

On the Total Variation Distance of Semi-Markov Chains*

Giorgio Bacci, Giovanni Bacci, Kim Guldstrand Larsen, and Radu Mardare

Department of Computer Science, Aalborg University, Denmark
{grbacci, giovbacci, kgl, mardare}@cs.aau.dk

Abstract. Semi-Markov chains (SMCs) are continuous-time probabilistic transition systems where the residence time on states is governed by generic distributions on the positive real line.

This paper shows the *tight* relation between the total variation distance on SMCs and their model checking problem over linear real-time specifications. Specifically, we prove that the total variation between two SMCs *coincides* with the maximal difference w.r.t. the likelihood of satisfying arbitrary MTL formulas or ω -languages recognized by timed automata.

Computing this distance (i.e., solving its threshold problem) is NP-hard and its decidability is an open problem. Nevertheless, we propose an algorithm for approximating it with arbitrary precision.

1 Introduction

The growing interest in quantitative aspects in real world applications motivated the introduction of quantitative models and formal methods for studying their behaviors. Classically, the behavior of two models is compared by means of an equivalence (e.g., bisimilarity, trace equivalence, logical equivalence, etc.). However, when the models depend on numerical values that are subject to error estimates or obtained from statistical samplings, any notion of equivalence is too strong a concept. This motivated the study of *behavioral distances*. The idea is to generalize the concept of equivalence with that of *pseudometric*, aiming at measuring the behavioral dissimilarities between nonequivalent models.

Given a suitably large set of properties Φ , containing all the properties of interest, the behavioral dissimilarities of two states s, s' of a quantitative model are naturally measured by the pseudometric $d(s, s') = \sup_{\phi \in \Phi} |\phi(s) - \phi(s')|$, where $\phi(s)$ denotes the value of ϕ at s . This has been the leading idea for several proposals of behavioral distances, the first one given by Desharnais et al. [12] on probabilistic systems, and further developed by De Alfaro, van Breugel, Worrell, and others [10,11,18,15].

* Work supported by the European Union 7th Framework Programme (FP7/2007-2013) under Grants Agreement nr. 318490 (SENSATION), nr. 601148 (CASSTING) and by the Sino-Danish Basic Research Center IDEA4CPS funded by the Danish National Research Foundation and the National Science Foundation China.

For probabilistic models $\phi(s)$ may represent the probability of satisfaction of a modal formula ϕ measured at s , hence relating the distance d to the *probabilistic model checking problem*. In this context an immediate application is that the probability $\phi(s)$ of satisfying the formula ϕ at s can be approximated by $\phi(s')$ with an error bounded by $d(s, s')$, for any $\phi \in \Phi$. This may lead to savings in the overall cost of model checking.

In this paper we study the total variation distance of probabilistic systems, a popular distance used in a number of domains such as networks security and artificial intelligence, that measures the maximal difference in the probabilities of two systems of realizing the same event. We show that it is a genuine behavioral distance in the above sense by relating it to the probabilistic model checking problem over linear real-time specifications. Specifically, we prove that the total variation distance on semi-Markov chains coincides with the maximal difference in the probability of satisfying the same property, expressed either as an MTL formula [2,3] or an ω -language accepted by a timed automaton (TA) [1].

Semi-Markov chains (SMCs) are continuous-time probabilistic transition systems where the residence time on states is governed by generic distributions on the positive real line. SMCs subsume many probabilistic models, e.g., Markov chains (MCs) and continuous-time Markov Chains (CTMCs). Our attention on linear real-time properties is motivated by applications where the system to be modeled cannot be internally accessed but only tested via observations performed over a set of random executions. For instance, this is mostly common in domains such as systems biology, modeling/testing and machine learning, where real-time features are important e.g. for performance evaluation of cyber-physical systems or dependability analysis.

The total variation distance was already known to be a bound for the maximal difference w.r.t. the probability of satisfying linear-time formulas; our result guarantees that it is the tightest one. Since SMCs and MTL subsume MCs and LTL, respectively, the result holds also in the discrete-time case.

This further motivates the study of efficient methods for computing the total variation. Unfortunately, in [14,9] the threshold problem for the total variation distance is proven to be NP-hard in the case of MCs, and to the best of our knowledge, its decidability is still an open problem. Nevertheless, we prove that the problem of approximating the total variation distance with arbitrary precision is computable. This is done providing two effective sequences that converge from below and above to the total variation distance. This result generalizes that of [9] to the real-time setting. Our approach, however, is different, as it is based on a duality that characterizes the total variation between two measures as the minimal discrepancy associated with their couplings.

The technical contributions of the paper can be summarized as follows.

1. We solved the open problem of how tight is the upper-bound given by the total variation distance w.r.t. the variational distance ranging over MTL formulas and TA specifications, respectively. This has been made possible due to a more general result (Theorem 6) that entails many other nontrivial characterizations of the total variation distance on SMCs.

2. We provided sufficient conditions to construct sequences that converge, from below and above, to the total variation distance. Differently from [9], these conditions are not specific to the probabilistic transition system at hand, but the results hold for probability measures on an arbitrary measurable space.

3. Lastly, we proved the computability of the converging sequences of the previous point. This yields a decidable procedure to approximate the total variation distance with arbitrary precision.

An extended version of the paper containing all the proofs is available at [5].

2 Preliminaries

The set of functions from X to Y is denoted by Y^X and for $f \in Y^X$, let $\equiv_f = \{(x, x') \mid f(x) = f(x')\}$. Given an equivalence relation $R \subseteq X \times X$, X/R denotes the set of R -equivalence classes and $[x]_R$ the equivalence class of $x \in X$.

Measure Theory. A *field* over a set X is a nonempty family $\Sigma \subseteq 2^X$ closed under complement and finite union. Σ is a σ -algebra if, in addition, it is closed under countable union; in this case (X, Σ) is called a *measurable space* and the elements of Σ *measurable sets*. The σ -algebra generated by $\Sigma \subseteq 2^X$, denoted by $\sigma(\Sigma)$, is the smallest σ -algebra containing Σ . Hereafter $(\mathbb{R}_+, \mathbb{B})$ denotes the measurable space of positive real numbers (including zero) with Borel algebra.

Given two measurable spaces (X, Σ) and (Y, Θ) , a function $f: X \rightarrow Y$ is *measurable* if for all $E \in \Theta$, $f^{-1}(E) = \{x \mid f(x) \in E\} \in \Sigma$. The *product space*, $(X, \Sigma) \otimes (Y, \Theta)$, is the measurable space $(X \times Y, \Sigma \otimes \Theta)$, where $\Sigma \otimes \Theta$ is the σ -algebra generated by the *rectangles* $E \times F$ for $E \in \Sigma$ and $F \in \Theta$.

A *measure* on (X, Σ) is a function $\mu: \Sigma \rightarrow \mathbb{R}_+$ s.t. $\mu(\bigcup_{E \in \mathcal{F}} E) = \sum_{E \in \mathcal{F}} \mu(E)$ for all countable families \mathcal{F} of pairwise disjoint measurable sets (σ -*additive*); it is a *probability measure* if, in addition, $\mu(X) = 1$. In what follows $\Delta(X, \Sigma)$ denotes the set of probability measures on (X, Σ) and let $\mathcal{D}(X) = \Delta(X, 2^X)$.

Given a measurable function $f: (X, \Sigma) \rightarrow (Y, \Theta)$, any measure μ on (X, Σ) defines a measure $\mu[f]$ on (Y, Θ) by $\mu[f](E) = \mu(f^{-1}(E))$, for all $E \in \Theta$; it is called the *push forward of μ under f* .

Given μ and ν measures on (X, Σ) and (Y, Θ) , respectively, the *product measure* $\mu \times \nu$ on $(X, \Sigma) \otimes (Y, \Theta)$ is *uniquely* defined by $(\mu \times \nu)(E \times F) = \mu(E) \cdot \nu(F)$, for all $(E, F) \in \Sigma \times \Theta$.

A measure ω on $(X, \Sigma) \otimes (Y, \Theta)$ is a *coupling* for (μ, ν) if for all $E \in \Sigma$ and $F \in \Theta$, $\omega(E \times Y) = \mu(E)$ and $\omega(X \times F) = \nu(F)$ (μ is the *left* and ν the *right marginals* of ω). We denote by $\Omega(\mu, \nu)$ the set of couplings for (μ, ν) .

Metric Spaces. Given a set X , $d: X \times X \rightarrow \mathbb{R}_+$ is a *pseudometric* on X if for arbitrary $x, y, z \in X$, $d(x, x) = 0$, $d(x, y) = d(y, x)$ and $d(x, y) + d(y, z) \geq d(x, z)$; d is a *metric* if, in addition, $d(x, y) = 0$ implies $x = y$. If d is a (pseudo)metric on X , (X, d) is called a (pseudo)metric space.

Given a measurable space (X, Σ) , the set of measures $\Delta(X, \Sigma)$ is metrized by the *total variation distance*, defined by $\|\mu - \nu\| = \sup_{E \in \Sigma} |\mu(E) - \nu(E)|$.

The Space of Timed Paths. A *timed path* over a set X is an infinite sequence $\pi = x_0, t_0, x_1, t_1 \dots$, where $x_i \in X$ and $t_i \in \mathbb{R}_+$; t_i are called *time*




Fig. 1. Two SMCs. (left) the differences are only in the residence time distributions; (right) the behavioral differences arise only from their transition distributions.

delays. For any $i \in \mathbb{N}$, let $\pi[i] = x_i$, $\pi\langle i \rangle = t_i$, $\pi^i = x_0, t_0, \dots, t_{i-1}, x_i$, and $\pi|_i = x_i, t_i, x_{i+1}, t_{i+1}, \dots$. Let $\Pi(X)$ denote the set of timed paths on X .

The *cylinder set* (of rank n) for $X_i \subseteq X$ and $R_i \subseteq \mathbb{R}_+$, $i = 0..n$ is the set $\mathcal{C}(X_0, R_0, \dots, R_{n-1}, X_n) = \{\pi \in \Pi(X) \mid \pi|_n \in X_0 \times R_0 \times \dots \times R_{n-1} \times X_n\}$. For $\mathcal{F} \subseteq 2^X$ and $\mathcal{I} \subseteq 2^{\mathbb{R}_+}$, let $\mathcal{C}_n(\mathcal{F}, \mathcal{I}) = \{\mathcal{C}(X_0, R_0, \dots, R_{n-1}, X_n) \mid X_i \in \mathcal{F}, R_i \in \mathcal{I}\}$, for $n \in \mathbb{N}$, and $\mathcal{C}(\mathcal{F}, \mathcal{I}) = \bigcup_{n \in \mathbb{N}} \mathcal{C}_n(\mathcal{F}, \mathcal{I})$.

If (X, Σ) is a measurable space, $\Pi(X, \Sigma)$ denotes the measurable space of timed paths with σ -algebra generated by $\mathcal{C}(\Sigma, \mathbb{B})$. If $\Sigma = \sigma(\mathcal{F})$ and $\mathbb{B} = \sigma(\mathcal{I})$, then $\sigma(\mathcal{C}(\Sigma, \mathbb{B})) = \sigma(\mathcal{C}(\mathcal{F}, \mathcal{I}))$. Moreover, if both \mathcal{F} and \mathcal{I} are fields, so is $\mathcal{C}(\mathcal{F}, \mathcal{I})$.

Any function $f: X \rightarrow Y$ can be stepwise extended to $f^\omega: \Pi(X) \rightarrow \Pi(Y)$. Note that if f is measurable, so is f^ω .

3 Semi-Markov Chains and Trace Distance

In this section we recall labelled *semi-Markov chains* (SMCs), models that subsume most of the space-finite Markovian models including Markov chains (MCs) and continuous-time Markov chains (CTMCs). We define the total variation distance between SMCs, called *trace distance*, which measures the difference between two SMCs w.r.t. their probabilities of generating labelled timed traces.

In what follows we fix a countable set \mathbb{A} of atomic properties.

Definition 1 (Semi-Markov Chains). A labelled semi-Markov chain is a tuple $\mathcal{M} = (S, \tau, \rho, \ell)$ consisting of a finite set S of states, a transition probability function $\tau: S \rightarrow \mathcal{D}(S)$, a residence-time probability function $\rho: S \rightarrow \Delta(\mathbb{R}_+)$, and a labelling function $\ell: S \rightarrow 2^{\mathbb{A}}$.

In what follows we use $\mathcal{M} = (S, \tau, \rho, \ell)$ to range over the class of SMCs.

Intuitively, if \mathcal{M} is in the state s , it moves to an arbitrary $s' \in S$ within time $t \in \mathbb{R}_+$ with probability $\rho(s)([0, t]) \cdot \tau(s)(s')$. For example, in Fig. 1(right) the SMC moves from s_1 to s_2 before time $t > 0$ with probability $\frac{1}{4} \cdot U[1, 2]([0, t])$, where $U[i, j]$ is the uniform distribution on $[i, j]$. An atomic proposition $p \in \mathbb{A}$ is said to hold in s if $p \in \ell(s)$.

Notice that MCs are the SMCs s.t. for all $s \in S$, $\rho(s)$ is the Dirac measure at 0 (transitions happen instantaneously); while CTMCs are the SMCs s.t. for all $s \in S$, $\rho(s) = \text{Exp}(\lambda)$ —the exponential distribution with rate $\lambda > 0$.

An SMC in an initial state is a stochastic process generating timed paths. They are distributed as in the next definition.

Definition 2. Given $s \in S$ state in \mathcal{M} , let \mathbb{P}_s be the unique probability measure¹ on $\Pi(S)$ such that for all $s_i \in S$ and $R_i \in \mathbb{B}$, $i = 0..n$,

$$\mathbb{P}_s(\mathfrak{C}(\{s_0\}, R_0, \dots, R_{n-1}, \{s_n\})) = \mathbb{1}_{\{s\}}(s_0) \cdot \prod_{i=0}^{n-1} P(s_i, R_i, s_{i+1}),$$

where $\mathbb{1}_A$ is the indicator function of A and $P(u, R, v) = \rho(u)(R) \cdot \tau(u)(v)$.

Since the only things that we observe in a state are the atomic properties (labels), timed paths are considered up to label equivalence. This leads to the definition of *trace cylinders*, which are elements in $\mathfrak{C}(S/\equiv_\epsilon, \mathbb{B})$, and to the following equivalence between states.

Definition 3 (Trace Equivalence). For arbitrary $\mathcal{M} = (S, \tau, \rho, \ell)$, $s, s' \in S$ are trace equivalent, written $s \approx s'$, if for all $T \in \mathfrak{C}(S/\equiv_\epsilon, \mathbb{B})$, $\mathbb{P}_s(T) = \mathbb{P}_{s'}(T)$.

Hereafter, we use \mathcal{T} to denote the set $\mathfrak{C}(S/\equiv_\epsilon, \mathbb{B})$ of trace cylinders.

If two states of an SMCs are *not* trace equivalent, then their difference is usually measured by the total variation distance between their corresponding probabilities restricted to events generated by labelled traces.

Definition 4 (Trace Pseudometric). Given $\mathcal{M} = (S, \tau, \rho, \ell)$, the trace pseudometric $\delta: S \times S \rightarrow [0, 1]$ is defined, for arbitrary $s, s' \in S$, by

$$\delta(s, s') = \sup_{E \in \sigma(\mathcal{T})} |\mathbb{P}_s(E) - \mathbb{P}_{s'}(E)|.$$

It is not difficult to observe that two states $s, s' \in S$ are trace equivalent if and only if $\delta(s, s') = 0$. This demonstrates that the trace equivalence is a *behavioural distance*.

4 Trace Distance and Probabilistic Model Checking

In this section we investigate the connections between the trace distance and model checking SMCs over linear real-time specifications. We show that the variational distance over measurable sets expressed either as Metric Temporal Logic (MTL) formulas or as languages accepted by Timed Automata (TAs) coincides with the trace distance introduced in the previous section. Both these results are instances of a more general result (Theorem 6), which also entails other similar nontrivial characterizations of the trace distance.

A measure μ on (X, Σ) induces the so-called *Fréchet-Nikodym pseudometric* on Σ , $d_\mu: \Sigma \times \Sigma \rightarrow \mathbb{R}_+$ defined for arbitrary $E, F \in \Sigma$, by $d_\mu(E, F) = \mu(E \Delta F)$, where $E \Delta F := (E \setminus F) \cup (F \setminus E)$ is the symmetric difference between sets.

Recall that in a (pseudo)metric space a subset D is dense if its closure \overline{D} (i.e., the set of all the points arbitrarily close to D) coincides with the entire space. In order to prove the aforementioned general result, we need firstly to provide a sufficient condition for a family of measurable sets to be dense w.r.t. the Fréchet-Nikodym pseudometric for some finite measure.

¹ Existence and uniqueness of \mathbb{P}_s is guaranteed by the Hahn-Kolmogorov extension theorem and by the fact that, for all $s \in S$, $\tau(s)$ and $\rho(s)$ are finite measures.

Lemma 5. *Let (X, Σ) be a measurable space and μ be a measure on it. If Σ is generated by a field \mathcal{F} , then \mathcal{F} is dense in the pseudometric space (Σ, d_μ) .*

Proof (sketch). We show that $\overline{\mathcal{F}} := \{E \in \Sigma \mid \forall \varepsilon > 0. \exists F \in \mathcal{F}. d_\mu(E, F) < \varepsilon\} = \Sigma$. To prove $\Sigma \subseteq \overline{\mathcal{F}}$, it is sufficient to show that $\overline{\mathcal{F}}$ is a σ -algebra. The closure under complement follows from $E \triangle F = (X \setminus E) \triangle (X \setminus F)$. The closure under countable union follows from monotonicity, additivity and ω -continuity from below of μ given that for any $\{E_i \mid i \in \mathbb{N}\} \subseteq \overline{\mathcal{F}}$ and $\varepsilon > 0$ the following hold:

- a) there exists $k \in \mathbb{N}$, such that $d_\mu(\bigcup_{i \in \mathbb{N}} E_i, \bigcup_{i=0}^k E_i) < \frac{\varepsilon}{2}$;
 - b) for all $n \in \mathbb{N}$, there exist $F_0, \dots, F_n \in \mathcal{F}$, such that $d_\mu(\bigcup_{i=0}^n E_i, \bigcup_{i=0}^n F_i) < \frac{\varepsilon}{2}$.
- Indeed, by triangular inequality, for arbitrary $F_0, \dots, F_k \in \mathcal{F}$,

$$d_\mu(\bigcup_{i \in \mathbb{N}} E_i, \bigcup_{i=0}^k F_i) \leq d_\mu(\bigcup_{i \in \mathbb{N}} E_i, \bigcup_{i=0}^k E_i) + d_\mu(\bigcup_{i=0}^k E_i, \bigcup_{i=0}^k F_i) < \varepsilon.$$

Then, the lemma follows since \mathcal{F} is a field. \square

With this result in hands we can state the main theorem of this section.

Theorem 6. *Let (X, Σ) be a measurable space and μ, ν be two finite measures on it. If Σ is generated by a field \mathcal{F} , then $\|\mu - \nu\| = \sup_{E \in \mathcal{F}} |\mu(E) - \nu(E)|$.*

Proof. For $Y \neq \emptyset$ and $f: Y \rightarrow \mathbb{R}$ bounded and continuous, if $D \subseteq Y$ is dense then $\sup f(D) = \sup f(Y)$. By Lemma 5, \mathcal{F} is dense in $(\Sigma, d_{\mu+\nu})$. We show that $|\mu - \nu|: \Sigma \rightarrow \mathbb{R}$ is bounded and continuous. Boundedness follows since μ and ν are finite. By monotonicity, positivity, and additivity of the measures one can show that μ and ν are 1-Lipschitz continuous, so $|\mu - \nu|$ is continuous. \square

4.1 Model Checking for MTL Formulas

Metric Temporal Logic [2] has been introduced as a formalism for reasoning on sequences of events in a real-time setting. The grammar of formulas is as follows

$$\varphi ::= p \mid \perp \mid \varphi \rightarrow \varphi \mid \mathbf{X}^{[t, t']} \varphi \mid \varphi \mathbf{U}^{[t, t']} \varphi,$$

where $p \in \mathbb{A}$ and $[t, t']$ are positive-reals intervals with rational boundaries.

The formal semantics² of MTL is given by means of a satisfiability relation defined, for an arbitrary SMC \mathcal{M} and a timed path $\pi \in \Pi(S)$, as follows [16].

$$\begin{array}{ll} \mathcal{M}, \pi \models p & \text{if } p \in \ell(\pi[0]), \\ \mathcal{M}, \pi \models \perp & \text{never,} \\ \mathcal{M}, \pi \models \varphi \rightarrow \psi & \text{if } \mathcal{M}, \pi \models \psi \text{ whenever } \mathcal{M}, \pi \models \varphi, \\ \mathcal{M}, \pi \models \mathbf{X}^{[t, t']} \varphi & \text{if } \pi\langle 0 \rangle \in [t, t'], \text{ and } \mathcal{M}, \pi|_1 \models \varphi, \\ \mathcal{M}, \pi \models \varphi \mathbf{U}^{[t, t']} \psi & \text{if } \exists i > 0 \text{ such that } \sum_{k=0}^{i-1} \pi\langle k \rangle \in [t, t'], \mathcal{M}, \pi|_i \models \psi, \\ & \text{and } \mathcal{M}, \pi|_j \models \varphi \text{ whenever } 0 \leq j < i. \end{array}$$

² This is known as the *point-based* semantics, since the connectives quantify over a countable set of positions in the path; it differs from the *interval-based* semantics, adopted in [7,17], which associates a state with each point in the real line, and let the temporal connectives quantify over intervals with uncountable many points.

Having fixed an SMC \mathcal{M} , let $\llbracket \varphi \rrbracket = \{\pi \mid \mathcal{M}, \pi \models \varphi\}$ and $\llbracket \mathcal{L} \rrbracket = \{\llbracket \varphi \rrbracket \mid \varphi \in \mathcal{L}\}$, for any $\mathcal{L} \subseteq \text{MTL}$. Let MTL^- be the fragment of MTL without until operator.

Lemma 7. (i) $\llbracket \text{MTL} \rrbracket \subseteq \sigma(\mathcal{T})$ and (ii) $\mathcal{T} \subseteq \sigma(\llbracket \text{MTL}^- \rrbracket)$.

Lemma 7 states that (i) MTL formulas describe events in the σ -algebra generated by the trace cylinders; and (ii) the trace cylinders are measurable sets generated by MTL formulas without until operator. Consequently, the probabilistic model checking problem for SMC, which is to determine the probability $\mathbb{P}_s(\llbracket \varphi \rrbracket)$ given the initial state s of \mathcal{M} , is well defined. Moreover, for any $\mathcal{L} \subseteq \text{MTL}$,

$$\delta_{\mathcal{L}}(s, s') = \sup_{\varphi \in \mathcal{L}} |\mathbb{P}_s(\llbracket \varphi \rrbracket) - \mathbb{P}_{s'}(\llbracket \varphi \rrbracket)|$$

is a well-defined pseudometric that distinguishes states w.r.t. their maximal difference in the likelihood of satisfying formulas in \mathcal{L} .

Obviously, the trace distance δ is an upper bound of $\delta_{\mathcal{L}}$; however, Theorem 6 reveals a set of conditions on \mathcal{L} guaranteeing that the two actually coincide.

Corollary 8 (Logical Characterization). *Let \mathcal{L} be a Boolean-closed fragment of MTL s.t. $\mathcal{T} \subseteq \sigma(\llbracket \mathcal{L} \rrbracket)$. Then, $\delta_{\mathcal{L}} = \delta$. In particular, $\delta_{\text{MTL}} = \delta_{\text{MTL}^-} = \delta$.*

Remark 9. The supremum in the definition of δ_{MTL} is not a maximum. Fig.1 shows two examples. The SMC on the right is taken from [9, Example 1]³, where it is proven that $\delta(s_1, s_4)$ has a maximizing event that is not an ω -regular language, hence not describable by an LTL formula. As for the SMC on the left, the maximizing event corresponding to $\delta(u, v)$ should have the form $X^I \top$ for $I = [0, \log(3) - \log(2)]$. However the previous is not an MTL formula since I has an irrational endpoint. ■

4.2 Model Checking for Timed Automata

Timed Automata (TAs) [1] have been introduced to model the behavior of real-time systems over time. Here we consider TAs without location invariants.

Let \mathcal{X} be a finite set of variables (*clocks*) and $\mathcal{V}(\mathcal{X})$ the set of *valuations* $v: \mathcal{X} \rightarrow \mathbb{R}_+$. As usual, for $v \in \mathcal{V}(\mathcal{X})$, $t \in \mathbb{R}_+$ and $X \subseteq \mathcal{X}$, we denote by $\mathbf{0}$ the null valuation, by $v + t$ the t -delay of v and by $v[X := t]$ the update of X in v .

A *clock guard* $g \in \mathcal{G}(\mathcal{X})$ over \mathcal{X} is a finite set of expressions of the form $x \bowtie q$, for $x \in \mathcal{X}$, $q \in \mathbb{Q}_+$ and $\bowtie \in \{<, \leq, >, \geq\}$. We say that a valuation $v \in \mathcal{V}(\mathcal{X})$ *satisfies* a clock guard $g \in \mathcal{G}(\mathcal{X})$, written $v \models g$, if $v(x) \bowtie n$ holds, for all $x \bowtie q \in g$. Two clock guards $g, g' \in \mathcal{G}(\mathcal{X})$ are *orthogonal* (or *non-overlapping*), written $g \perp g'$, if there is no $v \in \mathcal{V}(\mathcal{X})$ such that $v \models g$ and $v \models g'$.

Definition 10 (Timed Automaton). *A timed (Muller) automaton over a set of clocks \mathcal{X} is a tuple $\mathcal{A} = (Q, L, q_0, F, \rightarrow)$ consisting of a finite set Q of locations, a set L of input symbols, an initial location $q_0 \in Q$, a family $F \subseteq 2^Q$ of final sets of locations, and a transition relation $\rightarrow \subseteq Q \times L \times \mathcal{G}(\mathcal{X}) \times 2^{\mathcal{X}} \times Q$.*

\mathcal{A} is deterministic if $(q, a, g, X, q'), (q, a, g', X', q'') \in \rightarrow$ and $g \neq g'$ implies $g \perp g'$; it is resetting if $(q, a, g, X, q') \in \rightarrow$ implies $X = \mathcal{X}$.

³ The SMC has been adapted to the current setting where the labels are in the state, instead of in the transitions.

A run of $\mathcal{A} = (Q, L, q_0, F, \rightarrow)$ over a timed path $\pi = a_0, t_0, a_1, t_1, \dots$ is an infinite sequence

$$(q_0, v_0) \xrightarrow{a_0, t_0} (q_1, v_1) \xrightarrow{a_1, t_1} (q_2, v_2) \xrightarrow{a_2, t_2} \dots$$

with $q_i \in Q$ and $v_i \in \mathcal{V}(\mathcal{X})$ for all $i \geq 0$, satisfying the following requirements: (*initialization*) $v_0 = \mathbf{0}$; (*consecution*) for all $i \geq 0$, exists $(q_i, a_i, g_i, X_i, q_{i+1}) \in \rightarrow$ such that $v_{i+1} = (v_i + t_i)[X_i := 0]$ and $v_i + t_i \models g_i$.

A run over π is *accepting* (π is *accepted* by \mathcal{A}) if the set of locations visited infinitely often is in F . Let $\mathcal{L}(\mathcal{A})$ be the set of timed paths accepted by \mathcal{A} .

A deterministic TA (DTA) has at most one accepting run over a given timed path in $\Pi(L)$. With respect to TAs, which are only closed under finite union and intersection, DTAs are also closed under complement [1].

To relate TAs and SMCs, consider $\mathcal{M} = (S, \tau, \rho, \ell)$ and a TA \mathcal{A} that uses the labels of \mathcal{M} as input symbols. Let $\llbracket \mathcal{A} \rrbracket = \{\pi \mid \ell^\omega(\pi) \in \mathcal{L}(\mathcal{A})\}$ be the set of timed paths in \mathcal{M} accepted by \mathcal{A} and $\llbracket \mathcal{F} \rrbracket = \{\llbracket \mathcal{A} \rrbracket \mid \mathcal{A} \in \mathcal{F}\}$ for any set $\mathcal{F} \in \text{TA}$.

Lemma 11. (i) $\llbracket \text{TA} \rrbracket \subseteq \sigma(\mathcal{T})$ and (ii) $\mathcal{T} \subseteq \sigma(\llbracket \text{DTA} \rrbracket)$.

Lemma 11 states that the model checking problem of an SMC \mathcal{M} against a TA \mathcal{A} , which is to determine the probability $\mathbb{P}_s(\llbracket \mathcal{A} \rrbracket)$ given the initial state s of \mathcal{M} , is well defined and for any $\Phi \subseteq \text{TA}$ we can define the pseudometric

$$\delta_\Phi(s, s') = \sup_{\mathcal{A} \in \Phi} |\mathbb{P}_s(\llbracket \mathcal{A} \rrbracket) - \mathbb{P}_{s'}(\llbracket \mathcal{A} \rrbracket)|$$

that distinguishes states looking at a specific subclass Φ of TA specifications. For a generic $\Phi \subseteq \text{TA}$, the trace distance is an upper bound of δ_Φ . However, Theorem 6 provides conditions that guarantee the equality of the two distances.

Corollary 12. Let $\Phi \subseteq \text{TA}$ be closed under Boolean operations and such that $\mathcal{T} \subseteq \sigma(\llbracket \Phi \rrbracket)$. Then, $\delta_\Phi = \delta$. In particular, $\delta_{\text{TA}} = \delta_{\text{DTA}} = \delta$.

Single-Clock Resetting DTAs. The decidability of model checking CTMCs against TA specifications is open, even for the subclass of DTAs. Recently, Chen et al. [8] provided a decidable algorithm for the case of *single-clock* DTAs (1-DTAs). In this context, an alternative characterization of the trace distance in terms of 1-DTAs is appealing. Notice however that Corollary 12 cannot be applied, since 1-DTAs are not closed under union. We show that the *resetting* 1-DTAs (1-RDTA) satisfy the requirements, hence $\delta_{1\text{-DTA}} = \delta_{1\text{-RDTA}} = \delta$.

Lemma 13. (i) $\llbracket 1\text{-RDTA} \rrbracket$ is a field and (ii) $\mathcal{T} \subseteq \sigma(\llbracket 1\text{-RDTA} \rrbracket)$.

5 General Convergence Criteria

In this section we provide sufficient conditions to construct sequences that converge, from below and from above, to the total variation distance between a generic pair of probability measures. Eventually, we instantiate these results to the specific case of the trace distance on SMCs.

Convergence from Below. To define a converging sequence of under-approximations of the total variation distance we exploit Theorem 6 as follows.

Theorem 14. *Let (X, Σ) be a measurable space and μ, ν be probability measures on it. Let $\mathcal{F}_0 \subseteq \mathcal{F}_1 \subseteq \mathcal{F}_2 \subseteq \dots$ be a sequence s.t. $\mathcal{F} = \bigcup_{i \in \mathbb{N}} \mathcal{F}_i$ is a field that generates Σ and*

$$l_i = \sup \{ |\mu(E) - \nu(E)| \mid E \in \mathcal{F}_i \} .$$

Then, $l_i \leq l_{i+1}$ and $\sup_{i \in \mathbb{N}} l_i = \|\mu - \nu\|$, for all $i \in \mathbb{N}$.

Proof. $l_i \leq l_{i+1}$ follows from $\mathcal{F}_i \subseteq \mathcal{F}_{i+1}$. Because \mathcal{F} is a field s.t. $\sigma(\mathcal{F}) = \Sigma$, μ and ν are finite measures and $\sup_{i \in \mathbb{N}} l_i = \sup_{E \in \mathcal{F}} |\mu(E) - \nu(E)|$, Theorem 6 concludes our proof. \square

According to Theorem 14, to approximate the trace distance δ from below, we just need to find an increasing sequence of collections of measurable sets of timed paths whose union is a field generating $\sigma(\mathcal{T})$. We define it as follows.

For $k \in \mathbb{N}$, let \mathcal{E}_k be the set of all finite unions of cylinders in $\mathfrak{C}_k(S/\equiv_\ell, \mathfrak{R}_k)$, where $\mathfrak{R}_k = \{ [\frac{n}{2^k}, \frac{n+1}{2^k}) \mid 0 \leq n < k2^k \} \cup \{ [k, \infty) \}$. Note that, these cylinders are pairwise disjoint and, in particular, they form a $\sigma(\mathcal{T})$ -measurable partition of $\Pi(S)$. The choice is justified by the following result.

Lemma 15. *For all $k \in \mathbb{N}$, $\mathcal{E}_k \subseteq \mathcal{E}_{k+1}$ and $\bigcup_{k \in \mathbb{N}} \mathcal{E}_k$ is a field generating $\sigma(\mathcal{T})$.*

Given an SMC \mathcal{M} , a sequence of under-approximations of the trace distance δ is given, for $k \in \mathbb{N}$, by $\delta \uparrow_k : S \times S \rightarrow [0, 1]$ defined by

$$\delta \uparrow_k(s, s') = \sup \{ |\mathbb{P}_s(E) - \mathbb{P}_{s'}(E)| \mid E \in \mathcal{E}_k \} . \tag{1}$$

The next result is an immediate consequence of Lemma 15 and Theorem 14.

Corollary 16. *For all $k \in \mathbb{N}$, $\delta \uparrow_k \leq \delta \uparrow_{k+1}$ and $\delta = \sup_{k \in \mathbb{N}} \delta \uparrow_k$.*

Remark 17 (A logical convergence). Note that Theorem 14 suggests alternative constructions of convergent sequences. For example, as lower-approximations of δ one can use the pseudometrics $\delta_{\text{MTL}_k^-}$, where MTL_k^- is the set of MTL^- formulas with *modal depth* at most $k \in \mathbb{N}$. \blacksquare

Convergence from Above. The construction of the converging sequence of over-approximations of the total variation is based on a classic duality result asserting that *the total variation of two measures corresponds to the minimal discrepancy measured among all their possible couplings* [13].

Recall that a coupling $\omega \in \Omega(\mu, \nu)$ for two probability measures μ, ν on (X, Σ) is a measure in the product space $(X, \Sigma) \otimes (X, \Sigma)$ whose left and right marginals are μ and ν , respectively. The *discrepancy* associated with ω is the value $\omega(\not\cong)$, where $\not\cong = \bigcap_{E \in \Sigma} \{ (x, y) \mid x \in E \text{ iff } y \in E \}$ is the *inseparability relation* w.r.t. measurable sets in Σ . Then, the following duality holds.

Lemma 18 ([13, Th.5.2]). *Let μ, ν be probability measures on (X, Σ) . Then, provided that $\not\cong$ is measurable in $\Sigma \otimes \Sigma$, $\|\mu - \nu\| = \min \{ \omega(\not\cong) \mid \omega \in \Omega(\mu, \nu) \}$.*

Given the above result, we can state a second general converging criterion to approach the total variation distance from above.

Theorem 19. *Let (X, Σ) be a measurable space s.t. $\cong \in \Sigma \otimes \Sigma$ and μ, ν be probability measures on it. Let $\Omega_0 \subseteq \Omega_1 \subseteq \Omega_2 \dots$ be an increasing sequence s.t. $\bigcup_{i \in \mathbb{N}} \Omega_i$ is dense in $\Omega(\mu, \nu)$ w.r.t. the total variation distance and define*

$$u_i = \inf \{ \omega(\not\cong) \mid \omega \in \Omega_i \} .$$

Then, $u_i \geq u_{i+1}$ and $\inf_{i \in \mathbb{N}} u_i = \|\mu - \nu\|$, for all $i \in \mathbb{N}$.

Proof. $u_i \geq u_{i+1}$ follows from $\Omega_i \subseteq \Omega_{i+1}$. To prove $\inf_{i \in \mathbb{N}} u_i = \|\mu - \nu\|$, recall that for $Y \neq \emptyset$ and $f: Y \rightarrow \mathbb{R}$ bounded and continuous, if $D \subseteq Y$ is dense then $\inf f(D) = \inf f(Y)$. By hypothesis $\bigcup_{i \in \mathbb{N}} \Omega_i \subseteq \Omega(\mu, \nu)$ is dense; moreover, $\mu \times \nu \in \Omega(\mu, \nu) \neq \emptyset$. We show that $ev_{\not\cong}: \Omega(\mu, \nu) \rightarrow \mathbb{R}$, defined by $ev_{\not\cong}(\omega) = \omega(\not\cong)$ is bounded and continuous. It is bounded since all $\omega \in \Omega(\mu, \nu)$ are finite measures. It is continuous because $\|\omega - \omega'\| \geq |\omega(\not\cong) - \omega'(\not\cong)| = |ev_{\not\cong}(\omega) - ev_{\not\cong}(\omega')|$ (1-Lipschitz continuity). Now, applying Lemma 18, we derive our result. \square

To conclude this section, we define a sequence of sets of couplings that, according to Theorem 19, characterizes the trace distance δ on SMCs.

Observe that the inseparability relation w.r.t. the σ -algebra generated by trace cylinders is measurable and it can be characterized as follows.

Lemma 20. $\equiv_{\ell\omega} = \bigcap_{E \in \sigma(\mathcal{T})} \{ (\pi, \pi') \mid \pi \in E \text{ iff } \pi' \in E \} \in \sigma(\mathcal{T}) \otimes \sigma(\mathcal{T})$.

Next we introduce the notion of *coupling structure* for an SMC. Let $\Pi^k(S) = \{s_0, t_0, \dots, t_{k-1}, s_k \mid s_i \in S, t_i \in \mathbb{R}_+\}$ be the measurable space with σ -algebra generated by $\mathcal{R}_k = \{ \{s_0\} \times R_0 \times \dots \times R_{k-1} \times \{s_k\} \mid s_i \in S, R_i \in \mathbb{B} \}$. Note that, the prefix function $(\cdot)^{|k}: \Pi(S) \rightarrow \Pi^k(S)$ is measurable, hence, the push forward w.r.t. it on $\mu \in \Delta(\Pi(S))$, denoted by $\mu^{|k}$, is a measure in $\Pi^k(S)$.

Definition 21 (Coupling Structure). *A coupling structure of rank $k \in \mathbb{N}$ for an SMC \mathcal{M} is a function $\mathcal{C}: S \times S \rightarrow \Delta(\Pi^k(S) \times \Pi^k(S))$ such that, for all states $s, s' \in S$, $\mathcal{C}(s, s') \in \Omega(\mathbb{P}_s^{|k}, \mathbb{P}_{s'}^{|k})$.*

The set of coupling structures of rank k for \mathcal{M} is denoted by $\mathbb{C}_k(\mathcal{M})$.

A coupling structure of rank k together with a distinguished initial pair of states, can be intuitively seen as a stochastic process generating pairs of timed paths divided in multi-steps of length k and distributed according to the following probability.

Definition 22. *For $k \in \mathbb{N}$, $s, s' \in S$ states in \mathcal{M} and $\mathcal{C} \in \mathbb{C}_k(\mathcal{M})$, let $\mathbb{P}_{s,s'}^{\mathcal{C}}$ be the unique probability measure⁴ on $\Pi(S) \otimes \Pi(S)$ such that, for all $n \in \mathbb{N}$ and $E = \{u_0\} \times R_0 \times \dots \times R_{nk-1} \times \{u_{nk}\}$, $F = \{v_0\} \times H_0 \times \dots \times H_{nk-1} \times \{v_{nk}\} \in \mathcal{R}_{nk}$*

$$\mathbb{P}_{s,s'}^{\mathcal{C}}(\mathfrak{C}(E) \times \mathfrak{C}(F)) = \mathbb{1}_{\{(s,s')\}}(u_0, v_0) \cdot \prod_{h=0}^{n-1} \mathcal{C}(u_{hk}, v_{hk})(E_h \times F_h) ,$$

where $\mathfrak{C}(E)$ denotes the cylinder obtained as the pre-image under $(\cdot)^{|nk}$ of E and $E_h = \{u_{hk}\} \times R_{hk} \times \dots \times R_{(h+1)k-1} \times \{u_{(h+1)k}\}$ (similarly for F).

⁴ The existence and the uniqueness of this measure follow by Hahn-Kolmogorov extension theorem and the fact that any cylinder of rank k can always be represented as a disjoint union of cylinders of rank $k' \geq k$ (see e.g., [6, pp.29–32]).

The name “coupling structure” is justified by the following result.

Lemma 23. *Let \mathcal{C} be a coupling structure for \mathcal{M} , then $\mathbb{P}_{s,s'}^{\mathcal{C}} \in \Omega(\mathbb{P}_s, \mathbb{P}_{s'})$.*

We are finally ready to describe a decreasing sequence that converges to the trace distance on SMCs. Given \mathcal{M} , let $\delta_{\downarrow k}: S \times S \rightarrow [0, 1]$ for $k \in \mathbb{N}$, be

$$\delta_{\downarrow k}(s, s') = \min \{ \mathbb{P}_{s,s'}^{\mathcal{C}}(\neq_{\ell^\omega}) \mid \mathcal{C} \in \mathbb{C}_{2k}(\mathcal{M}) \} . \tag{2}$$

According to Theorem 14 the following suffices to prove the convergence.

Lemma 24. *Let $s, s' \in S$ be a pair of states of an SMC \mathcal{M} . Then,*

- (i) *for all $k \in \mathbb{N}$, $\{ \mathbb{P}_{s,s'}^{\mathcal{C}} \mid \mathcal{C} \in \mathbb{C}_k(\mathcal{M}) \} \subseteq \{ \mathbb{P}_{s,s'}^{\mathcal{C}} \mid \mathcal{C} \in \mathbb{C}_{2k}(\mathcal{M}) \}$;*
- (ii) *$\bigcup_{k \in \mathbb{N}} \{ \mathbb{P}_{s,s'}^{\mathcal{C}} \mid \mathcal{C} \in \mathbb{C}_{2k}(\mathcal{M}) \}$ is dense in $\Omega(\mathbb{P}_s, \mathbb{P}_{s'})$ w.r.t. the total variation.*

Proof (sketch). (i) Let $k > 0$ and $\mathcal{C} \in \mathbb{C}_k(\mathcal{M})$. Define $\mathcal{D}(s, s')$ as the unique measure on $\Pi^{2k}(S) \otimes \Pi^{2k}(S)$ s.t., for all $E = \{u_0\} \times R_0 \times \dots \times R_{2k-1} \times \{u_{2k}\}$ and $F = \{v_0\} \times H_0 \times \dots \times H_{2k-1} \times \{v_{2k}\}$ in \mathcal{R}_{2k}

$$\mathcal{D}(s, s')(E \times F) = \mathcal{C}(s, s')(E' \times F') \cdot \mathcal{C}(u_k, v_k)(E'' \times F''),$$

where $E' = \{u_0\} \times R_0 \times \dots \times R_{k-1} \times \{u_k\}$ and $E'' = \{u_k\} \times R_k \times \dots \times R_{2k-1} \times \{u_{2k}\}$ (similarly for F). One can check that $\mathcal{D} \in \mathbb{C}_{2k}(\mathcal{M})$ and $\mathbb{P}_{s,s'}^{\mathcal{C}} = \mathbb{P}_{s,s'}^{\mathcal{D}}$.

(ii) Let $\Omega = \bigcup_{k \in \mathbb{N}} \{ \mathbb{P}_{s,s'}^{\mathcal{C}} \mid \mathcal{C} \in \mathbb{C}_{2k}(\mathcal{M}) \}$. Let \mathcal{F}_k be the collection of all finite union of sets of the form $\mathfrak{C}(E) \times \mathfrak{C}(F)$, for $E, F \in \mathcal{R}_k$. Note that $\mathcal{F} = \bigcup_{k \in \mathbb{N}} \mathcal{F}_k$ is a field generating the σ -algebra of $\Pi(S) \otimes \Pi(S)$. By Lemma 5 and Definition 22, to prove that Ω is dense it suffices that for all $\mu \in \Omega(\mathbb{P}_s, \mathbb{P}_{s'})$, $k \in \mathbb{N}$ and $F \in \mathcal{F}_k$, there exists $\omega \in \Omega$ s.t. $\omega(F) = \mu(F)$. One can check that $\omega = \mathbb{P}_{s,s'}^{\mathcal{C}}$, where $\mathcal{C} \in \mathbb{C}_{2k}(\mathcal{M})$ is s.t. $\mathcal{C}(s, s') = \mu[|\cdot|^{2k} \times |\cdot|^{2k}]$ (i.e., the push forward of μ along the function $(\pi, \pi') \mapsto (\pi|^{2k}, \pi'|^{2k})$) has the desired property. \square

The following corollary derives from Lemma 24 and Theorem 19.

Corollary 25. *For all $k \in \mathbb{N}$, $\delta_{\downarrow k} \geq \delta_{\downarrow k+1}$ and $\delta = \inf_{k \in \mathbb{N}} \delta_{\downarrow k}$.*

6 An Approximation Algorithm

This section exploits the aforementioned results to propose a decidable procedure for approximating the trace distance δ on SMCs with arbitrary precision.

Let $\varepsilon > 0$ and consider the sequences $\{\delta_{\uparrow k}\}_{k \in \mathbb{N}}$ and $\{\delta_{\downarrow k}\}_{k \in \mathbb{N}}$ from Section 5. The procedure proceeds step-wise (increasing $k \geq 0$) by computing the point-wise difference $\delta_{\downarrow k} - \delta_{\uparrow k}$ until is smaller then ε . Termination and correctness is ensured by the convergence of the sequences from above and below to δ .

Theorem 26. *Let \mathcal{M} be a SMC. There exists an algorithm that, given a rational number $\varepsilon > 0$, computes a function $d: S \times S \rightarrow [0, 1] \cap \mathbb{Q}_+$ such that $|d - \delta| < \varepsilon$.*

We prove this theorem under two reasonable assumptions regarding SMCs:

- A1.** For all $s \in S$ and $q, q' \in \mathbb{Q}_+$, $\rho(s)([q, q'])$ is computable;
- A2.** For all $s, s' \in S$, $\|\rho(s) - \rho(s')\|$ is computable.

In the above $\rho(s)([q, q'])$ and $\|\rho(s) - \rho(s')\|$ may assume real values, and with the term “compute” we mean that there exists an effective Cauchy sequence of rationals that converges to the value.

Lemma 27. *Assuming A1, $\delta \uparrow_k$ is computable for all $k \in \mathbb{N}$.*

Proof (sketch). For each $k \in \mathbb{N}$, the set \mathcal{E}_k is finite. Moreover, for each $s \in S$ and $E \in \mathcal{E}_k$, $\mathbb{P}_s(E)$ is computable thanks to its additivity and the hypothesis A1. \square

The computability of the sequence $\{\delta \downarrow_k\}_{k \in \mathbb{N}}$ is less trivial. Equation (2) suggests to look for a coupling structure $\mathcal{C} \in \mathbb{C}_{2^k}(\mathcal{M})$ that minimizes the discrepancy $\mathbb{P}_{s, s'}^{\mathcal{C}}(\neq_{\ell^\omega})$. This is done by following a searching strategy similar to the one in [4] and structured as follows: (i) we provide an alternative characterization of the discrepancy associated with a coupling structure (Section 6.1); (ii) we describe how to construct an optimal coupling structure and show that its associated discrepancy is computable (Section 6.2).

6.1 Fixed Point Characterization of the Discrepancy

We characterize the discrepancy associated with a coupling structure \mathcal{C} by means of the least fixed point of a suitable operator parametric in \mathcal{C} . To define the fixed point operator it is convenient to split a coupling structure into two “projections”: on discrete state transitions (regardless of time delays); and on residence times (given that a sequence of transitions has occurred). To this end define $\mathbb{S}^k: S \rightarrow \mathcal{D}(S^{k+1})$ and $\mathbb{T}^k: S^k \rightarrow \Delta(\mathbb{R}_+^k)$ as follows

$$\mathbb{S}^k(s)(u_0..u_k) = \mathbb{1}_s(u_0) \cdot \prod_{i=0}^{k-1} \tau(u_i)(u_{i+1}), \quad \mathbb{T}^k(v_1..v_k) = \rho(v_1) \times \dots \times \rho(v_k).$$

Lemma 28. *The set $\mathbb{C}_k(\mathcal{M})$ is in bijection with the set of pairs of functions $\tau_{\mathcal{C}}: S \times S \rightarrow \mathcal{D}(S^{k+1} \times S^{k+1})$ and $\rho_{\mathcal{C}}: S^k \times S^k \rightarrow \Delta(\mathbb{R}_+^k \times \mathbb{R}_+^k)$ such that*

$$\tau_{\mathcal{C}}(u, v) \in \Omega(\mathbb{S}^k(u), \mathbb{S}^k(v)) \quad \text{and} \quad \rho_{\mathcal{C}}(u_1..u_k, v_1..v_k) \in \Omega(\mathbb{T}^k(u_1..u_k), \mathbb{T}^k(v_1..v_k)).$$

Hereafter we identify the coupling structure \mathcal{C} with its bijective image $(\tau_{\mathcal{C}}, \rho_{\mathcal{C}})$.

Intuitively, $\tau_{\mathcal{C}}(u, v)(u_0..u_k, v_0..v_k)$ is the probability that two copies of \mathcal{M} , scheduled according to \mathcal{C} , have respectively generated the sequences of states $u_0..u_k$ and $v_0..v_k$ starting from u and v ; while $\rho(u_0..u_{k-1}, v_0..v_{k-1})(R \times R')$ is the probability that, having observed $u_0..u_{k-1}$ and $v_0..v_{k-1}$, the generated sequence of time delays are in $R, R' \subseteq \mathbb{R}_+^k$, respectively.

For a coupling structure $\mathcal{C} = (\tau_{\mathcal{C}}, \rho_{\mathcal{C}}) \in \mathbb{C}_k(\mathcal{M})$, define the self-map $\Gamma^{\mathcal{C}}$ over $[0, 1]$ -valued functions on $S^{k+1} \times S^{k+1}$ as follows⁵

$$\Gamma^{\mathcal{C}}(d)(u_0..u_k, v_0..v_k) = \begin{cases} 0 & \text{if } \alpha = 0 \\ 1 & \text{if } \alpha \neq 0, \exists i. u_i \neq_{\ell} v_i \\ \beta + (1 - \beta) \cdot \int d \, d\tau_{\mathcal{C}}(u_k, v_k) & \text{otherwise} \end{cases}$$

⁵ Since, for all $u, v \in S$, $\tau_{\mathcal{C}}(u, v)$ is a discrete measure on a finite space, the Lebesgue integral $\int d \, d\tau_{\mathcal{C}}(u, v)$ in the definition of $\Gamma^{\mathcal{C}}$ is $\sum_{x, y \in S^{k+1}} d(x, y) \cdot \tau_{\mathcal{C}}(u, v)(x, y)$.

where $\beta = \rho_{\mathcal{C}}(u_0..u_{k-1}, v_0..v_{k-1})(\neq)$ and $\alpha = \tau_{\mathcal{C}}(u_0, v_0)(u_0..u_k, v_0..v_k)$.

The operator $\Gamma^{\mathcal{C}}$ is monotonic w.r.t. the point-wise order on $[0, 1]$ -valued functions. Hence, by Tarski’s fixed point theorem, $\Gamma^{\mathcal{C}}$ has a least fixed point, which we denote by $\gamma^{\mathcal{C}}$. The next result shows that $\gamma^{\mathcal{C}}$ is closely related to the discrepancy associated with the coupling structure \mathcal{C} , and this will eventually be used to compute it.

Lemma 29. *For any coupling structure \mathcal{C} , $\mathbb{P}_{s, s'}^{\mathcal{C}}(\neq_{\ell\omega}) = \int \gamma^{\mathcal{C}} \, d\tau_{\mathcal{C}}(s, s')$.*

6.2 Construction of an Optimal Coupling Structure

In this subsection we construct an optimal coupling structure by iterating successive updates of a given coupling structure. We provide necessary and sufficient conditions for a coupling structure \mathcal{C} to ensure that $\delta_{\downarrow k}$ is obtained from $\gamma^{\mathcal{C}}$.

To this end, we first introduce the notion of update for a coupling structure.

Definition 30 (Update). *Let $\mathcal{C} = (\tau_{\mathcal{C}}, \rho_{\mathcal{C}}) \in \mathbb{C}_k(\mathcal{M})$. For $\mu \in \Omega(\mathbb{S}^k(u), \mathbb{S}^k(v))$ and $\nu \in \Omega(\mathbb{T}^k(u_1..u_k), \mathbb{T}^k(v_1..v_k))$, define*

- *transition update:* $\mathcal{C}[(u, v)/\mu] = (\tau_{\mathcal{C}}[(u, v) \mapsto \mu], \rho_{\mathcal{C}})$;
- *delay update:* $\mathcal{C}[(u_1..u_k, v_1..v_k)/\nu] = (\tau_{\mathcal{C}}, \rho_{\mathcal{C}}[(u_1..u_k, v_1..v_k) \mapsto \nu])$.

where, for a function $f: X \rightarrow Y$, $f[x \mapsto y]$ denotes the update of f at x with y .

Our update strategy relies on the following result.

Lemma 31 (Update Criteria). *Let $\mathcal{C} = (\tau_{\mathcal{C}}, \rho_{\mathcal{C}}) \in \mathbb{C}_k(\mathcal{M})$ be a coupling structure and $u_0..u_k, v_0..v_k \in S$ such that $\tau_{\mathcal{C}}(u_0..u_k, v_0..v_k) > 0$ and, for all $i \leq k$, $u_i \equiv_{\ell} v_i$. Then, for $\mu \in \Omega(\mathbb{S}^k(u_k), \mathbb{S}^k(v_k))$, $\nu \in \Omega(\mathbb{T}^k(u_0..u_{k-1}), \mathbb{T}^k(v_0..v_{k-1}))$ and $\mathcal{D} = \mathcal{C}[(u_k, v_k)/\mu] \ll (\nu)$, it holds $\gamma^{\mathcal{D}} < \gamma^{\mathcal{C}}$ whenever*

- (i) $\nu(\neq) < \rho_{\mathcal{C}}(u_0..u_{k-1}, v_1..v_{k-1})(\neq)$ and $\int \gamma^{\mathcal{C}} \, d\mu \leq \int \gamma^{\mathcal{C}} \, d\tau_{\mathcal{C}}(u_k, v_k)$, or
- (ii) $\nu(\neq) \leq \rho_{\mathcal{C}}(u_0..u_{k-1}, v_1..v_{k-1})(\neq)$ and $\int \gamma^{\mathcal{C}} \, d\mu < \int \gamma^{\mathcal{C}} \, d\tau_{\mathcal{C}}(u_k, v_k)$.

Condition (i) in Lemma 31 ensures that any $\mathcal{C} = (\tau_{\mathcal{C}}, \rho_{\mathcal{C}}) \in \mathbb{C}_k(\mathcal{M})$ is improved by replacing $\rho_{\mathcal{C}}$ with the function $\rho^*: S^k \times S^k \rightarrow \Delta(\mathbb{R}_+^k \times \mathbb{R}_+^k)$ defined as

$$\begin{aligned} \rho^*(u_0..u_{k-1}, v_1..v_{k-1}) &= \min \{ \nu(\neq) \mid \nu \in \Omega(\mathbb{T}^k(u_0..u_{k-1}), \mathbb{T}^k(v_0..v_{k-1})) \} \\ &= \|\mathbb{T}^k(u_0..u_{k-1}) - \mathbb{T}^k(v_0..v_{k-1})\| \quad (\text{Lemma 18}) \\ &= 1 - \prod_{i=0}^{k-1} (1 - \|\rho(u_i) - \rho(v_i)\|) = \beta^*, \end{aligned}$$

where the last equality follows by the definition of $\mathbb{T}^k(u_0..u_{k-1})$ and $\mathbb{T}^k(v_0..v_{k-1})$ as product measures. Notice that, assuming A2, the above is computable. By replacing β in the definition of $\Gamma^{\mathcal{C}}$ with β^* , $\gamma^{\mathcal{C}}$ can be computed as the least solution of the linear equation system induced by the definition of $\Gamma^{\mathcal{C}}$.

Condition (ii) of Lemma 31 suggests to improve \mathcal{C} with $\mathcal{C}[(u_k, v_k)/\mu^*]$ where

$$\begin{aligned} \mu^* &= \arg \min \{ \int \gamma^{\mathcal{C}} \, d\mu \mid \mu \in \Omega(\mathbb{S}^k(u_k), \mathbb{S}^k(v_k)) \} \\ &= \arg \min \left\{ \sum_{x, y \in S^{k+1}} \gamma^{\mathcal{C}}(x, y) \cdot \mu(x, y) \mid \mu \in \Omega(\mathbb{S}^k(u_k), \mathbb{S}^k(v_k)) \right\}. \end{aligned}$$

The above is a linear program (a.k.a. *transportation problem*), hence computable.

The sufficient conditions for termination is provided by the following lemma.

Lemma 32. *Let $\mathcal{C} = (\tau_{\mathcal{C}}, \rho^*) \in \mathbb{C}_{2^k}(\mathcal{M})$ be such that $\delta_{\downarrow k}(u, v) \neq \int \gamma^{\mathcal{C}} d\tau_{\mathcal{C}}(u, v)$ for some $u, v \in S$. Then there exist $u', v' \in S$ and $\mu \in \Omega(\mathbb{S}^{2^k}(u'), \mathbb{S}^{2^k}(v'))$ such that $\int \gamma^{\mathcal{C}} d\mu < \int \gamma^{\mathcal{C}} d\tau_{\mathcal{C}}(u', v')$.*

Intuitively, the above ensures that, unless \mathcal{C} is an optimal coupling structure, (ii) in Lemma 31 is satisfied, so that, we can further improve \mathcal{C} as aforesaid.

Proposition 33. *Assuming A2, $\delta_{\downarrow k}$ is computable for all $k \in \mathbb{N}$.*

Proof (sketch). The aforementioned strategy ensures that the updated couplings are chosen from the vertices of the polytopes $\Omega(\mathbb{S}^k(u), \mathbb{S}^k(v))$, for $u, v \in S$. Since these polytopes have finitely many vertexes, the procedure eventually terminates. By Lemma 32, the last coupling describes $\delta_{\downarrow k}$. \square

7 Conclusions and Future Work

In this paper we showed that the total variation distance of SMCs (i.e., the trace distance) is the appropriate behavioral distance to reason about linear real-time properties. This has been done by giving characterizations in terms of MTL formulas or timed ω -regular languages that arise naturally in the context of linear real-time probabilistic model checking. Notably, the technique that has been proposed to prove this result is more general and allows for many more interesting characterizations. We showed, for instance, that the distance can be characterized by considering strictly less expressive fragments of MTL, namely MTL^- ; analogously, it suffices to consider only the subclass of ω -languages recognized by single-clock always resetting DTAs.

Moreover, we studied the problem of approximating the trace distance within any absolute error. We showed that the problem is computable by approximating the total variation distance both from above and below by means of the sequences $\{\delta_{\downarrow k}\}_k$ and $\{\delta_{\uparrow k}\}_k$, that are proved to be effective. This both extends the result of [9] to the real-time setting and gives an alternative way to approximate the total variation distance on MCs.

As a future work we consider to further explore the potentiality of the presented results by studying how fast the sequences converge to the total variation distance. Moreover, we would like to see if similar results can be used to link different behavioral distances, such as the Kantorovich-based bisimilarity distance and the total variation (for which the former is know to be an upper bound of the latter), opening for the possibility of “bridging the gap” between trace and branching-based behavioral distances.

From a computational perspective, also motivated by our previous work [4] on MCs, we would like to implement an on-the-fly algorithm for computing tight over-approximations of the trace distance.

References

1. Alur, R., Dill, D.: Automata for Modeling real-time Systems. In: Paterson, M. S. (ed.) ICALP 1990. LNCS, vol. 443, pp. 322–335. Springer, Heidelberg (1990)
2. Alur, R., Henzinger, T.A.: Real-Time Logics: Complexity and Expressiveness. *Information and Computation* 104(1), 35–77 (1993)
3. Alur, R., Henzinger, T.A.: A Really Temporal Logic. *Journal of the ACM* 41(1), 181–204 (1994)
4. Bacci, G., Bacci, G., Larsen, K.G., Mardare, R.: On-the-Fly Exact Computation of Bisimilarity Distances. In: Piterman, N., Smolka, S.A. (eds.) TACAS 2013. LNCS, vol. 7795, pp. 1–15. Springer, Heidelberg (2013)
5. Bacci, G., Bacci, G., Larsen, K.G., Mardare, R.: On the Total Variation Distance of Semi-Markov Chains. Technical report, AAU, DK (2014), http://people.cs.aau.dk/~grbacci/Papers/tvsmc_ext.pdf
6. Billingsley, P.: Probability and Measure, 3rd edn. Wiley, New York (1995)
7. Chen, T., Daciolla, M., Kwiatkowska, M., Mereacre, A.: Time-Bounded Verification of CTMCs against Real-Time Specifications. In: Fahrenberg, U., Tripakis, S. (eds.) FORMATS 2011. LNCS, vol. 6919, pp. 26–42. Springer, Heidelberg (2011)
8. Chen, T., Han, T., Katoen, J.-P., Mereacre, A.: Model checking of continuous-time markov chains against timed automata specifications. *Logical Methods in Computer Science* 7(1) (2011)
9. Chen, T., Kiefer, S.: On the Total Variation Distance of Labelled Markov Chains. In: Proc. of CSL-LICS 2014, p. 33:1–33:10. ACM, New York (2014)
10. de Alfaro, L., Faella, M., Stoelinga, M.: Linear and branching metrics for quantitative transition systems. In: Díaz, J., Karhumäki, J., Lepistö, A., Sannella, D. (eds.) ICALP 2004. LNCS, vol. 3142, pp. 97–109. Springer, Heidelberg (2004)
11. de Alfaro, L., Majumdar, R., Raman, V., Stoelinga, M.: Game Relations and Metrics. In: LICS, pp. 99–108 (July 2007)
12. Desharnais, J., Gupta, V., Jagadeesan, R., Panangaden, P.: Metrics for labelled Markov processes. *Theoretical Computer Science* 318(3), 323–354 (2004)
13. Lindvall, T.: Lectures on the Coupling Method. Wiley Series in Probability and Mathematical Statistics. John Wiley, New York (1992)
14. Lyngsø, R.B., Pedersen, C.N.: The consensus string problem and the complexity of comparing hidden Markov models. *Journal of Computer and System Sciences* 65(3), 545–569 (2002), Special Issue on Computational Biology 2002
15. Mio, M.: Upper-Expectation Bisimilarity and Łukasiewicz μ -Calculus. In: Muscholl, A. (ed.) FOSSACS 2014. LNCS, vol. 8412, pp. 335–350. Springer, Heidelberg (2014)
16. Ouaknine, J., Worrell, J.: On the decidability and complexity of Metric Temporal Logic over finite words. *Logical Methods in Computer Science* 3(8) (2007)
17. Sharma, A., Katoen, J.-P.: Weighted Lumpability on Markov Chains. In: Clarke, E., Virbitskaite, I., Voronkov, A. (eds.) PSI 2011. LNCS, vol. 7162, pp. 322–339. Springer, Heidelberg (2012)
18. van Breugel, F., Worrell, J.: Approximating and computing behavioural distances in probabilistic transition systems. *Theoretical Computer Science* 360(3), 373–385 (2006)