_x: X \rightarrow [0, 1]$ defined by $1_x(y) = 1$ if $x = y$, 0 otherwise.

For $a_1, \dots, a_n \in [0, 1]$ such that $\sum_{i=1}^n a_i = 1$, we denote by $\sum_{i=1}^n a_i \cdot \mu_i$ the convex combination of $\mu_1, \dots, \mu_n \in \mathcal{D}(X)$, defined as

$$\left(\sum_{i=1}^n a_i \cdot \mu_i\right)(x) = \sum_{i=1}^n a_i \cdot \mu_i(x).$$

The support of a probability distribution $\mu \in \mathcal{D}(X)$ is defined by

$$\text{support}(\mu) = \{x \in X \mid \mu(x) > 0\}.$$

For example, given $x, y \in X$, $\text{support}\left(\frac{1}{3} \cdot 1_x + \frac{2}{3} \cdot 1_y\right) = \{x, y\}$.

In what follows we fix a countable set L of labels.

Definition 1 (Markov Chain). A *Markov chain* is a tuple $\mathcal{M} = (M, \tau, \ell)$ consisting of a finite nonempty *set of states* M , a *transition distribution function* $\tau: M \rightarrow \mathcal{D}(M)$, and a *labelling function* $\ell: M \rightarrow L$.

Intuitively, if \mathcal{M} is in state m it moves to state m' with probability $\tau(m)(m')$. Labels represent atomic properties that hold in certain states. The set of labels of \mathcal{M} is denoted by $L(\mathcal{M}) = \{\ell(m) \mid m \in M\}$. Hereafter, we use $\mathcal{M} = (M, \tau, \ell)$ and $\mathcal{N} = (N, \theta, \alpha)$ to range over MCs and we refer to their constituents implicitly.

Definition 2 (Probabilistic Bisimulation [8]). An equivalence relation $R \subseteq M \times M$ is a *probabilistic bisimulation* on \mathcal{M} if whenever $m R n$, then

1. $\ell(m) = \ell(n)$, and
2. for all $C \in M/R$, $\tau(m)(C) = \tau(n)(C)$.

Two states $m, n \in M$ are *probabilistic bisimilar w.r.t. \mathcal{M}* , written $m \sim_{\mathcal{M}} n$ if they are related by some probabilistic bisimulation on \mathcal{M} . In fact, probabilistic bisimilarity is the greatest probabilistic bisimulation.

Any bisimulation R on \mathcal{M} induces a quotient construction, the *R -quotient of \mathcal{M}* , denoted $\mathcal{M}/R = (M/R, \tau/R, \ell/R)$, having R -equivalence classes as states, transition function $\tau/R([m]_R)([n]_R) = \sum_{u \in [n]_R} \tau(m)(u)$, and labelling function $\ell/R([m]_R) = \ell(m)$. An MC \mathcal{M} is said *minimal* if it is isomorphic to its quotient w.r.t. probabilistic bisimilarity.

A 1-bounded *pseudometric* on X is a function $d: X \times X \rightarrow [0, 1]$ such that, for any $x, y, z \in X$, $d(x, x) = 0$, $d(x, y) = d(y, x)$, and $d(x, y) + d(y, z) \geq d(x, z)$. 1-bounded pseudometrics on X forms a complete lattice under the point-wise partial order $d \sqsubseteq d'$ iff, for all $x, y \in X$, $d(x, y) \leq d'(x, y)$.

A pseudometric is said to lift an equivalence relation if it enjoys the property that two points are at distance zero iff they are related by the equivalence. A lifting for the probabilistic bisimilarity is provided by the *bisimilarity distance* of Desharnais et al. [26]. Its definition is based on the *Kantorovich (pseudo)metric* on probability distributions over a finite set X , defined as

$$\mathcal{K}(d)(\mu, \nu) = \min \left\{ \sum_{x, y \in X} d(x, y) \cdot \omega(x, y) \mid \omega \in \Omega(\mu, \nu) \right\}, \quad (1)$$

where d is a (pseudo)metric on X and $\Omega(\mu, \nu)$ denotes the set of *couplings* for (μ, ν) , i.e., distributions $\omega \in \mathcal{D}(X \times X)$ such that,

$$\begin{aligned} \sum_{x \in X} \omega(x, y) &= \nu(y) && \text{for all } y \in X \\ \sum_{y \in X} \omega(x, y) &= \mu(x) && \text{for all } x \in X. \end{aligned}$$

The above condition can be equivalently stated as for all $E \subseteq X$, $\omega(E \times X) = \mu(E)$ and $\omega(X \times E) = \nu(E)$, and μ (resp. ν) will be called left (resp. right) marginal of ω .

Remark 1 (Kantorovich as a Transportation Problem). The Kantorovich metric has an intuitive interpretation as the solution of an optimization problem usually referred to as (homogeneous) *mass transportation problem*. Assume you are given a pile of sand and a hole we have to completely fill up with the sand. Obviously, the pile and the hole must have the same volume. Both the pile and the hole are modeled by probability measures $\mu, \nu \in \mathcal{D}(X)$, with $\mu(E)$ giving a measure of how much sand is located in the pile in location $E \subseteq X$ and, $\nu(F)$ how much sand can be piled in $F \subseteq X$. Moving the sand around should be done by minimizing the traveling distance. In this respect a coupling can be interpreted as a transportation plan (or schedule).


A convenient way for visualizing a coupling $\omega \in \mathcal{D}(X \times X)$ for (μ, ν) is by means of the so called *transportation table*, with coordinates given by the supports of the measures μ and ν , and cells containing the value $\omega(x, y)$ at coordinate (x, y) (for convenience, when $\omega(x, y) = 0$ the cell is left blank). In this representation the condition of ω of being a coupling can be easily checked by summing up the values of the cell in each row (resp., column) and checking that it equals the value of the corresponding marginal.

To make explicit the cost of the transportation from x to y , the cell (x, y) is further decorated with the distance $d(x, y)$, displayed in the top-left corner of the cell.

As an example, consider the MC $\mathcal{M} = (M, \tau, \ell)$ in Figure 1 and the transition probabilities $\tau(m_1) = \frac{1}{3} \cdot 1_{m_3} + \frac{2}{3} \cdot 1_{m_4}$ and $\tau(m_2) = \frac{1}{2} \cdot 1_{m_4} + \frac{1}{2} \cdot 1_{m_5}$. The coupling $\omega \in \Omega(\tau(m_1), \tau(m_2))$ defined as

$$\omega = \frac{1}{2} \cdot 1_{(m_4, m_4)} + \frac{1}{3} \cdot 1_{(m_3, m_5)} + \frac{1}{6} \cdot 1_{(m_4, m_5)}$$

is equivalently represented as a transportation table as follows:



Here we consider the distance $d: M \times M \rightarrow [0, 1]$ defined as $d(m, n) = 0$ if $m = n$, 1 otherwise. Note that the above is an optimal coupling in the sense of (1).

Definition 3 (Bisimilarity Distance). Let $\lambda \in (0, 1]$. The λ -discounted *bisimilarity pseudometric* on \mathcal{M} , denoted by δ_λ , is the least fixed-point of the following functional operator on 1-bounded pseudometrics over M (ordered point-wise)

$$\Psi_\lambda(d)(m, n) = \begin{cases} 1 & \text{if } \ell(m) \neq \ell(n) \\ \lambda \cdot \mathcal{K}(d)(\tau(m), \tau(n)) & \text{otherwise.} \end{cases}$$

The operator Ψ_λ is monotonic with respect to \sqsubseteq , hence, by Knaster-Tarski's fixed-point theorem, δ_λ is well defined.

Intuitively, if two states have different labels δ_λ considers them as “incomparable” (i.e., at distance 1), otherwise their distance is given by the Kantorovich distance w.r.t. δ_λ between their transition distributions. The *discount factor* $\lambda \in (0, 1]$ controls the significance of the future steps in the measurement of the distance; if $\lambda = 1$, the distance is said *undiscounted*.

The distance δ_λ has also a characterization based on the notion of coupling structure.

Definition 4 (Coupling Structure). A function $\mathcal{C}: M \times M \rightarrow \mathcal{D}(M \times M)$ is a *coupling structure* for \mathcal{M} if for all $m, n \in M$, $\mathcal{C}(m, n) \in \Omega(\tau(m), \tau(n))$.

Intuitively, a coupling structure can be thought of as an MC on the cartesian product $M \times M$, obtained as the probabilistic combination of two copies of \mathcal{M} .

Given a coupling structure \mathcal{C} for \mathcal{M} and $\lambda \in (0, 1]$, let $\gamma_\lambda^\mathcal{C}$ be the least fixed-point of the following operator on $[0, 1]$ -valued functions $d: M \times M \rightarrow [0, 1]$ (ordered point-wise)

$$\Gamma_\lambda^\mathcal{C}(d)(m, n) = \begin{cases} 1 & \text{if } \ell(m) \neq \ell(n) \\ \lambda \sum_{u, v \in M} d(u, v) \cdot \mathcal{C}(m, n)(u, v) & \text{otherwise.} \end{cases}$$

The function $\gamma_\lambda^\mathcal{C}$ is called λ -discounted *discrepancy* of \mathcal{C} , and the value $\gamma_\lambda^\mathcal{C}(m, n)$ is the λ -discounted probability of hitting from (m, n) a pair of states with different labels in \mathcal{C} .

Theorem 1 (Minimal coupling criterion [18]). For arbitrary MCs \mathcal{M} and discount factors $\lambda \in (0, 1]$, $\delta_\lambda = \min \{ \gamma_\lambda^\mathcal{C} \mid \mathcal{C} \text{ coupling structure for } \mathcal{M} \}$.

As originally noted in [18, Lemma 10], the (undiscounted) discrepancy can be used to bound the variational distance between trace distributions. The following lemma generalizes this result for arbitrary discount values. In the lemma we use $\mathbb{P}_{\mathcal{M},m}(A)$ to denote the probability that a run of the Markov chain \mathcal{M} starting in state m is in the set $A \subseteq L^\omega$. For a formal definition of $\mathbb{P}_{\mathcal{M},m}(A)$ and a definition of measurable subset of the set L^ω of infinite sequences over L , we refer the reader to, e.g., [27, Chapter 10].

For $\lambda \in (0, 1]$ we use \mathcal{M}_λ to denote the MC $(M \uplus \{\perp\}, \tau_\lambda, \ell_\lambda)$ obtained by adding to the MC \mathcal{M} a ‘sink’ state \perp to which all states in \mathcal{M} go with probability $(1 - \lambda)$, that is, for all $m \in M$, $\tau_\lambda(m) = (1 - \lambda)\mathbf{1}_\perp + \lambda\tau(m)$, and $\tau_\lambda(\perp) = \mathbf{1}_\perp$. The ‘sink’ state \perp has label different from all other states, that is, for all $m \in M$, $\ell_\lambda(m) = \ell(m)$, and $\ell_\lambda(\perp) \neq \ell_\lambda(m)$.

Lemma 2. *Let \mathcal{C} be a coupling structure for the MC \mathcal{M} . Then, for any measurable set $A \subseteq L^\omega$, discount factor $\lambda \in (0, 1]$, and $m, n \in M$, $|\mathbb{P}_{\mathcal{M}_\lambda, m}(A) - \mathbb{P}_{\mathcal{M}_\lambda, n}(A)| \leq \gamma_\lambda^{\mathcal{C}}(m, n)$.*

PROOF. Let define $\mathcal{C}_\lambda: (M \uplus \{\perp\})^2 \rightarrow \mathcal{D}((M \uplus \{\perp\})^2)$, for arbitrary $m, n \in M$, as

$$\begin{aligned} \mathcal{C}_\lambda(m, n) &= (1 - \lambda) \cdot \mathbf{1}_{(\perp, \perp)} + \lambda \cdot \mathcal{C}(m, n) & \mathcal{C}_\lambda(\perp, \perp) &= \mathbf{1}_{(\perp, \perp)} \\ \mathcal{C}_\lambda(m, \perp) &= \sum_{u \in M} \tau(m)(u) \cdot \mathbf{1}_{(m, \perp)} & \mathcal{C}_\lambda(\perp, n) &= \sum_{v \in M} \tau(n)(v) \cdot \mathbf{1}_{(\perp, n)}. \end{aligned}$$

One can easily verify that \mathcal{C}_λ is a coupling structure for \mathcal{M}_λ . By [18, Lemma 10],

$$|\mathbb{P}_{\mathcal{M}_\lambda, m}(A) - \mathbb{P}_{\mathcal{M}_\lambda, n}(A)| \leq \gamma_1^{\mathcal{C}_\lambda}(m, n). \quad (2)$$

We conclude the proof by showing that for any $m, n \in M$, $\gamma_1^{\mathcal{C}_\lambda}(m, n) \leq \gamma_\lambda^{\mathcal{C}}(m, n)$.

Let $d: (M \uplus \{\perp\}) \rightarrow [0, 1]$ be defined, for $m, n \in M$, as $d(m, n) = \gamma_\lambda^{\mathcal{C}}(m, n)$, $d(m, \perp) = d(\perp, n) = 1$, and $d(\perp, \perp) = 0$. We show that d is a fixed point for $\Gamma_1^{\mathcal{C}_\lambda}$, that is $\Gamma_1^{\mathcal{C}_\lambda}(d)(m, n) = d(m, n)$ for all $m, n \in M \uplus \{\perp\}$. We show only the cases when m and n have the same label —the other cases are immediate by definition of $\Gamma_1^{\mathcal{C}_\lambda}$ and d .

Case $m = \perp$ and $n = \perp$.

$$\begin{aligned} \Gamma_1^{\mathcal{C}_\lambda}(d)(\perp, \perp) &= \sum_{u, v \in M \uplus \{\perp\}} d(u, v) \cdot \mathcal{C}_\lambda(\perp, \perp)(u, v) & (\text{def. } \Gamma_1^{\mathcal{C}_\lambda}) \\ &= d(\perp, \perp). & (\mathcal{C}_\lambda(\perp, \perp) = \mathbf{1}_{(\perp, \perp)}) \end{aligned}$$

Case $m, n \in M$ and $\ell(m) = \ell(n)$.

$$\begin{aligned} \Gamma_1^{\mathcal{C}_\lambda}(d)(m, n) &= \sum_{u, v \in M \uplus \{\perp\}} d(u, v) \cdot \mathcal{C}_\lambda(m, n)(u, v) & (\text{def. } \Gamma_1^{\mathcal{C}_\lambda}) \\ &= \sum_{u, v \in M} d(u, v) \cdot \mathcal{C}_\lambda(m, n)(u, v) & (\text{def. } \mathcal{C}_\lambda \text{ and } d(\perp, \perp) = 0) \\ &= \lambda \sum_{u, v \in M} d(u, v) \cdot \mathcal{C}(m, n)(u, v) & (\text{def. } \mathcal{C}_\lambda) \\ &= \lambda \sum_{u, v \in M} \gamma_\lambda^{\mathcal{C}}(u, v) \cdot \mathcal{C}(m, n)(u, v) & (\text{def. } d) \\ &= \Gamma_\lambda^{\mathcal{C}}(\gamma_\lambda^{\mathcal{C}})(m, n) & (\text{def. } \Gamma_\lambda^{\mathcal{C}}) \\ &= \gamma_\lambda^{\mathcal{C}}(m, n) & (\text{def. } \gamma_\lambda^{\mathcal{C}}) \\ &= d(m, n). & (\text{def. } d \text{ and } m, n \in M) \end{aligned}$$

By definition, $\gamma_1^{\mathcal{C}_\lambda}$ is the least fixed point of $\Gamma_1^{\mathcal{C}_\lambda}$, therefore, for all $m, n \in M$,

$$\gamma_1^{\mathcal{C}_\lambda}(m, n) \leq d(m, n) = \gamma_\lambda^{\mathcal{C}}(m, n).$$

By the above inequality and (2) we conclude $|\mathbb{P}_{\mathcal{M}_\lambda, m}(A) - \mathbb{P}_{\mathcal{M}_\lambda, n}(A)| \leq \gamma_\lambda^{\mathcal{C}}(m, n)$. \square

The following is a generalization of [18, Corollary 10] for arbitrary discount factors.

Corollary 3. *For any measurable set $A \subseteq L^\omega$, discount factor $\lambda \in (0, 1]$, and $m, n \in M$,*

$$|\mathbb{P}_{\mathcal{M}_\lambda, m}(A) - \mathbb{P}_{\mathcal{M}_\lambda, n}(A)| \leq \delta_\lambda(m, n).$$

PROOF. Immediate consequence of Theorem 1 and Lemma 2. \square

So far we have considered a single Markov chain and described a pseudometric space over its states. The above definitions can be naturally extended to reason about the distance between two MCs by considering the distance induced over their *disjoint union*.

Given two MCs $\mathcal{M} = (M, \tau, \ell)$ and $\mathcal{N} = (N, \theta, \alpha)$ with $M \cap N = \emptyset$, their disjoint union, denoted as $\mathcal{M} \oplus \mathcal{N}$, is the Markov chain having state space $M \cup N$, probability transition function $\vartheta \in (M \cup N) \rightarrow \mathcal{D}(M \cup N)$ and labelling function $l: (M \cup N) \rightarrow L$ respectively defined as

$$\vartheta(x) = \begin{cases} \tau(x) & \text{if } x \in M \\ \theta(x) & \text{if } x \in N \end{cases} \quad l(x) = \begin{cases} \ell(x) & \text{if } x \in M \\ \alpha(x) & \text{if } x \in N. \end{cases}$$

Usually, MCs are associated with an initial state to be thought of as their initial configurations. In the rest of the paper when we talk about the distance between two MCs, written $\delta_\lambda(\mathcal{M}, \mathcal{N})$, we implicitly refer to the distance between their initial states computed over the disjoint union of their MCs¹. Analogously, we may simply write $\mathbb{P}_{\mathcal{M}}(A)$ instead of $\mathbb{P}_{\mathcal{M}, m}(A)$ when m is the initial state of \mathcal{M} .

3. The Closest Bounded Approximant Problem

In this section we introduce the *Closest Bounded Approximant* problem w.r.t. δ_λ (CBA- λ), and give a characterization of it as a bilinear optimization problem.

Definition 5 (Closest Bounded Approximant). Let $k \in \mathbb{N}$ and $\lambda \in (0, 1]$. The *closest bounded approximant problem* w.r.t. δ_λ for an MC \mathcal{M} is the problem of finding an MC \mathcal{N} with at most k states minimizing $\delta_\lambda(\mathcal{M}, \mathcal{N})$.

Clearly, when k is greater than or equal to the number of bisimilarity classes of \mathcal{M} , an optimal solution of CBA- λ is the bisimilarity quotient. Therefore, without loss of generality, we will assume $1 \leq k < |M|$ and \mathcal{M} to be minimal. Note that, under these assumptions \mathcal{M} must have at least two nodes with different labels.

Let $\text{MC}(k)$ denote the set of MCs with at most k states and $\text{MC}_A(k)$ its restriction to those using only labels in $A \subseteq L$. Using this notation, the optimization problem CBA- λ on the instance $\langle \mathcal{M}, k \rangle$ can be reformulated as finding an MC \mathcal{N}^* such that

$$\delta_\lambda(\mathcal{M}, \mathcal{N}^*) = \min \{ \delta_\lambda(\mathcal{M}, \mathcal{N}) \mid \mathcal{N} \in \text{MC}(k) \}. \quad (3)$$

¹Here we implicitly assume that different MCs have disjoint sets of states. This can be done without loss of generality since the behaviors that we observe are equal up-to isomorphism. Therefore, when we write e.g., $\delta_\lambda(\mathcal{M}, \mathcal{M})$, we actually compare \mathcal{M} with another MC \mathcal{M}' isomorphic to \mathcal{M} but whose state space is disjoint from that of \mathcal{M} .

mimimize d_{m_0, n_0}		
such that $\lambda \sum_{(u,v) \in M \times N} c_{u,v}^{m,n} \cdot d_{u,v} \leq d_{m,n}$	$m \in M, n \in N$	(7)
$1 - \alpha_{n,l} \leq d_{m,n} \leq 1$	$n \in N, l \in L(\mathcal{M}), \ell(m) \neq l$	(8)
$\alpha_{n,l} \cdot \alpha_{n,l'} = 0$	$n \in N, l, l' \in L(\mathcal{M}), l \neq l'$	(9)
$\sum_{l \in L(\mathcal{M})} \alpha_{n,l} = 1$	$n \in N$	(10)
$\sum_{v \in N} c_{u,v}^{m,n} = \tau(m)(u)$	$m, u \in M, n \in N$	(11)
$\sum_{u \in M} c_{u,v}^{m,n} = \theta_{n,v}$	$m \in M, n, v \in N$	(12)
$c_{u,v}^{m,n} \geq 0$	$m, u \in M, n, v \in N$	(13)

Figure 2: Characterization of CBA- λ as a bilinear optimization problem.

In general, it is not obvious that for arbitrary instances $\langle \mathcal{M}, k \rangle$ a minimum in (3) exists. At the end of the section, we will show that such a minimum always exists (Corollary 6).

A useful property of CBA- λ is that an optimal solution can be found among the MCs using labels from the given MC.

Lemma 4 (Meaningful labels). *Let \mathcal{M} be an MC. Then, for any $\mathcal{N}' \in \text{MC}(k)$ there exists $\mathcal{N} \in \text{MC}_{L(\mathcal{M})}(k)$ such that $\delta_\lambda(\mathcal{M}, \mathcal{N}) \leq \delta_\lambda(\mathcal{M}, \mathcal{N}')$.*

PROOF. Let $\mathcal{N}' = (N', \theta', \alpha')$. If $L(\mathcal{N}') \subseteq L(\mathcal{M})$, take $\mathcal{N} = \mathcal{N}'$. Otherwise, define $\mathcal{N} = (N, \theta, \alpha)$ as follows: $N = N'$, $\theta = \theta'$, and $\alpha(n) = \alpha'(n)$ if $\alpha'(n) \in L(\mathcal{M})$, otherwise $\alpha(n) = \ell(m_0)$, where m_0 is the initial state of \mathcal{M} . The initial state of \mathcal{N} is the one of \mathcal{N}' . Clearly, $\mathcal{N} \in \text{MC}(k)$ and $L(\mathcal{N}) \subseteq L(\mathcal{M})$.

Let $\mathcal{A} = \mathcal{M} \oplus \mathcal{N}$ and $\mathcal{B} = \mathcal{M} \oplus \mathcal{N}'$, we prove $\delta_\lambda(\mathcal{M}, \mathcal{N}) \leq \delta_\lambda(\mathcal{M}, \mathcal{N}')$ by showing a stronger statement: $\delta_\lambda^{\mathcal{A}} \sqsubseteq \delta_\lambda^{\mathcal{B}}$. By Knaster-Tarski's fixed-point theorem, it suffices to show $\Psi_\lambda^{\mathcal{A}}(\delta_\lambda^{\mathcal{B}}) \sqsubseteq \delta_\lambda^{\mathcal{B}}$. Let $u, v \in M \uplus N$. When u and v have different labels in \mathcal{B} , then, $\Psi_\lambda^{\mathcal{A}}(\delta_\lambda^{\mathcal{B}})(u, v) \leq 1 = \delta_\lambda^{\mathcal{B}}(u, v)$ follows by definition of Ψ_λ and the fact that $\delta_\lambda^{\mathcal{B}} = \Psi_\lambda^{\mathcal{B}}(\delta_\lambda^{\mathcal{B}})$. Assume that u and v have the same label in \mathcal{B} . Then, by construction of \mathcal{N} (i.e. by definition of α'), u and v have the same label in \mathcal{A} . Since \mathcal{N} and \mathcal{N}' have the same transition distribution function, we obtain $\Psi_\lambda^{\mathcal{A}}(\delta_\lambda^{\mathcal{B}})(u, v) = \delta_\lambda^{\mathcal{B}}(u, v)$. \square

In the following, fix $\langle \mathcal{M}, k \rangle$ as instance of CBA- λ , let $m_0 \in M$ be the initial state of \mathcal{M} . By Lemma 4, Theorem 1 and Knaster-Tarski fixed-point theorem

$$\inf \{ \delta_\lambda(\mathcal{M}, \mathcal{N}) \mid \mathcal{N} \in \text{MC}(k) \} = \tag{4}$$

$$= \inf \{ \gamma_\lambda^{\mathcal{C}}(\mathcal{M}, \mathcal{N}) \mid \mathcal{N} \in \text{MC}_{L(\mathcal{M})}(k) \text{ and } \mathcal{C} \in \Omega(\mathcal{M}, \mathcal{N}) \} \tag{5}$$

$$= \inf \{ d(\mathcal{M}, \mathcal{N}) \mid \mathcal{N} \in \text{MC}_{L(\mathcal{M})}(k), \mathcal{C} \in \Omega(\mathcal{M}, \mathcal{N}), \text{ and } \Gamma_\lambda^{\mathcal{C}}(d) \sqsubseteq d \} , \tag{6}$$

where $\Omega(\mathcal{M}, \mathcal{N})$ denotes the set of all coupling structures for the disjoint union of \mathcal{M} and \mathcal{N} . This simple change in perspective yields a translation of the problem of computing the optimal value of CBA- λ to the bilinear program in Figure 2.

In our encoding, $N = \{n_0, \dots, n_{k-1}\}$ are the states of an arbitrary Markov chain with k states $\mathcal{N} = (N, \theta, \alpha) \in \text{MC}(k)$. and n_0 is the initial one. The variable $\theta_{n,v}$ is used to

encode the transition probability $\theta(n)(v)$. Hence, a feasible solution satisfying (11–13) will have the variable $c_{u,v}^{m,n}$ representing the value $\mathcal{C}(m,n)(u,v)$ for a coupling structure $\mathcal{C} \in \Omega(\mathcal{M}, \mathcal{N})$. An assignment for the variables $\alpha_{n,l}$ satisfying (9–10) encodes (uniquely) a labeling function $\alpha: N \rightarrow L(\mathcal{M})$ satisfying the following property:

$$\text{for all } n \in N, l \in L(\mathcal{M}) \quad \alpha_{n,l} = 1 \quad \text{iff} \quad \alpha(n) = l. \quad (14)$$

The constraint (9) models the fact that each node $n \in N$ is assigned at most to one label $l \in L(\mathcal{M})$, and the constraint (10) ensures that each node is assigned to at least one label in $L(\mathcal{M})$. Conversely, any labeling $\alpha: N \rightarrow L(\mathcal{M})$ admits an assignment of the variables $\alpha_{n,l}$ satisfying (9–10) and (14). Finally, an assignment for the variables $d_{m,n}$ satisfying the constraints (7–8) represents a prefix point of $\Gamma_\lambda^{\mathcal{C}}$. Note that (8) guarantees that $d_{m,n} = 1$ whenever $\alpha(n) \neq \ell(m)$ —indeed, by (9), $\alpha_{n,l} = 0$ iff $\alpha(n) \neq \ell(m)$.

Let $F_\lambda \langle \mathcal{M}, k \rangle$ denote the bilinear optimization problem in Figure 2. Directly from the arguments stated above we obtain the following result.


Theorem 5. $\inf \{ \delta_\lambda(\mathcal{M}, \mathcal{N}) \mid \mathcal{N} \in \text{MC}(k) \}$ is the optimal value of $F_\lambda \langle \mathcal{M}, k \rangle$.

Corollary 6. Any instance of CBA- λ admits an optimal solution.

PROOF. We show that $\delta_\lambda(\mathcal{M}, \mathcal{N}^*) = \inf \{ \delta_\lambda(\mathcal{M}, \mathcal{N}) \mid \mathcal{N} \in \text{MC}(k) \}$ for some $\mathcal{N}^* \in \text{MC}(k)$. Let h be the number of variables in $F_\lambda \langle \mathcal{M}, k \rangle$. The constraints (8–13) describe a compact subset of \mathbb{R}^h —it is an intersection of closed sets bounded by $[0, 1]^h$. The objective function of $F_\lambda \langle \mathcal{M}, k \rangle$ is linear, hence the infimum is attained by a feasible solution. The thesis follows by Theorem 5. \square

The next example shows that even by starting with an MC with rational transition probabilities, optimal solutions for CBA- λ may have irrational transition probabilities.

Example 1. Consider the MC \mathcal{M} depicted below, with initial state m_0 and labels represented by distinct colors. We claim that the MC \mathcal{N}_{xy} depicted below, with initial state n_0 and parameters $x = \frac{1}{30}(10 + \sqrt{163})$, $y = \frac{21}{100}$, is an optimal solution of CBA-1 on input $\langle \mathcal{M}, 3 \rangle$.



Since the distance $\delta_1(\mathcal{M}, \mathcal{N}_{xy}) = \frac{436}{675} - \frac{163\sqrt{163}}{13500} \approx 0.49$ is irrational, by [18, Proposition 13], any optimal solution must have some irrational transition probability.

Now we prove that the above is indeed an optimal solution for CBA-1 on input $\langle \mathcal{M}, 3 \rangle$. Assume by contradiction that there exists $\mathcal{N}^* \in \text{MC}(3)$ s.t. $\delta_1(\mathcal{M}, \mathcal{N}^*) < \delta_1(\mathcal{M}, \mathcal{N}_{xy})$. Without loss of generality, by Lemma 4 we can assume that $L(\mathcal{N}^*) \subseteq L(\mathcal{M})$.

If $L(\mathcal{N}^*) = L(\mathcal{M})$, then \mathcal{N}^* must be an MC of the form \mathcal{N}_{zw} for some of $z, w \in [0, 1]$ such that $z + w \leq 1$. Thus,

$$\begin{aligned} \delta_1(\mathcal{M}, \mathcal{N}^*) &\geq \min \{ \delta_1(\mathcal{M}, \mathcal{N}_{zw}) \mid z, w \geq 0, z + w \leq 1 \} \\ &= \min \{ \gamma_1^{\mathcal{C}}(m_0, n_0) \mid \mathcal{C} \in \Omega(\mathcal{M}, \mathcal{N}_{zw}), z, w \geq 0, z + w \leq 1 \}. \quad (\text{Thm. 1}) \end{aligned} \quad (15)$$




Figure 3: Family of coupling structures for \mathcal{M} and $\mathcal{N}_{z,w}$ with $z \in [0, \frac{79}{100}]$ and $w = \frac{21}{100}$ (cf. Example 1).

Consider an arbitrary coupling structure $\mathcal{C} \in \Omega(\mathcal{M}, \mathcal{N}_{z,w})$ for some $z, w \geq 0$ such that $z + w \leq 1$. By definition of $\gamma_1^{\mathcal{C}}$, for any $i \in \{0, 1, 2\}$,

$$\gamma_1^{\mathcal{C}}(m_i, n_0) = A_i + \sum_{j=0}^1 \gamma_1^{\mathcal{C}}(m_{i+1}, n_j) \cdot \mathcal{C}(m_i, n_0)(m_{i+1}, n_j), \quad (16)$$

where $A_i = \mathcal{C}(m_i, n_0)(m_{i+1}, n_2) + \mathcal{C}(m_i, n_0)(m_4, n_0) + \mathcal{C}(m_i, n_0)(m_4, n_1)$. The constraints on the marginals require that, for all $i \in \{0, 1, 2\}$,

$$w = \mathcal{C}(m_i, n_0)(m_{i+1}, n_2) + \mathcal{C}(m_i, n_0)(m_4, n_2), \text{ and}$$

$$\frac{21}{100} = \mathcal{C}(m_i, n_0)(m_4, n_0) + \mathcal{C}(m_i, n_0)(m_4, n_1) + \mathcal{C}(m_i, n_0)(m_4, n_2).$$

Consequently, for minimizing $\gamma_1^{\mathcal{C}}(m_i, n_0)$, we shall choose the values of $\mathcal{C}(m_i, n_0)$ in a way that makes $A_i = 0$ (hence $\mathcal{C}(m_i, n_0)(m_4, n_2) = \frac{21}{100}$). Therefore, we can restrict our attention to the family of coupling structures depicted in Figure 1

$$\delta_1(\mathcal{M}, \mathcal{N}^*) \geq \min \left\{ \gamma_1^{\mathcal{C}}(m_0, n_0) \mid \mathcal{C} \in \Omega(\mathcal{M}, \mathcal{N}_{z,w}), 0 \leq z \leq \frac{79}{100}, w = \frac{21}{100} \right\}. \quad (17)$$

We compute $\gamma_1^{\mathcal{C}}(m_0, n_0)$ backwards

$$\begin{aligned} \gamma_1^{\mathcal{C}}(m_4, n_2) &= \gamma_1^{\mathcal{C}}(m_3, n_1) = 0 \\ \gamma_1^{\mathcal{C}}(m_2, n_0) &= z \\ \gamma_1^{\mathcal{C}}(m_1, n_0) &= z \cdot \gamma_1^{\mathcal{C}}(m_2, n_0) + 1 - z - w \\ \gamma_1^{\mathcal{C}}(m_0, n_0) &= z \cdot \gamma_1^{\mathcal{C}}(m_1, n_0) + 1 - z - w \end{aligned}$$

thus, the inequality (17) simplifies to

$$\delta_1(\mathcal{M}, \mathcal{N}^*) \geq \min \left\{ z^3 - z^2 - \frac{21}{100}z + \frac{79}{100} \mid 0 \leq z \leq \frac{79}{100} \right\}$$

The minimum value of the above is achieved at $z = \frac{1}{30}(10 + \sqrt{163})$. This contradicts the initial assumption $\delta_1(\mathcal{M}, \mathcal{N}^*) < \delta_1(\mathcal{M}, \mathcal{N}_{xy})$ excluding the hypothesis $L(\mathcal{N}^*) = L(\mathcal{M})$.

Let's consider the case $L(\mathcal{N}^*) \subsetneq L(\mathcal{M})$ and analyse three possible sub-cases.

If $\ell(m_0) \notin L(\mathcal{N}^*)$, then no state in \mathcal{N} has the same label as the initial state of \mathcal{M} . Therefore, by definition of δ_1 we have $\delta_1(\mathcal{M}, \mathcal{N}^*) = 1$, which is greater than $\delta_1(\mathcal{M}, \mathcal{N}_{xy})$.

If $\ell(m_3) \notin L(\mathcal{N}^*)$ we have

$$\begin{aligned} \delta_1(\mathcal{M}, \mathcal{N}^*) &\geq |\mathbb{P}_{\mathcal{M}}(\ell(m_3)L^\omega) - \mathbb{P}_{\mathcal{N}^*}(\ell(m_3)L^\omega)| && \text{(Corollary 3)} \\ &= \mathbb{P}_{\mathcal{M}}(L^*\ell(m_3)L^\omega) && (\mathbb{P}_{\mathcal{N}^*}(\ell(m_3)L^\omega) = 0) \\ &= \left(\frac{79}{100}\right)^3 && \text{(def. } \mathbb{P}_{\mathcal{M}}) \\ &> \delta_1(\mathcal{M}, \mathcal{N}_{xy}). \end{aligned}$$

Analogously, if $\ell(m_4) \notin L(\mathcal{N}^*)$,

$$\delta_1(\mathcal{M}, \mathcal{N}^*) \geq \mathbb{P}_{\mathcal{M}}(L^*\ell(m_4)L^\omega) = \frac{21}{100} \sum_{i=0}^2 \left(\frac{79}{100}\right)^i > \delta_1(\mathcal{M}, \mathcal{N}_{xy}).$$

Therefore \mathcal{N}_{xy} is an optimal solution. \square

4. The Bounded Approximant Threshold Problem

The *Bounded Approximant problem* w.r.t. δ_λ (BA- λ) is the threshold decision problem of CBA- λ , that, given MC \mathcal{M} , integer $k \geq 1$, and rational $\epsilon \geq 0$, asks whether there exists $\mathcal{N} \in \text{MC}(k)$ such that $\delta_\lambda(\mathcal{M}, \mathcal{N}) \leq \epsilon$.

From the characterization of CBA- λ as a bilinear optimization problem (Theorem 5) we immediately get the following complexity upper-bound for BA- λ .

Theorem 7. *For any $\lambda \in (0, 1]$, BA- λ is in PSPACE.*

PROOF. By Theorem 5, deciding an instance $\langle \mathcal{M}, k, \epsilon \rangle$ of BA- λ can be encoded as a decision problem for the existential theory of the reals, namely, checking the feasibility of the constraints (8–13) in conjunction with $d_{m_0, n_0} \leq \epsilon$. The encoding is polynomial in the size of $\langle \mathcal{M}, k, \epsilon \rangle$, thus it can be solved in PSPACE (cf. Canny [28]). \square

In the rest of the section we provide a complexity lower-bound for BA- λ , by showing that BA- λ is NP-hard via a reduction from VERTEX COVER. Recall that, a vertex cover of an undirected graph G is a subset C of vertices such that every edge in G has at least one endpoint in C . Given a graph G and a positive integer h , the VERTEX COVER problem asks if G has a cover of size at most h .

The following lemma provides a lower-bound on the λ -discounted bisimilarity distance between \mathcal{M} and any $\mathcal{N} \in \text{MC}(k)$.

Lemma 8. *For all $m \in M$ and $n \in N$, $\delta_\lambda(m, n) \geq \lambda \cdot \tau(m)(\{u \in M \mid \ell(u) \notin L(\mathcal{N})\})$.*

PROOF. The thesis holds trivially when $\ell(m) \neq \alpha(n)$, since $\delta_\lambda(m, n) = 1$.




Figure 4: (Left) An undirected graph G ; (Center) The MC \mathcal{M}_G associated to the graph G ; (Right) The MC \mathcal{M}_C associated to the vertex cover $C = \{2, 3\}$ of G . (see Thm. 9).

Let $\ell(m) = \alpha(n)$, and $M' = \{u \in M \mid \ell(u) \notin L(\mathcal{N})\}$, then the following hold

$$\begin{aligned}
\delta_\lambda(m, n) &= \lambda \sum_{u \in M} \sum_{v \in N} \delta_\lambda(u, v) \cdot \omega(u, v) && \text{(for some } \omega \in \Omega(\tau(m), \theta(n))\text{)} \\
&\geq \lambda \sum_{u \in M'} \sum_{v \in N} \delta_\lambda(u, v) \cdot \omega(u, v) && \text{(by } M' \subseteq M\text{)} \\
&= \lambda \sum_{u \in M'} \sum_{v \in N} \omega(u, v) && (\delta_\lambda(u, v) = 1 \text{ for all } u \in M' \text{ and } n \in N) \\
&= \lambda \cdot \tau(m)(M') && \text{(by } \omega \in \Omega(\tau(m), \theta(n))\text{)}
\end{aligned}$$

□

We are now ready to present the main result of this session.

Theorem 9. *For any $\lambda \in (0, 1]$, BA- λ is NP-hard.*

PROOF. We provide a polynomial-time many-one reduction from VERTEX COVER.

Let $\langle G = (V, E), h \rangle$ be an instance of VERTEX COVER and let $m = |E|$. Without loss of generality we assume $m \geq 2$.

From G we construct the MC $\mathcal{M}_G = (M, \tau, \ell)$ having the following states: a root state r (thought of as the initial state); a sink state s ; a state v for each vertex in V ; and two ‘twin’ states e and e' for each edge $e \in E$. In \mathcal{M}_G each pair of twin edge states have the same label, all other nodes have pairwise distinct labels. The sink state s and all vertex states loop to themselves with probability 1; the root state r goes with uniform probability $\frac{1}{2m}$ to each edge state; for each edge $e = (v_1, v_2)$ in G , the state e (resp. e') in \mathcal{M}_G goes with probability $\frac{1}{m}$ to v_1 (resp. v_2) and probability $1 - \frac{1}{m}$ to the sink state s (cf. Figure 4 for an example of the construction of \mathcal{M}_G). Next we show that

$$\langle G, h \rangle \in \text{VERTEX COVER} \quad \text{iff} \quad \langle \mathcal{M}_G, m + h + 2, \frac{\lambda^2}{2m} \rangle \in \text{BA-}\lambda.$$

(\Rightarrow) Let C be a h -vertex cover of G . Construct $\mathcal{M}_C \in \text{MC}(m + h + 2)$ by taking a copy of \mathcal{M}_G , removing all vertex states in $V \setminus C$, then removing one twin edge state for each edge in E making sure to keep those which are going with probability greater than zero to some vertex state in C —if both endpoints of the edge are in the cover C ,

just pick one twin edge state at random. Finally, redistribute uniformly the transition probabilities of the root state over the remaining edge states (*cf.* Figure 4).

Next we show that $\delta_\lambda(\mathcal{M}_G, \mathcal{M}_C) \leq \frac{\lambda^2}{2m}$. For convenience, the states in \mathcal{M}_C will be marked with a bar to distinguish them from their counterpart in \mathcal{M}_G . By construction of $\mathcal{M}_G, \mathcal{M}_C$, for each $e \in E$, $\delta_\lambda(e, \bar{e}) + \delta_\lambda(e', \bar{e}) = \frac{\lambda}{m}$ since either e or e' is at distance 0 from \bar{e} while the other state differs from \bar{e} only for the transition to a vertex state. Therefore,

$$\delta_\lambda(\mathcal{M}_G, \mathcal{M}_C) = \delta_\lambda(r, \bar{r}) = \frac{\lambda}{2m} \sum_{e \in E} (\delta_\lambda(e, \bar{e}) + \delta_\lambda(e', \bar{e})) = \frac{\lambda^2}{2m}.$$

(\Leftarrow) Assume that there exists $\mathcal{N} = (N, \theta, \alpha) \in \text{MC}(m+h+2)$ s.t. $\delta_\lambda(\mathcal{M}_G, \mathcal{N}) \leq \frac{\lambda^2}{2m}$. We claim that G has a vertex cover of size h . Without loss of generality we may assume that the following hold for \mathcal{N} :

1. $L(\mathcal{N}) \subseteq L(\mathcal{M}_G)$;
2. \mathcal{N} has initial state \bar{r} with the same label than r , i.e., $\alpha(\bar{r}) = \ell(r)$;
3. $\ell(s) \in L(\mathcal{N})$; and
4. $\{\ell(e) \mid e \in E\} \subseteq L(\mathcal{N})$.
5. \mathcal{N} has minimal size, i.e., $|N|$ is minimal

Assumption (1) follows by Lemma 4. Assumptions (2), (3) and (4) are necessary to ensure $\delta_\lambda(\mathcal{M}_G, \mathcal{N}) \leq \frac{\lambda^2}{2m}$. Indeed, if $\alpha(\bar{r}) \neq \ell(r)$, $\delta_\lambda(\mathcal{M}_G, \mathcal{N}) = 1 > \frac{\lambda^2}{2m}$. If $\ell(s) \notin L(\mathcal{N})$, then we get the following contradiction

$$\begin{aligned} \delta_\lambda(\mathcal{M}_G, \mathcal{N}) &\geq |\mathbb{P}_{(\mathcal{M}_G)_\lambda}(L^* \ell(s) L^\omega) - \mathbb{P}_{(\mathcal{N})_\lambda}(L^* \ell(s) L^\omega)| && \text{(Corollary 3)} \\ &= \mathbb{P}_{(\mathcal{M}_G)_\lambda}(L^* \ell(s) L^\omega) && (\mathbb{P}_{(\mathcal{N})_\lambda}(L^* \ell(s) L^\omega) = 0) \\ &= \lambda^2 \left(1 - \frac{1}{m}\right) && \text{(definition of } \mathbb{P}) \\ &> \frac{\lambda^2}{2m}. && (m \geq 2) \end{aligned}$$

Finally, if $\ell(e) \notin L(\mathcal{N})$, for some $e \in E$, then we get the following contradiction

$$\begin{aligned} \delta_\lambda(\mathcal{M}_G, \mathcal{N}) &= \delta_\lambda(r, \bar{r}) && (r \text{ and } \bar{r} \text{ are the initial states)} \\ &\geq \lambda \cdot \tau(r)(\{e, e'\}) && \text{(Lemma 8)} \\ &= \frac{\lambda}{m} && \text{(definition of } \mathcal{M}_G) \\ &> \frac{\lambda^2}{2m}. && (\lambda \in (0, 1]) \end{aligned}$$

By assumption (5) we have that the sink state s in \mathcal{M}_G requires exactly one sink state \bar{s} in \mathcal{N} with same label than s and having a self loop with probability 1. Similarly, for vertex states v in \mathcal{M} , \mathcal{N} requires at most one vertex state \bar{v} with same label than v and having a self loop with probability 1. Lastly, for each edge $e \in E$, \mathcal{N} requires at most




Figure 5: (Left) Couplings $\omega \in \Omega(\tau(r), \mu)$ parametric in $\epsilon, \epsilon' \in [0, \frac{1}{2m}]^m$, for $m = 2$; (Right) Couplings $\omega \in \Omega(\tau(e), \nu_x)$ and $\tilde{\omega} \in \Omega(\tau(e'), \nu_x)$ parametric in $x \in [0, \frac{1}{m}]$ (cf. Theorem 9).

two edge states \bar{e} and \bar{e}' to respectively represent the edge states e and e' in \mathcal{M}_G ; clearly $\ell(e) = \alpha(\bar{e})$ and $\ell(e') = \alpha(\bar{e}')$.

By the assumptions made above, we have that $\{\bar{r}, \bar{s}\} \subseteq N$ and $\{\bar{e}, \bar{e}'\} \cap N \neq \emptyset$, for each edge $e \in E$. By Theorem 1, there exists $\mathcal{C} \in \Omega(\mathcal{M}, \mathcal{N})$ such that $\delta_\lambda = \gamma_\lambda^{\mathcal{C}}$. Then,

$$\begin{aligned}
\delta_\lambda(\mathcal{M}_G, \mathcal{N}) &= \delta_\lambda(r, \bar{r}) && (r \text{ and } \bar{r} \text{ are the initial states}) \\
&= \lambda \sum_{m \in M, n \in N} \delta_\lambda(m, n) \cdot \mathcal{C}(r, \bar{r})(m, n) && (\text{def. } \Gamma_\lambda^{\mathcal{C}} \text{ and } \gamma_\lambda^{\mathcal{C}}) \\
&\geq \lambda \cdot \min \left\{ \sum_{m \in M, n \in N} \delta_\lambda(m, n) \cdot \omega(m, n) \mid \begin{array}{l} \omega \in \Omega(\tau(r), \mu) \\ \mu \in \mathcal{D}(N) \end{array} \right\} \\
&= \lambda \sum_{e \in E} \min \left\{ f_e(\epsilon_e, \epsilon'_e) \mid \epsilon_e, \epsilon'_e \in \left[0, \frac{1}{2m}\right] \right\}, && (18)
\end{aligned}$$

where for arbitrary $e \in E$, f_e is a function in ϵ_e and ϵ'_e defined as

$$f_e(\epsilon_e, \epsilon'_e) = \left(\frac{1}{2m} - \epsilon_e\right) \cdot \delta_\lambda(e, \bar{e}) + \left(\frac{1}{2m} - \epsilon'_e\right) \cdot \delta_\lambda(e', \bar{e}) + \epsilon_e \cdot \delta_\lambda(e, \bar{e}') + \epsilon'_e \cdot \delta_\lambda(e', \bar{e}').$$

The equality (18) follows by the fact that an optimal coupling ω can be found among those with support included in $\bigcup_{e \in E} \{(e, \bar{e}), (e', \bar{e}), (e, \bar{e}'), (e', \bar{e}')\}$ (cf. Figure 5 (left)). Note that the above formulation is general enough to model the case when $\bar{e} \notin N$ (resp. $\bar{e}' \notin N$) in which case $\epsilon_e = \epsilon'_e = \frac{1}{2m}$ (resp. $\epsilon_e = \epsilon'_e = 0$).

Consider an arbitrary edge $e \in E$ with endpoints $u, v \in V$. Next, we lower-bounds for $\min \{f_e(\epsilon_e, \epsilon'_e) \mid \epsilon_e, \epsilon'_e \in [0, \frac{1}{2m}]\}$. The edge e may be covered by \bar{V} (i.e., \bar{u} or \bar{v} appear in \mathcal{N} , cf. Case B.) or not (i.e., neither \bar{u} nor \bar{v} appear in \mathcal{N} , cf. Case A.).

A. If neither \bar{u} nor \bar{v} appear in \mathcal{N} . Then, by Lemma 8, for arbitrary $p \in \{e, e'\}$ and

$q \in \{\bar{e}, \bar{e}'\}$, we have $\delta_\lambda(p, q) \geq \frac{\lambda}{m}$. Therefore, for arbitrary $\epsilon_e, \epsilon'_e \in [0, \frac{1}{2m}]$

$$f_e(\epsilon_e, \epsilon'_e) \geq \frac{\lambda}{m} \left(\frac{1}{2m} - \epsilon_e + \frac{1}{2m} - \epsilon'_e + \epsilon_e + \epsilon'_e \right) = \frac{\lambda}{m^2}.$$

B. If \bar{u} or \bar{v} appear in \mathcal{N} . We distinguish two subcases: when both \bar{e} and \bar{e}' appear in \mathcal{N} , or only one of them.

B.1. If $\bar{e}' \notin N$ (resp. $\bar{e} \notin N$), then $\epsilon_e = \epsilon'_e = 0$ (resp. $\epsilon_e = \epsilon'_e = \frac{1}{2m}$).

$$\begin{aligned} & f_e(\epsilon_e, \epsilon'_e) \\ &= \frac{1}{2m} (\delta_\lambda(e, \bar{e}) + \delta_\lambda(e', \bar{e})) \quad (\epsilon_e = \epsilon'_e = 0) \\ &= \frac{\lambda}{2m} \sum_{p \in M, q \in N} (\mathcal{C}(e, \bar{e})(p, q) + \mathcal{C}(e', \bar{e})(p, q)) \cdot \delta_\lambda(p, q) \quad (\text{def. } \mathcal{C} \text{ and } \gamma_\lambda^{\mathcal{C}}) \\ &\geq \frac{\lambda}{2m} \cdot \min \left\{ \sum_{p \in M, q \in N} (\omega(p, q) + \tilde{\omega}(p, q)) \cdot \delta_\lambda(p, q) \left| \begin{array}{l} \omega \in \Omega(\tau(e), \nu) \\ \tilde{\omega} \in \Omega(\tau(e'), \nu) \\ \nu \in \mathcal{D}(N) \end{array} \right. \right\} \\ &= \frac{\lambda}{2m} \cdot \min \left\{ x + \left(\frac{1}{m} - x \right) \left| 0 \leq x \leq 1 - \frac{1}{m} \right. \right\} \quad (\text{cf. Figure. 5 (right)}) \\ &= \frac{\lambda}{2m^2}. \end{aligned}$$

Analogously, the case $\bar{e} \notin N$ has the same lower-bound.

B.2. If both \bar{e} and \bar{e}' are in \mathcal{N} , then $0 < \epsilon_e + \epsilon'_e < \frac{1}{m}$. We further split in two subcases: when both \bar{u} and \bar{v} appear in \mathcal{N} , or only one of them.

B.2.1. If both \bar{u} and \bar{v} are in \mathcal{N} , then $f_e(\epsilon_e, \epsilon'_e) \geq 0$ because $\delta_\lambda(p, q) \geq 0$ for arbitrary $p \in M$ and $q \in N$.

B.2.2. If $\bar{u} \notin N$ (resp. $\bar{v} \notin N$), then, for any $\epsilon_e, \epsilon'_e \in [0, \frac{1}{m}]$

$$\begin{aligned} & f_e(\epsilon_e, \epsilon'_e) \\ &\geq \left(\frac{1}{2m} - \epsilon_e \right) \cdot \delta_\lambda(e, \bar{e}) + \epsilon_e \cdot \delta_\lambda(e, \bar{e}') \quad (\delta_\lambda(e', q) \geq 0 \text{ for any } q \in N) \\ &\geq \lambda \cdot \tau(e)(u) \cdot \left(\frac{1}{2m} - \epsilon_e + \epsilon_e \right) \quad (\text{Lemma 8}) \\ &= \frac{\lambda}{2m^2}. \quad (\tau(e)(u) = \frac{1}{m}) \end{aligned}$$

Similarly, the case $\bar{v} \notin N$ has the same lower-bound.

By assigning with each edge $e \in E$ any among the cases A, B.1, B.2.1, and B.2.2 described above, we induce a selection of the states of \mathcal{N} and, at the same time, a selection of vertices $\bar{V} \subseteq V$ in the graph G , namely, $\bar{V} = \{v \in V \mid \bar{v} \in N\}$. Recall that at least m states in \mathcal{N} have to be edge vertices and two other states are reserved respectively for the sink \bar{s} and the initial state \bar{r} , therefore $|\bar{V}| \leq h$.



Figure 6: (Left) By removing one of the edge states we turn the assignment from type B.2.1 to type B.1 or B.2.1; (Right) By adding one state vertex corresponding to one endpoint of the edge we turn the assignment from type A to type B.1 or B.2.1.

Assume that the m edges in G are assigned the above case as follows:

$$m = \underbrace{m'}_{\text{B.2.1}} + \underbrace{m''}_{\text{A}} + \underbrace{m'''}_{\text{B.1 or B.2.2}} .$$

In the above assignment $m' + m'''$ edges are covered by \bar{V} whereas m'' edges are not covered by \bar{V} . Necessarily, $m'' \leq m'$. Otherwise, we have that

$$\begin{aligned} \delta_\lambda(\mathcal{M}, \mathcal{N}) &\geq \lambda \sum_{e \in E} \min \left\{ f_e(\epsilon_e, \epsilon'_e) \mid \epsilon_e, \epsilon'_e \in \left[0, \frac{1}{2m} \right] \right\} && \text{(by Equation (18))} \\ &\geq \lambda \left(m'' \cdot \frac{\lambda}{m^2} + m''' \cdot \frac{\lambda}{2m^2} \right) && \text{(lower-bounds proved before)} \\ &= \lambda \left(m'' \cdot \frac{\lambda}{m^2} + (m - m' - m''') \cdot \frac{\lambda}{2m^2} \right) && (m''' = m - m' - m'') \\ &= \frac{\lambda^2}{2m^2} \cdot (m - m' + m'') \\ &> \frac{\lambda^2}{2m} && (m'' > m') \end{aligned}$$

which contradicts the assumption that $\delta_\lambda(\mathcal{M}, \mathcal{N}) \leq \frac{\lambda^2}{2m}$.

If $m'' = 0$, then \bar{V} is a vertex cover for G and we are done. Otherwise, we claim that there is another assignment of the edges where no edge is of type A inducing a vertex cover. Such assignment is obtained from the previous one, leveraging from the fact that $m'' \leq m'$. We proceed by turning all m'' edges of type A to edges of type B.1 or B.2.2 removing exactly one of the two edge states from (at most m'') edges of type B.2.1 (cf. Fig. 6 (left)). This gives enough room to add vertex states in a way that all edges of type A are turned into edges of type B.1 or B.2.2 (cf. Fig. 6 (right)).

This concludes the proof of the reduction from VERTEX COVER. \square

5. Minimum Significant Approximant Bound

Recall that, having two MCs that are at distance 1 from each other means that there is no significant similarity between their behaviors. Accordingly, we say that an MC \mathcal{N} is a *significant approximant* for the MC \mathcal{M} w.r.t. δ_λ if, and only if, $\delta_\lambda(\mathcal{M}, \mathcal{N}) < 1$.


The *Minimum Significant Approximant Bound* problem w.r.t. δ_λ (MSAB- λ) looks for the smallest positive integer k such that $\mathcal{N} \in \text{MC}(k)$ is a significant approximant for

a given an MC \mathcal{M} . The decision version of this problem is called *Significant Bounded Approximant problem* w.r.t. δ_λ (SBA- λ), and asks whether, for a given positive integer k , there exists $\mathcal{N} \in \text{MC}(k)$ such that $\delta_\lambda(\mathcal{M}, \mathcal{N}) < 1$.

When the distance δ_λ is discounted (i.e., $\lambda < 1$), the two problems above turn out to be trivial because for any MCs \mathcal{M}, \mathcal{N} with initial states labelled with same label, $\delta_\lambda(\mathcal{M}, \mathcal{N}) \leq \lambda$, thus the minimum size for a significant approximant is always 1. In contrast, we show that when the distance δ_λ is undiscounted ($\lambda = 1$) the same two problems are NP-complete. The NP-completeness result is obtained via a characterization of SBA-1 as a combinatorial problem in graph theory on vertex-labelled directed graphs.

A *vertex-labelled directed graph* is a directed graph $G = (V, E)$ with a *vertex labelling function* associating with each vertex in V a label. For a Markov chain $\mathcal{M} = (M, \tau, \ell)$, we denote by $G(\mathcal{M})$ its underlying vertex-labelled directed graph, having M as set of vertices labelled by ℓ and directed edges (m, m') if and only if $\tau(m)(m') > 0$.

Definition 6 (Reflected paths). A path $v_0 \dots v_n$ in a vertex-labelled directed graph G , is *reflected* in a subgraph G' of G if there exists a path $v'_0 \dots v'_n$ in G' such that $v_n = v'_n$ and, for all $0 \leq i \leq n$, v_i and v'_i have the same label.



We show that deciding $\langle \mathcal{M}, k \rangle \in \text{SBA-1}$ is equivalent to checking whether the underlying graph of \mathcal{M} has a bottom strongly connected component (BSCC) and a path reaching it from the initial state satisfying the following condition: the size of the BSCC plus the number of labels occurring in a prefix of such path, obtained by removing the part that is reflected in the BSCC, does not exceed k .

Lemma 10. *Let \mathcal{M} be a minimal MC with initial state m_0 . Then, $\langle \mathcal{M}, k \rangle \in \text{SBA-1}$ iff $G(\mathcal{M})$ has a BSCC $G' = (V, E)$ and a path $m_0 \dots m_h$, such that, for some $0 \leq p \leq h$,*

1. *the path $m_p \dots m_h$ is reflected in G' , and*
2. $|\{\ell(m_i) \mid i < p\}| + |V| \leq k$.

PROOF. (\Rightarrow) By hypothesis there exists $\mathcal{N} = (N, \theta, \alpha) \in \text{MC}(k)$ such that $\delta_1(\mathcal{M}, \mathcal{N}) < 1$. We assume that \mathcal{N} is minimal (otherwise one can replace it with its bisimilarity quotient). By Lemma 12 and Theorem 1, there exists $\mathcal{C} \in \Omega(\mathcal{M}, \mathcal{N})$ such that $\beta_1^{\mathcal{C}}(\mathcal{M}, \mathcal{N}) > 0$. Therefore there exists a path $(m_0, n_0) \dots (m_p, n_p)$ in $G(\mathcal{C})$, such that $\ell(m_i) = \alpha(n_i)$ (for $i = 0..p$) and $\gamma_1^{\mathcal{C}}(m_p, n_p) = 0$.

Note that for arbitrary $m \in M$ and $n \in N$ such that $\gamma_1^{\mathcal{C}}(m, n) = 0$, the following hold

$$0 = \gamma_1^{\mathcal{C}}(m, n) = \sum_{u \in M} \sum_{v \in N} \gamma_1^{\mathcal{C}}(u, v) \cdot \mathcal{C}(m, n)(u, v).$$

Therefore, for any $u \in M$ and $v \in N$ we have that $\mathcal{C}(m, n)(u, v) > 0$ implies $\gamma_1^{\mathcal{C}}(u, v) = 0$.

Let $R \subseteq M \times N$ be the set of states reachable from (m_p, n_p) in $G(\mathcal{C}_1)$. By $\gamma_1^{\mathcal{C}}(m_p, n_p) = 0$ and what have been said before we have that $(m, n) \in R$ implies that $m \sim n$. Let $G = (V, E)$ be a bottom strongly connected component of $G(\mathcal{C})$ such that $V \subseteq R$.

Consider now the graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ where

$$V_1 = \{m \mid (m, n) \in V\} \quad E_1 = \{(m, u) \mid \tau(m)(u) > 0 \text{ and } m, u \in V_1\}, \quad (19)$$

$$V_2 = \{n \mid (m, n) \in V\} \quad E_2 = \{(n, v) \mid \theta(n)(v) > 0 \text{ and } n, v \in V_1\}. \quad (20)$$

Since \mathcal{M} and \mathcal{N} are minimal and $m \sim n$ for all $(m, n) \in V$ we have that for all $(m, n), (u, v) \in V$, $C(m, n)(u, v) = \tau(m)(u) = \theta(n)(v)$. Therefore G_1 and G_2 are bottom strongly connected components of $G(\mathcal{M})$ and $G(\mathcal{N})$ respectively, and are isomorphic with each other.

Let now take the path $(m_0, n_0) \dots (m_h, n_h)$ in $G(\mathcal{C}_1)$ such that $(m_h, n_h) \in V$ obtained by appending the path $(m_0, n_0) \dots (m_p, n_p)$ with the path $(m_p, n_p) \dots (m_h, n_h)$. Note that such a path exists since $(m_h, n_h) \in R$ and, $\ell(m_i) = \alpha(n_i)$ for all $0 \leq i \leq h$.

There are two possible cases:

Case 1: if $n_i \notin V_2$ for all $0 \leq i < h$ we have that the zero-step path m_h is trivially reflected in G_1 and

$$|\{\ell(m_i) \mid i < h\}| + |V_1| = |\{\ell(n_i) \mid i < h\}| + |V_2| \leq |\{n_i \mid i < h\}| + |V_2| \leq |N| \leq k.$$

Case 2: If $n_i \in V_2$ for some $0 \leq i < h$. Let $q < h$ be the smallest index such that $n_q \in V_2$. Since $C(m_q, n_q)(n_{q+1}, n_{q+1}) > 0$ implies $\theta(n_q)(n_{q+1}) > 0$ and G_2 is a bottom strongly connected component, we have that also $n_{q+1} \in V_2$. This shows that $n_q \dots n_h$ is a path in G_2 . Since the isomorphism between G_2 and G_1 preserves the labels (indeed, any $n \in V_2$ is mapped with the unique state $m \in V_1$ such that $m \sim n$) we can see that there exists a path $v_p \dots v_{h-1} m_h$ in G_1 such that $\ell(m_i) = \ell(v_i)$ for all $p \leq i < h$. Therefore we have that the path $m_p \dots m_h$ is reflected in G_1 and

$$|\{\ell(m_i) \mid i < p\}| + |V_1| = |\{\ell(n_i) \mid i < p\}| + |V_2| \leq |\{n_i \mid i < p\}| + |V_2| \leq |N| \leq k.$$

(\Leftarrow) Let $M' = \{m_0, \dots, m_h\}$. Assume w.l.o.g. that $M' \cap V = \{m_h\}$ (otherwise one can consider a prefix of the path that enjoys the assumption). By hypothesis we have that $m_p \dots m_h$ is reflected in G' , therefore there exists a path $v_p \dots v_h$ in G' with $v_h = m_h$ and $\ell(m_j) = \ell(v_j)$ for all $p \leq j \leq h$.

Let $\{l_0, \dots, l_q\} = \{\ell(m_i) \mid i < p\}$ and $N' = \{n_0, \dots, n_q\}$. Consider the Markov chain $\mathcal{N} = (N, \theta, \alpha)$ where $N = N' \cup V$, and

$$\theta(n) = \begin{cases} \tau(n) & \text{if } n \in V \\ \sum_{v \in N} (1/|N|) \cdot 1_v & \text{if } n \in N' \end{cases} \quad \alpha(n) = \begin{cases} \ell(n) & \text{if } n \in V \\ l_i & \text{if } n = n_i \text{ for some } 0 \leq i \leq q \end{cases}$$

Note that θ is well defined because the support of $\tau(v)$ is included in V for all $v \in V$.

By construction, states in N' have pairwise distinct labels, therefore we can define the function $f: M' \rightarrow N$ as $f(m_i) = n$ if $0 \leq i < p$ and $n \in N'$ such that $\alpha(n) = \ell(m_i)$; and $f(m_i) = v_i$ if $p \leq i \leq h$. In the following we will prove that for all $m_i \in M'$, $\delta_1(m_i, f(m_i)) < 1$. We proceed by induction on $r = h - i$.

BASE CASE ($i = h$): One can readily show that $\delta_1(m_h, m_h) = 0$.




Figure 7: (Left) The MC \mathcal{M}_G associated to the graph G in Figure 4 and (right) an MC \mathcal{N} associated to the vertex cover $C = \{2, 3\}$ of G such that $\delta_1(\mathcal{M}_G, \mathcal{N}) < 1$ (cf. Theorem 11).

INDUCTIVE STEP ($i < h$): Let $n = f(m_i)$ and $n' = f(m_{i+1})$ then the following hold

$$\begin{aligned}
\delta_1(m_i, f(m_i)) &= \mathcal{K}(\delta_1)(\tau(m_i), \theta(n)) && (\ell(m_i) = \alpha(n)) \\
&= \min \left\{ \sum_{u \in M} \sum_{v \in N} \omega(u, v) \cdot \delta_1(u, v) \mid \omega \in \Omega(\tau(m_i), \theta(n)) \right\} && (\text{def. } \mathcal{K} \text{ cf. (1)}) \\
&\leq \sum_{u \in M} \sum_{v \in N} \tau(m_i)(u) \cdot \theta(n)(v) \cdot \delta_1(u, v) && (*) \\
&\leq \tau(m_i)(m_{i+1}) \cdot \theta(n)(n') \cdot \delta_1(m_{i+1}, n') + (1 - \tau(m_i)(m_{i+1}) \cdot \theta(n)(n')) && (\delta_1 \sqsubseteq 1) \\
&< 1 && (\delta_1(m_{i+1}, f(m_{i+1})) < 1 \text{ and } \tau(m_i)(m_{i+1}) \cdot \theta(n)(n') > 0)
\end{aligned}$$

Inequality (*) is due to the fact that given $\mu, \nu \in \mathcal{D}(X)$, their ‘parallel synchronisation’ $\mu \odot \nu \in \mathcal{D}(X \times X)$ defined as $(\mu \odot \nu)(x, y) = \mu(x) \cdot \nu(y)$ is a coupling for (μ, ν) .

Therefore $\delta_1(\mathcal{M}, \mathcal{N}) = \delta_1(m_0, f(m_0)) < 1$. By construction $|N| \leq k$, therefore $\langle \mathcal{M}, k \rangle \in \text{SBA-1}$. \square

Theorem 11. *SBA-1 is NP-complete.*

PROOF. The membership in NP is easily proved by using the characterization in Lemma 10 and exploiting Tarjan’s algorithm for generating bottom SCCs. As for the NP-hardness, we provide a polynomial-time many-one reduction from VERTEX COVER. Let $G = (V, E)$ be a graph with $E = \{e_1, \dots, e_n\}$. We construct the MC \mathcal{M}_G as follows. The set of states is given by the set of edges E along with two states e_i^1 and e_i^2 , for each edge $e_i \in E$, representing the two endpoints of e_i and an extra sink state e_0 . The initial state is e_n . The transition probabilities are given as follows. The sink state e_0 loops with probability 1 to itself. Each edge $e_i \in E$ goes with probability $\frac{1}{2}$ to e_i^1 and e_i^2 , respectively. For $1 \leq i \leq n$, the states e_i^1 and e_i^2 go with probability 1 to the state e_{i-1} . The edge states and the sink state are labelled by pairwise distinct labels, while the endpoint states e_i^1 and e_i^2 are labelled by the node in V they represent. An example of construction for \mathcal{M}_G is shown in Figure 7.

Next we show the following equivalence:

$$\langle G, h \rangle \in \text{VERTEX COVER} \quad \text{iff} \quad \langle \mathcal{M}_G, h + n + 1 \rangle \in \text{SBA-1} \quad (21)$$

By construction, \mathcal{M}_G is minimal and its underlying graph H has a unique bottom strongly connected component, namely the self-loop in e_0 . Each path $p = e_n \rightsquigarrow e_0$ in H passes through all edge states, and the set of labels of the endpoint states in p is a vertex cover of G . Since e_0, \dots, e_n have pairwise distinct labels, we have that G has a vertex cover of size at most h iff there exists a path in H from e_n to e_0 that has at most $n + 1 + h$ different labels. Thus, (21) follows by Lemma 10.

Algorithm 1 Approximate Minimization – Expectation Maximization-like heuristic

Input: $\mathcal{M} = (M, \tau, \ell)$, $\mathcal{N}_0 = (N, \theta_0, \alpha)$, and $h \in \mathbb{N}$.

1. $i \leftarrow 0$
 2. **repeat**
 3. $i \leftarrow i + 1$
 4. compute $\mathcal{C} \in \Omega(\mathcal{M}, \mathcal{N}_{i-1})$ such that $\delta_\lambda(\mathcal{M}, \mathcal{N}_{i-1}) = \gamma_\lambda^\mathcal{C}(\mathcal{M}, \mathcal{N}_{i-1})$
 5. $\theta_i \leftarrow \text{UPDATETRANSITION}(\theta_{i-1}, \mathcal{C})$
 6. $\mathcal{N}_i \leftarrow (N, \theta_i, \alpha)$
 7. **until** $\delta_\lambda(\mathcal{M}, \mathcal{N}_i) > \delta_\lambda(\mathcal{M}, \mathcal{N}_{i-1})$ or $i \geq h$
 8. **return** \mathcal{N}_{i-1}
-

6. An Expectation Maximization-like Heuristic

In this section we describe an approximation algorithm for computing near-optimal solutions of CBA- λ for an arbitrary instance $\langle \mathcal{M}, k \rangle$.

Given an initial approximant $\mathcal{N}_0 \in \text{MC}(k)$, the algorithm produces a sequence of MCs $\mathcal{N}_0, \mathcal{N}_1, \dots$ in $\text{MC}(k)$ having successively decreased distance from \mathcal{M} . We defer until later a discussion of how the initial MC \mathcal{N}_0 is chosen. The procedure is described in Algorithm 1.

The intuitive idea of the algorithm is to iteratively update the initial MC by assigning relatively greater probability to transitions that are most representative of the behavior of the MC \mathcal{M} w.r.t. δ_λ . The procedure stops when the last iteration has not yield an improved approximant w.r.t. the preceding one. The input also includes a parameter $h \in \mathbb{N}$ that bounds the number of iterations. Furthermore, to simplify the exposition and avoid computational issues, we assume that $\delta_\lambda(\mathcal{M}, \mathcal{N}_0) < 1$, \mathcal{M} is minimal², $|M| > |N|$ and both τ and θ_0 are rational transition functions.

The rest of the section explains two heuristics used in the UPDATETRANSITION function invoked at line 5. This function shall return the transition probabilities for the successive approximant (see line 6). The two heuristics are both based upon an analysis of the coupling structure $\mathcal{C} \in \Omega(\mathcal{M}, \mathcal{N}_{i-1})$ constructed at line 4.

Define $\beta_\lambda^\mathcal{C}$ to be the least fixed-point of the following functional operator on 1-bounded real-valued functions $d: M \times N \rightarrow [0, 1]$ (ordered point-wise):

$$B_\lambda^\mathcal{C}(d)(m, n) = \begin{cases} 1 & \text{if } \gamma_\lambda^\mathcal{C}(m, n) = 0 \\ 0 & \text{if } \ell(m) \neq \alpha(n) \\ (1 - \lambda) + \lambda \sum_{u \in M, v \in N} d(u, v) \cdot \mathcal{C}(m, n)(u, v) & \text{otherwise.} \end{cases} \quad (22)$$

By Theorem 1, the relation $R_\mathcal{C} = \{(m, n) \mid \gamma_\lambda^\mathcal{C}(m, n) = 0\}$ is easily shown to be a bisimulation, specifically, the greatest bisimulation induced by \mathcal{C} .

Let \mathcal{C}_λ be the MC obtained by augmenting \mathcal{C} with an ‘sink’ state \perp to which any other state moves with probability $(1 - \lambda)$. Intuitively, the value $\beta_\lambda^\mathcal{C}(m, n)$ can be interpreted as the reachability probability in \mathcal{C}_λ of either hitting the sink state or a pair of bisimilar states in $R_\mathcal{C}$ along a path formed only by pairs of states with identical labels starting from (m, n) .

²In case \mathcal{M} is not minimal, one can efficiently replaced it with its bisimilarity quotient [7, 29].

Lemma 12. For all $m \in M$ and $n \in N$, $\beta_\lambda^{\mathcal{C}}(m, n) = 1 - \gamma_\lambda^{\mathcal{C}}(m, n)$.

PROOF. We prove the equivalent statement $\gamma_\lambda^{\mathcal{C}} = \mathbf{1} - \beta_\lambda^{\mathcal{C}}$. Consider the following operator

$$G_\lambda^{\mathcal{C}}(d)(m, n) = \begin{cases} 0 & \text{if } \gamma_\lambda^{\mathcal{C}}(m, n) = 0 \\ 1 & \text{if } \ell(m) \neq \ell(n) \\ \lambda \sum_{u \in M, v \in N} d(u, v) \cdot \mathcal{C}(m, n)(u, v) & \text{otherwise.} \end{cases}$$

One can easily show that $G_\lambda^{\mathcal{C}}$ is monotonic, thus by Knaster-Tarski's fixed-point theorem it admits least and greatest fixed points, say d and d' respectively. We prove that $G_\lambda^{\mathcal{C}}$ has a unique fixed showing that $d = d'$. We proceed by contradiction. Assume that $d \sqsubset d'$.

Consider the set $M = \{(m, n) \mid d'(m, n) - d(m, n) = \|d' - d\|_\infty\}$. By definition of M and $G_\lambda^{\mathcal{C}}$ we have that $(m, n) \in M$ implies that $\gamma_\lambda^{\mathcal{C}}(m, n) \neq 0$ and $\ell(m) = \ell(n)$. Consider $(m, n) \in M$ then we have

$$\begin{aligned} \|d' - d\|_\infty &= d'(m, n) - d(m, n) && \text{(def. } M) \\ &= \lambda \sum_{u \in M, v \in N} (d'(u, v) - d(u, v)) \cdot \mathcal{C}(m, n)(u, v) && \text{(def. } G_\lambda^{\mathcal{C}}) \\ &\leq \lambda \sum_{u \in M, v \in N} \|d' - d\|_\infty \cdot \mathcal{C}(m, n)(u, v) && \text{(def. } \|d' - d\|_\infty) \\ &= \lambda \cdot \|d' - d\|_\infty. && (\mathcal{C}(m, n) \in \mathcal{D}(M \times N)) \end{aligned}$$

There are two possible cases. If $\lambda < 1$ then the above inequality implies that $\|d' - d\|_\infty = 0$ which contradicts $d \neq d'$. If $\lambda = 1$, the above inequality implies that the support of $\mathcal{C}(m, n)$ is included in M . By the generality of m, n , we have that $(m, n) \in M$ implies $\text{support}(\mathcal{C}(m, n)) \subseteq M$, but this implies that $\gamma_\lambda^{\mathcal{C}}(m, n) = 0$, leading to a contradiction.

By definition of $\Gamma_\lambda^{\mathcal{C}}$, it's immediate to see that $\gamma_\lambda^{\mathcal{C}} = G_\lambda^{\mathcal{C}}(\gamma_\lambda^{\mathcal{C}})$. We will complete the proof by showing that also $\mathbf{1} - \beta_\lambda^{\mathcal{C}}$ is a fixed point for $G_\lambda^{\mathcal{C}}$. Let $m \in M$ and $n \in N$, if $\gamma_\lambda^{\mathcal{C}}(m, n) = 0$ or $\ell(m) \neq \alpha(n)$ it trivially holds that $\mathbf{1} - \beta_\lambda^{\mathcal{C}}(m, n) = G_\lambda^{\mathcal{C}}(\mathbf{1} - \beta_\lambda^{\mathcal{C}})(m, n)$. If $\gamma_\lambda^{\mathcal{C}}(m, n) > 0$ and $\ell(m) = \alpha(n)$, then the following equalities hold

$$\begin{aligned} G_\lambda^{\mathcal{C}}(\mathbf{1} - \beta_\lambda^{\mathcal{C}})(m, n) &= \lambda \sum_{u \in M, v \in N} (1 - \beta_\lambda^{\mathcal{C}}(u, v)) \mathcal{C}(m, n)(u, v) && \text{(by def. } G_\lambda^{\mathcal{C}}) \\ &= \lambda - \lambda \sum_{u \in M, v \in N} \beta_\lambda^{\mathcal{C}}(u, v) \cdot \mathcal{C}(m, n)(u, v) && (\mathcal{C}(m, n) \in \mathcal{D}(M \times N)) \\ &= 1 - \beta_\lambda^{\mathcal{C}}(m, n) && (\beta_\lambda^{\mathcal{C}}(m, n) = B_\lambda^{\mathcal{C}}(\beta_\lambda^{\mathcal{C}})(m, n)) \end{aligned}$$

□

From equation (5) and Lemma 12, we can turn the problem CBA- λ as

$$\text{argmax} \{ \beta_\lambda^{\mathcal{C}}(\mathcal{M}, \mathcal{N}) \mid \mathcal{N} \in \text{MC}_{L(\mathcal{M})}(k), \mathcal{C} \in \Omega(\mathcal{M}, \mathcal{N}) \}. \quad (23)$$

Equation (23) says that a solution of CBA- λ is the right marginal of a coupling structure \mathcal{C} such that \mathcal{C}_λ maximizes the probability of generating paths with prefix in $\cong^*(R_{\mathcal{C}} \cup \perp)$ starting from the pair (m_0, n_0) of initial states³, where $\cong = \{(m, n) \notin R_{\mathcal{C}} \mid \ell(m) = \alpha(n)\}$.

In the rest of the section we assume $\mathcal{N}_{i-1} \in \text{MC}(k)$ to be the current approximant with associated coupling structure $\mathcal{C} \in \Omega(\mathcal{M}, \mathcal{N}_{i-1})$ as in line 4 in Algorithm 1.

³We borrowed notation from regular expressions, such as language union, concatenation, and Kleene star, to express the set of finite paths $\cong^*(R_{\mathcal{C}} \cup \perp)$ as a language over the alphabet $(M \times N) \cup \perp$.

The ‘‘Averaged Marginal’’ Heuristic. The first heuristic is inspired by the Expectation Maximization (EM) algorithm described in [23]. The idea is to count the expected number of occurrences of the transitions in \mathcal{C} in the set of paths $\cong^* R_{\mathcal{C}}$ and, in accordance with (23), updating \mathcal{C} by increasing the probability of the transitions that were contributing the most.

For each $m, u \in M$ and $n, v \in N$ let $Z_{u,v}^{m,n}: ((M \times N) \cup \perp)^\omega \rightarrow \mathbb{N}$ be the random variable that counts the number of occurrences of the edge $((m, n)(u, v))$ in a prefix in $\cong^*(R_{\mathcal{C}} \cup \perp)$ of the given path. We denote by $\mathbf{E}[Z_{u,v}^{m,n} | \mathcal{C}]$ the expected value of $Z_{u,v}^{m,n}$ with respect to the probability distribution induced by \mathcal{C}_λ . Using these values we define the optimization problem $\text{EM}\langle \mathcal{N}, \mathcal{C} \rangle$:

$$\begin{aligned} & \text{maximize} && \sum_{m,u \in M} \sum_{n,v \in N} \mathbf{E}[Z_{u,v}^{m,n} | \mathcal{C}] \cdot \ln(c_{u,v}^{m,n}) \\ & \text{such that} && \sum_{v \in N} c_{u,v}^{m,n} = \tau(m)(u) && m, u \in M, n \in N \end{aligned} \quad (24)$$

$$\begin{aligned} & && \sum_{u \in M} c_{u,v}^{m,n} = \theta_{n,v} && m \in M, n, v \in N \\ & && c_{u,v}^{m,n} \geq 0 && m, u \in M, n, v \in N \end{aligned} \quad (25)$$

A solution of $\text{EM}\langle \mathcal{N}, \mathcal{C} \rangle$ describes a Markov chain $\mathcal{N}' = (N, \theta', \alpha)$ and a coupling structure $\mathcal{C}' \in \Omega(\mathcal{M}, \mathcal{N}')$ where, for arbitrary $m \in M$ and $n \in N$

$$\theta'(n) = \sum_{v \in N} \theta_{n,v} \cdot \mathbf{1}_v, \quad \mathcal{C}'(m, n) = \sum_{(u,v) \in M \times N} c_{u,v}^{m,n} \cdot \mathbf{1}_{(u,v)}.$$

The above can be used to improve a pair $\langle \mathcal{N}, \mathcal{C} \rangle$ in the sense of (23).

Theorem 13. *If $\beta_\lambda^{\mathcal{C}}(\mathcal{M}, \mathcal{N}) > 0$, then an optimal solution for $\text{EM}\langle \mathcal{N}, \mathcal{C} \rangle$ describes a Markov chain $\mathcal{N}' \in \text{MC}(k)$ and a coupling structure $\mathcal{C}' \in \Omega(\mathcal{M}, \mathcal{N}')$ satisfying the inequality $\beta_\lambda^{\mathcal{C}'}(\mathcal{M}, \mathcal{N}') \geq \beta_\lambda^{\mathcal{C}}(\mathcal{M}, \mathcal{N})$.*

PROOF. For any measurable set $A \subseteq (M \times N)^\omega$, and $\mathcal{C} \in \Omega(\mathcal{M}, \mathcal{N})$, we denote by $\mathbb{P}_{\mathcal{C}_\lambda}(A)$ the probability that a run of the chain \mathcal{C}_λ belongs to A . To shorten the notation, for $B \subseteq (M \times N)^*$ (resp. $\pi \in (M \times N)^*$) we write $\mathbb{P}_{\mathcal{C}}(B)$ (resp. $\mathbb{P}_{\mathcal{C}}(\pi)$) to indicate $\mathbb{P}_{\mathcal{C}_\lambda}(B(M \times N)^\omega)$ (resp. $\mathbb{P}_{\mathcal{C}_\lambda}(\pi(M \times N)^\omega)$).

Recall that, $\beta_\lambda^{\mathcal{C}}(m_0, n_0) = \mathbb{P}_{\mathcal{C}}(G)$ where $G = (\cong^*(R_{\mathcal{C}} \cup \perp))$ that is the probability that \mathcal{C}_λ generates a path with prefix in $\cong^* R_{\mathcal{C}}$ or $\cong^* \perp$ starting from (m_0, n_0) .

Consider the following inequalities

$$\begin{aligned}
& \ln \beta_\lambda^{\mathcal{C}'}(m_0, n_0) - \ln \beta_\lambda^{\mathcal{C}}(m_0, n_0) \\
&= \ln \frac{\mathbb{P}_{\mathcal{C}'}(G)}{\mathbb{P}_{\mathcal{C}}(G)} \\
&= \ln \frac{\sum_{\pi \in G} \mathbb{P}_{\mathcal{C}'}(G | \pi) \cdot \mathbb{P}_{\mathcal{C}'}(\pi)}{\mathbb{P}_{\mathcal{C}}(G)} \\
&= \ln \sum_{\pi \in G} \frac{\mathbb{P}_{\mathcal{C}'}(G | \pi) \cdot \mathbb{P}_{\mathcal{C}'}(\pi)}{\mathbb{P}_{\mathcal{C}}(G)} \cdot \frac{\mathbb{P}_{\mathcal{C}}(\pi | G)}{\mathbb{P}_{\mathcal{C}}(\pi | G)} \\
&\geq \sum_{\pi \in G} \mathbb{P}_{\mathcal{C}}(\pi | G) \cdot \ln \frac{\mathbb{P}_{\mathcal{C}'}(G | \pi) \cdot \mathbb{P}_{\mathcal{C}'}(\pi)}{\mathbb{P}_{\mathcal{C}}(G) \cdot \mathbb{P}_{\mathcal{C}}(\pi | G)} \quad (\text{by Jensen's inequality}) \\
&= \frac{1}{\beta_\lambda^{\mathcal{C}}(m_0, n_0)} \sum_{\pi \in G} \mathbb{P}_{\mathcal{C}}(\pi) \cdot \ln \frac{\mathbb{P}_{\mathcal{C}'}(\pi)}{\mathbb{P}_{\mathcal{C}}(\pi)} \\
&= \frac{1}{\beta_\lambda^{\mathcal{C}}(m_0, n_0)} (Q' - Q)
\end{aligned}$$

where $Q' = \sum_{\pi \in G} \mathbb{P}_{\mathcal{C}}(\pi) \cdot \ln \mathbb{P}_{\mathcal{C}'}(\pi)$ and $Q = \sum_{\pi \in G} \mathbb{P}_{\mathcal{C}}(\pi) \cdot \ln \mathbb{P}_{\mathcal{C}}(\pi)$. Rearranging we have

$$\beta_\lambda^{\mathcal{C}}(m_0, n_0) \cdot (\ln \beta_\lambda^{\mathcal{C}'}(m_0, n_0) - \ln \beta_\lambda^{\mathcal{C}}(m_0, n_0)) \geq Q' - Q. \quad (26)$$

The logarithm is a strictly increasing function, we have that

$$\ln \beta_\lambda^{\mathcal{C}'}(m_0, n_0) - \ln \beta_\lambda^{\mathcal{C}}(m_0, n_0) \geq 0 \quad \text{iff} \quad \beta_\lambda^{\mathcal{C}'}(m_0, n_0) \geq \beta_\lambda^{\mathcal{C}}(m_0, n_0),$$

therefore by (26) we conclude that $Q' \geq Q$ implies $\beta_\lambda^{\mathcal{C}'}(m_0, n_0) \geq \beta_\lambda^{\mathcal{C}}(m_0, n_0)$.

Thus, inequality 26 suggests that the best choice of \mathcal{C}' is that which maximizes Q' as a function of \mathcal{C}' . Expanding the definition of Q' we obtain

$$Q' = \sum_{\pi \in G} \mathbb{P}_{\mathcal{C}}(\pi) \cdot \ln \mathbb{P}_{\mathcal{C}'}(\pi) = \sum_{\pi \in G} \mathbb{P}_{\mathcal{C}}(\pi) \left(\ln \iota(\pi_0) + \sum_{i=0}^{|\pi|-1} \ln \mathcal{C}'_\lambda(\pi_i)(\pi_{i+1}) \right),$$

where ι is defined as $\iota(x) = 1$ if $x = (m_0, n_0)$; 0 otherwise.

Recall that $Z_{u,v}^{m,n}$ is the random variable that counts the number of occurrences of the edge $((m, n)(u, v))$ in a prefix in G for a given path. Then Q' can be rewritten as

$$\sum_{\pi \in G} \mathbb{P}_{\mathcal{C}}(\pi) \left(\ln(\iota(\pi_0)) + \sum_{m,u \in M} \sum_{n,v \in N} Z_{u,v}^{m,n}(\pi) \ln \mathcal{C}'(m, n)(u, v) \right).$$

Therefore the coupling structure \mathcal{C}' that maximizes the above is obtained as

$$\begin{aligned}
& \operatorname{argmax}_{\mathcal{C}'} \sum_{\pi \in G} \mathbb{P}_{\mathcal{C}}(\pi) \sum_{m,u \in M} \sum_{n,v \in N} Z_{u,v}^{m,n}(\pi) \ln \mathcal{C}'(m, n)(u, v) \quad (\text{eliminating constants}) \\
& \operatorname{argmax}_{\mathcal{C}'} \sum_{m,u \in M} \sum_{n,v \in N} \mathbf{E}[Z_{u,v}^{m,n} | \mathcal{C}] \cdot \ln \mathcal{C}'(m, n)(u, v)
\end{aligned}$$

Since \mathcal{C}' has to range among coupling structures of the form $\mathcal{C}' \in \Omega(\mathcal{M}, \mathcal{N}')$ for some chain \mathcal{N}' with the same states as \mathcal{N} we conclude that an optimal solution of $\text{EM}(\mathcal{N}, \mathcal{C})$ describes a coupling \mathcal{C}' such that $Q' \geq Q$. As above said, this implies $\beta_\lambda^{\mathcal{C}'}(\mathcal{M}, \mathcal{N}') \geq \beta_\lambda^{\mathcal{C}}(\mathcal{M}, \mathcal{N})$. \square

Unfortunately, $\text{EM}(\mathcal{N}, \mathcal{C})$ does not have an easy analytic solution and turns out to be inefficiently solved by nonlinear optimization methods. In contrast, the relaxed optimization problem obtained by dropping the constraints (25) has a simple analytic solution⁴:

$$c_{u,v}^{m,n} = \frac{\tau(m)(n) \cdot \mathbf{E}[Z_{u,v}^{m,n} | \mathcal{C}]}{\sum_{x \in N} \mathbf{E}[Z_{u,x}^{m,n} | \mathcal{C}]},$$

and the first heuristic at line 5, updates θ_i as follows

$$\theta_i(n)(v) = \begin{cases} \theta_{i-1}(n)(v) & \text{if } \exists m \in M. n R_{\mathcal{C}} m \\ \frac{\sum_{m,u \in M} c_{u,v}^{m,n}}{\sum_{x \in N} \sum_{m,u \in M} c_{u,x}^{m,n}} & \text{otherwise.} \end{cases}$$

Recall that the $c_{u,v}^{m,n}$ above may not describe a coupling structure because of the dropping of the constraints (25). Nevertheless we recover the transition probability θ_i , from it as the average of the right marginals.

The ‘‘Averaged Expectations’’ Heuristic. In contrast to the previous case, the second heuristic will update θ_i by directly averaging the expected values of $Z_{u,v}^{m,n}$ as follows

$$\theta_i(n)(v) = \begin{cases} \theta_{i-1}(n)(v) & \text{if } \exists m \in M. n R_{\mathcal{C}} m \\ \frac{\sum_{m,u \in M} \mathbf{E}[Z_{u,v}^{m,n} | \mathcal{C}]}{\sum_{x \in N} \sum_{m,u \in M} \mathbf{E}[Z_{u,x}^{m,n} | \mathcal{C}]} & \text{otherwise.} \end{cases}$$

Computing the Expected Values. We compute $\mathbf{E}[Z_{u,v}^{m,n} | \mathcal{C}]$ using a variant of the *forward-backward* algorithm for hidden Markov models. Let $Z^{m,n}: ((M \times N) \cup \perp)^\omega \rightarrow \mathbb{N}$ be the random variable that counts the number of occurrences of the pair (m, n) in a prefix in $\cong^*(R_{\mathcal{C}} \cup \perp)$ of the path. We compute the expected value of $Z^{m,n}$ w.r.t. the probability induced by \mathcal{C}_λ as the solution $z_{m,n}$ of the following system of equations

$$z_{m,n} = \begin{cases} 0 & \text{if } m \not\cong n \\ \iota(m, n) + \lambda \sum_{u,v} (z_{u,v} + 1) \cdot \mathcal{C}(u, v)(m, n) & \text{otherwise,} \end{cases} \quad (27)$$

where ι is defined as $\iota(x) = 1$ if $x = (m_0, n_0)$; 0 otherwise. Then, the expected value of $Z_{u,v}^{m,n}$ with respect to the probability distribution induced by \mathcal{C}_λ is given by

$$\mathbf{E}[Z_{u,v}^{m,n} | \mathcal{C}] = \lambda \cdot z_{m,n} \cdot \mathcal{C}(m, n)(u, v) \cdot \beta_\lambda^{\mathcal{C}}(u, v). \quad (28)$$

Complexity of the Heuristics. It is worth noting that Algorithm 1 runs in polynomial time in the size of its input. Computing the coupling structure \mathcal{C} in line 4, can be performed in polynomial time in the size of $\mathcal{M} \oplus \mathcal{N}_0$ [18, 30]⁵.

The update step requires one to compute $\mathbf{E}[Z_{u,v}^{m,n} | \mathcal{C}]$ for each $m, u \in M$ and $n, v \in N$. By Equation (28), this can be done by solving two systems of linear equations, namely (27) and (22), each with $O(|M| \cdot |N|)$ unknowns.

⁴By abusing the notation, whenever the nominator is 0, we consider the entire expression equal to 0, regardless of any division by 0. The same convention is used implicitly in the rest of the section.

⁵As pointed out in [31] the method proposed in [18] turns out to be slow in practice. In our implementation of Algorithm 1 the computation of the coupling structure \mathcal{C} in line 4 is performed by using the on-the-fly method proposed in [32]

Choosing the Initial Approximant. Similarly to EM algorithms, the choice of the initial approximant \mathcal{N}_0 may have a significant effect on the quality of the solution.

Notice that the requirement $\delta_\lambda(\mathcal{M}, \mathcal{N}_0) < 1$ is fundamental for both the heuristics, otherwise $\mathbf{E}[Z_{u,v}^{m,n} | \mathcal{C}] = 0$ for all $m, u \in M$ and $n, v \in N$. In case \mathcal{N}_0 does not satisfy the above condition, one should replace it with a better approximant. We have seen that this is trivial to do when $\lambda < 1$, but for $\lambda = 1$ it may be hard to find one (*cf.* Theorem 11), in which case one may need to increase the size of the initial approximant.

For the labeling of the states, one should follow Lemma 4. As for the choice of the underlying structure one shall be guided by Lemma 10. However, due to Theorem 9, it seems unlikely to have generic good strategies for selecting a starting approximant candidate. Nevertheless, good selections for the transition probabilities and the number of states may be suggested by looking at the problem instance.

Finally, by the assumption made on the rationality of the transition functions of \mathcal{M} and \mathcal{N}_0 , Algorithm 1 will always return a chain with rational transition function. In light of Example 1, this means that there are cases where Algorithm 1 will never return an (exact) optimal solution, regardless from the choice of (the rational) initial approximant. Nevertheless, we will see that the we are still able to provide good sub-optimal solutions.

Example 2. Consider the Markov chain \mathcal{M} from Example 1. We show some iterations of Algorithm 1 comparing averaged marginals (AM) and the averaged expectation (AE) heuristics. The tests have been performed starting from two instances of the parametric MC $\mathcal{N}_{x,y}$ described in Example 1.

Test 1				
Heur.	iter	x	y	$\delta_1(\mathcal{M}, \mathcal{N}_{x,y})$
AM	0	$\frac{1}{10}$	$\frac{1}{2}$	$\frac{7629}{10000} \approx 0.76$
	1	$\frac{79}{150}$	$\frac{21}{100}$	$\frac{231233}{421875} \approx \mathbf{0.55}$
	2	$\frac{79}{150}$	$\frac{21}{100}$	$\frac{231233}{421875} \approx 0.55$
AE	0	$\frac{1}{10}$	$\frac{1}{2}$	$\frac{7629}{10000} \approx \mathbf{0.762}$
	1	$\frac{131}{1215}$	$\frac{308}{405}$	$\frac{231233}{421875} \approx 0.763$

Test 2				
Heur.	iter	x	y	$\delta_1(\mathcal{M}, \mathcal{N}_{x,y})$
AM	0	$\frac{1}{10}$	$\frac{1}{8}$	$\frac{1707}{2000} \approx 0.85$
	1	$\frac{79}{150}$	$\frac{21}{100}$	$\frac{231233}{421875} \approx \mathbf{0.55}$
	2	$\frac{79}{150}$	$\frac{21}{100}$	$\frac{231233}{421875} \approx 0.55$
AE	0	$\frac{1}{10}$	$\frac{1}{8}$	$\frac{1707}{2000} \approx 0.85$
	1	$\frac{79}{595}$	$\frac{66}{119}$	$\frac{3171910751}{4212897500} \approx \mathbf{0.752}$
	2	$\frac{38983972}{287248825}$	$\frac{211382913}{287248825}$	$\frac{71603676151221064844103643}{94805770890127757886062500} \approx 0.755$

Table 1: Tests on the accuracy of Algorithm 1 on the Markov chain \mathcal{M} from Example 1.

Table 1 shows how the two heuristics update the initial approximant and how the

distance between \mathcal{M} and the current approximant evolves at each iteration. Notably, the averaged marginals heuristics sets the parameter y to $\frac{21}{100}$ in both tests, reaching the its theoretical optimal value. Regarding the parameter x the heuristics sets its value to $\frac{79}{150}$ reaching an absolute error of 0.23 from its (irrational) theoretical optimal value. As Table 1 shows, in this example, the averaged expectation heuristic suffers from its oversimplified update procedure. It is worth noting how the quality of the outcome may not be influenced by the quality of the initial estimate. \square

7. Experimental Results

We evaluate the performances of Algorithm 1 by comparing the two proposed heuristics on two classical case studies: the *IPv4 zeroconf protocol* from [27, Ex.10.5] (cf. Figure 8) and the *drunkard's walk* (cf. Figure 9).




Figure 8: Markov chain $IPv4(p, q, n)$ of the IPv4 zeroconf protocol for $n = 3$ probes. The parameter $p \in [0, 1]$ models the probability that no reply is received and the next probe is sent, while $q \in [0, 1]$ models the probability that the randomly chosen IP address is already used.




Figure 9: Markov chain $DrkW(p, n, k)$ of the drunkard's walk for $n = 1$ steps to home and $k = 2$ steps to the bar. The parameter $p \in [0, 1]$ models the probability of moving one step toward the bar.

Table 2 shows the results of our tests⁶. The experiments have been performed by running our algorithm on a number of instances $\langle \mathcal{M}, k \rangle$ of increasing size, where \mathcal{M} is an instance of either the IPv4 protocol or the drunkard's walk.

As initial approximant we use a suitably small instance of the same model, with parameters p and q chosen randomly in the open interval $(0, 1)$. For each experiment we report the discount factor λ ; the distance to the original model respectively from \mathcal{N}_0 and \mathcal{N}_h , where h is the total number of iterations; and execution time (in seconds). We compare the two heuristics, averaged marginals (AM) and averaged expectation (AE), on the same initial approximant.

The results obtained on the IPv4 protocol show significant improvements between the initial and the returned approximant. Notably, these are obtained in very few iterations of

⁶ The tests have been made using a prototype implementation coded in Mathematica[®] (available at people.cs.aau.dk/giovbacci/tools.html) running on an Intel Core-i5 2.5 GHz processor with 8GB of DDR3 RAM 1600MHz.

the update procedure. On this model, AM gives approximants of better quality compared with those obtained using AE; however AE seems to be slightly faster than AM. Both the heuristics can handle instances of size up to ~ 100 states. On the drunkard’s walk model, the two heuristics exhibit opposite results w.r.t. the previous experiment: AE provides the best solutions with fewer iterations and significantly lower execution times.

Case	$ M $	k	$\lambda = 1$				$\lambda = 0.8$			
			δ_λ -init	δ_λ -final	#	time	δ_λ -init	δ_λ -final	#	time
IPv4 (AM)	23	5	0.775	0.054	3	4.8	0.576	0.025	3	4.8
	53	5	0.856	0.062	3	25.7	0.667	0.029	3	25.9
	103	5	0.923	0.067	3	116.3	0.734	0.035	3	116.5
	53	6	0.757	0.030	3	39.4	0.544	0.011	3	39.4
	103	6	0.837	0.032	3	183.7	0.624	0.017	3	182.7
	203	6	–	–	–	TO	–	–	–	TO
IPv4 (AE)	23	5	0.775	0.109	2	2.7	0.576	0.049	3	4.2
	53	5	0.856	0.110	2	14.2	0.667	0.049	3	21.8
	103	5	0.923	0.110	2	67.1	0.734	0.049	3	100.4
	53	6	0.757	0.072	2	21.8	0.544	0.019	3	33.0
	103	6	0.837	0.072	2	105.9	0.624	0.019	3	159.5
	203	6	–	–	–	TO	–	–	–	TO
DrkW (AM)	39	7	0.565	0.466	14	259.3	0.432	0.323	14	252.8
	49	7	0.568	0.460	14	453.7	0.433	0.322	14	420.5
	59	8	0.646	–	–	TO	0.423	–	–	TO
DrkW (AE)	39	7	0.565	0.435	11	156.6	0.432	0.321	2	28.6
	49	7	0.568	0.434	10	247.7	0.433	0.316	2	46.2
	59	8	0.646	0.435	10	588.9	0.423	0.309	2	115.7

Table 2: Comparison of the performance of Algorithm 1 on the IPv4 zeroconf protocol and the classic Drunkard’s Walk w.r.t. the heuristics AM and AE.

8. Conclusions and Future Work

To the best of our knowledge, this is the first paper addressing the complexity of the *optimal* approximate minimization of MCs w.r.t. a behavioral metric semantics. Even though for a good evaluation of our heuristics more tests are needed, the current results seem promising. Moreover, in the light of [18, 30], relating the probabilistic bisimilarity distance to the LTL-model checking problem as $\delta_1(\mathcal{M}, \mathcal{N}) \geq |\mathbb{P}_{\mathcal{M}}(\varphi) - \mathbb{P}_{\mathcal{N}}(\varphi)|$, for all $\varphi \in \text{LTL}$, our results might be used to lead saving in the overall model checking time. A deeper study of this topic will be the focus of future work.

We close with an interesting open problem. Membership of BA- λ in NP is left open. However, by arguments analogous to [33, 34] and leveraging on the ideas that made us produce the MC in Example 1, we suspect that BA- λ is hard for the square-root-sum problem. The latter is known to be NP-hard and in PSPACE, but membership in NP

has been open since 1976. Allender et al. [35] showed that it can be decided in the 4th level of the counting hierarchy, thus it is unlikely its PSPACE-completeness.

Furthermore, in light of the relation of the bisimilarity distance δ_λ with the LTL model checking problem (*cf.* Corollary 3) and the PosSLP-hardness of the model checking problem for Interval Markov chains discovered in [23, Theorem 3] one may also consider to study the BA- λ problem as a good candidate for $\text{NP}_{\mathbb{R}}$ completeness in the Blum-Shub-Smale model of computation over the real field with order $(\mathbb{R}, <)$ [36].

Acknowledgements. We thank the anonymous reviewers for their comments on an earlier version of the article. We would also like to show our gratitude to Nathanaël Fijalkow who pointed out a flaw on an earlier version of the proof of Theorem 9.

References

- [1] G. Bacci, G. Bacci, K. G. Larsen, R. Mardare, On the Metric-Based Approximate Minimization of Markov Chains, in: I. Chatzigiannakis, P. Indyk, F. Kuhn, A. Muscholl (Eds.), 44th International Colloquium on Automata, Languages, and Programming (ICALP 2017), Vol. 80 of Leibniz International Proceedings in Informatics (LIPIcs), Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2017, pp. 104:1–104:14. doi:10.4230/LIPIcs.ICALP.2017.104.
- [2] E. F. Moore, Gedanken Experiments on Sequential Machines, in: Automata Studies, Princeton University, 1956, pp. 129–153.
- [3] J. Hopcroft, An $n \log n$ algorithm for minimizing states in a finite automaton, in: Z. Kohavi, A. Paz (Eds.), Theory of Machines and Computations, Academic Press, 1971, pp. 189–196. doi:10.1016/B978-0-12-417750-5.50022-1.
- [4] P. C. Kanellakis, S. A. Smolka, CCS expressions, finite state processes, and three problems of equivalence, in: Proceedings of the 2nd Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing, ACM, 1983, pp. 228–240. doi:10.1145/800221.806724.
- [5] P. C. Kanellakis, S. A. Smolka, CCS expressions, finite state processes, and three problems of equivalence, Information and Computation 86 (1) (1990) 43–68. doi:http://dx.doi.org/10.1016/0890-5401(90)90025-D.
- [6] R. Milner, A Calculus of Communicating Systems, Vol. 92 of Lecture Notes in Computer Science, Springer, 1980. doi:10.1007/3-540-10235-3.
- [7] C. Baier, Polynomial time algorithms for testing probabilistic bisimulation and simulation, in: CAV, Vol. 1102 of Lecture Notes in Computer Science, Springer, 1996, pp. 50–61. doi:10.1007/3-540-61474-5_57.
- [8] K. G. Larsen, A. Skou, Bisimulation through probabilistic testing, Information and Computation 94 (1) (1991) 1–28.
- [9] R. Alur, C. Courcoubetis, N. Halbwachs, D. L. Dill, H. Wong-Toi, Minimization of timed transition systems, in: CONCUR, Vol. 630 of Lecture Notes in Computer Science, Springer, 1992, pp. 340–354. doi:10.1007/BFb0084802.
- [10] M. Yannakakis, D. Lee, An efficient algorithm for minimizing real-time transition systems, Formal Methods in System Design 11 (2) (1997) 113–136. doi:10.1023/A:1008621829508.
- [11] S. Zhang, S. A. Smolka, Towards efficient parallelization of equivalence checking algorithms, in: FORTE, Vol. C-10 of IFIP Transactions, North-Holland, 1992, pp. 121–135.
- [12] S. Blom, S. Orzan, A distributed algorithm for strong bisimulation reduction of state spaces, International Journal on Software Tools for Technology Transfer 7 (1) (2005) 74–86. doi:10.1007/s10009-004-0159-4.
- [13] D. Lee, M. Yannakakis, Online minimization of transition systems (extended abstract), in: Annual ACM Symposium on Theory of Computing, ACM, 1992, pp. 264–274. doi:10.1145/129712.129738.
- [14] C.-C. Jou, S. A. Smolka, Equivalences, congruences, and complete axiomatizations for probabilistic processes, in: CONCUR’90 Theories of Concurrency: Unification and Extension, Vol. 458 of LNCS, 1990, pp. 367–383.
- [15] J. Desharnais, V. Gupta, R. Jagadeesan, P. Panangaden, Metrics for Labeled Markov Systems, in: CONCUR, Vol. 1664 of LNCS, Springer, 1999, pp. 258–273. doi:10.1007/3-540-48320-9_19.
- [16] F. van Breugel, J. Worrell, Towards Quantitative Verification of Probabilistic Transition Systems, in: ICALP, Vol. 2076 of LNCS, 2001, pp. 421–432.

- [17] F. van Breugel, J. Worrell, Approximating and computing behavioural distances in probabilistic transition systems, *Theoretical Computer Science* 360 (3) (2006) 373–385.
- [18] D. Chen, F. van Breugel, J. Worrell, On the Complexity of Computing Probabilistic Bisimilarity, in: *FoSSaCS*, Vol. 7213 of LNCS, Springer, 2012, pp. 437–451.
- [19] N. Ferns, P. Panangaden, D. Precup, *Metrics for finite Markov Decision Processes*, in: UAI, AUAI Press, 2004, pp. 162–169.
- [20] M. Kočvara, M. Stingl, PENNON: A code for convex nonlinear and semidefinite programming, *Optimization Methods and Software* 18 (3) (2003) 317–333. doi:10.1080/1055678031000098773.
- [21] M. Kočvara, M. Stingl, PENBMI 2.0, <http://www.penopt.com/penbmi.html>, accessed: 2016-08-28.
- [22] G. J. McLachlan, T. Krishnan, *The EM Algorithm and Extensions*, 2nd Edition, Wiley-Interscience, 2008.
- [23] M. Benedikt, R. Lenhardt, J. Worrell, LTL Model Checking of Interval Markov Chains, in: *TACAS*, Vol. 7795 of Lecture Notes in Computer Science, Springer, 2013, pp. 32–46. doi:10.1007/978-3-642-36742-7_3.
- [24] G. Franceschinis, R. R. Muntz, Bounds for quasi-lumpable markov chains, *Perform. Eval.* 20 (1-3) (1994) 223–243. doi:10.1016/0166-5316(94)90015-9.
- [25] B. Balle, P. Panangaden, D. Precup, A canonical form for weighted automata and applications to approximate minimization, in: *LICS*, IEEE Computer Society, 2015, pp. 701–712. doi:10.1109/LICS.2015.70.
- [26] J. Desharnais, V. Gupta, R. Jagadeesan, P. Panangaden, Metrics for labelled Markov processes, *Theoretical Computer Science* 318 (3) (2004) 323–354.
- [27] C. Baier, J.-P. Katoen, *Principles of Model Checking*, MIT Press, 2008.
- [28] J. F. Canny, Some Algebraic and Geometric Computations in PSPACE, in: *Proceedings of the 20th Annual ACM Symposium on Theory of Computing (STOC’88)*, ACM, 1988, pp. 460–467. doi:10.1145/62212.62257.
- [29] S. Derisavi, H. Hermanns, W. H. Sanders, Optimal state-space lumping in Markov chains, *Inf. Process. Lett.* 87 (6) (2003) 309–315. doi:10.1016/S0020-0190(03)00343-0.
- [30] G. Bacci, G. Bacci, K. G. Larsen, R. Mardare, Converging from Branching to Linear Metrics on Markov Chains, in: *ICTAC*, Vol. 9399 of LNCS, Springer, 2015, pp. 349–367. doi:10.1007/978-3-319-25150-9_21.
- [31] G. Bacci, G. Bacci, K. G. Larsen, R. Mardare, On-the-Fly Exact Computation of Bisimilarity Distances, in: *TACAS*, Vol. 7795 of LNCS, 2013, pp. 1–15.
- [32] Q. Tang, F. van Breugel, Algorithms to Compute Probabilistic Bisimilarity Distances for Labelled Markov Chains, in: R. Meyer, U. Nestmann (Eds.), *28th International Conference on Concurrency Theory (CONCUR 2017)*, Vol. 85 of Leibniz International Proceedings in Informatics (LIPIcs), Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2017, pp. 27:1–27:16. doi:10.4230/LIPIcs.CONCUR.2017.27.
- [33] T. Chen, S. Kiefer, On the Total Variation Distance of Labelled Markov Chains, in: *CSL-LICS’14*, ACM, 2014, pp. 33:1–33:10. doi:10.1145/2603088.2603099.
- [34] K. Etessami, M. Yannakakis, Recursive Markov chains, stochastic grammars, and monotone systems of nonlinear equations, *J. ACM* 56 (1) (2009) 1:1–1:66. doi:10.1145/1462153.1462154.
- [35] E. Allender, P. Bürgisser, J. Kjeldgaard-Pedersen, P. B. Miltersen, On the complexity of numerical analysis, *SIAM Journal on Computing* 38 (5) (2009) 1987–2006. doi:10.1137/070697926.
- [36] L. Blum, F. Cucker, M. Shub, S. Smale, *Complexity and Real Computation*, Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1998.