Weighted Lumpability on Markov Chains

Arpit Sharma ⋆ and Joost-Pieter Katoen ⋆⋆

RWTH Aachen University, Software Modeling and Verification Group, Germany
{arpit.sharma,katoen}@cs.rwth-aachen.de

Abstract. This paper reconsiders Bernardo's T-lumpability on continuous-time Markov chains (CTMCs). This notion allows for a more aggressive state-level aggregation than ordinary lumpability. We provide a novel structural definition of (what we refer to as) weighted lumpability, prove some elementary properties, and investigate its compatibility with linear real-time objectives. The main result is that the probability of satisfying a deterministic timed automaton specification coincides for a CTMC and its weigthed lumped analogue. The same holds for metric temporal logic formulas.

Keywords: continuous-time Markov chain, bisimulation, weighted lumpability, deterministic timed automaton, metric temporal logic.

1 Introduction

Continuous-time Markov chains (CTMCs) have a wide applicability ranging from classical performance evaluation to systems biology. Various branching-time relations on CTMCs have been defined such as weak and strong variants of bisimulation equivalence and simulation pre-orders. Strong bisimulation coincides with ordinary lumping equivalence [11]. Their compatibility to (fragments of) stochastic variants of CTL has been thoroughly investigated, cf. [4]. These relations allow for a state-space reduction prior to model checking; in particular, bisimulation minimisation yields considerable reductions and time savings [15] thanks to an efficient minimisation algorithm [13,19].

This paper focuses on a notion of lumpability that allows for a more aggressive state-space aggregation than ordinary lumpability. It originates by Bernardo [6] who considered Markovian testing equivalence over sequential Markovian process calculus (SMPC), and coined the term T-lumpability for the induced state-level aggregation where T stands for testing. His testing equivalence coincides with ready trace equivalence on CTMCs [20], it is a congruence w.r.t. parallel composition, and preserves transient as well as steady-state probabilities [6]. A logical characterisation via a variant of Hennessy-Milner logic has been given in [9,8] establishing the preservation of expected delays. Bernardo defines T-lumpability using four process-algebraic axioms, and alternatively, calls two states T-lumpable if their expected delays w.r.t. any testing process coincide. In this paper, we take a different route and start from a structural definition using first Markov chain principles. As so-called weighted rates are the key to this definition we baptize it weighted lumpability.

⋆ Supported by the European Commission under the India4EU project.
⋆⋆ Supported by the European Commission under the MoVeS project, FP7-ICT-2009-257005.

E. Clarke, I. Virbitskaite, and A. Voronkov (Eds.): PSI 2011, LNCS 7162, pp. 322–339, 2012.

Whereas ordinary lumpability compares states on the basis of their direct successors —the cumulative probability to directly move to any equivalence class must be equal— weighted lumpability (WL, for short) considers a *two-step* perspective. Before explaining the main principle of WL, let us recall that every transition of a CTMC is labeled with a positive real number λ. This parameter indicates the rate of the exponential distribution, i.e., the probability of a λ-labeled transition to be enabled within t time units equals $1 - e^{-\lambda \cdot t}$. In fact, the average residence time in a state is determined as the reciprocal of the sum of the rates of its outgoing transitions. Roughly speaking, two states s and s' are weighted lumpable if for each pair of their direct predecessors the weighted rate to directly move to any equivalence class via the equivalence class $[s] = [s']$ coincides. The main principle is captured in Fig. 1 where $\lambda_{1,1} + \lambda_{1,2} = \lambda_{2,1} + \lambda_{2,2}$, and $\lambda_{C_1} = p_1 \cdot \lambda_{1,1}$, $\lambda_{C_2} = p_1 \cdot \lambda_{1,2} + p_2 \cdot \lambda_{2,1}$, $\lambda_{C_3} = p_2 \cdot \lambda_{2,2}$ with $p_1 = \frac{\lambda_1}{\lambda_1 + \lambda_2}$ and $p_2 = \frac{\lambda_2}{\lambda_1 + \lambda_2}$. Here states s_1 and s_2 are weighted lumpable, as the probability to move from s_0 to all the states in the equivalence class C_i (for i=1, 2, 3) via all the states in $[s_1]$ is equal. This allows for the aggregation of s_1 and s_2, cf. the right CTMC in Fig. 1.

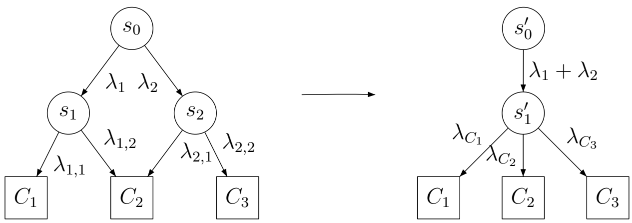

Fig. 1. CTMC aggregation under weighted lumpability

In this paper we define WL as a *structural* notion on CTMCs. We define the quotient under WL, and show that any CTMC is equivalent to its quotient under WL. Our structural definition allows for a simple proof that WL is (strictly) coarser than bisimulation, i.e., ordinary lumpability. Our main focus and motivation, however, is to investigate the preservation of *linear real-time objectives* under WL. We first show that the probability of satisfying a deterministic timed automaton (DTA) [1] specification for any CTMC coincides with that probability for its quotient. This allows for an a priori state-space reduction in linear real-time CTMC model checking [12,5], and implies the preservation of "flat" (i.e., unnested) timed reachability properties and CSL^{TA} formulas [14]. In addition, we study metric temporal logic (MTL) [16], a real-time variant of LTL that is typically used for timed automata (and not for CTMCs). DTA and MTL have incomparable expressiveness [17,3,10]. It is shown that WL-quotienting of CTMCs preserves the probability to satisfy any MTL formula. As a prerequisite result, we show that MTL formulas (interpreted on CTMCs) are measurable.

Organisation of the paper. Section 2 briefly recalls the main concepts of CTMCs. Section 3 defines weighted lumpability and treats some basic properties. Sections 4 and 5 discuss the preservation of DTA properties and MTL-formulas, respectively. Finally, Section 6 concludes the paper. All the proofs are contained in the appendix.

2 Continuous-Time Markov Chains

This section presents the necessary definitions and basic concepts related to continuous-time Markov chains that are needed for the understanding of the rest of this paper.

Definition 1 (CTMC). *A (labeled)* continuous-time Markov chain *(CTMC) is a tuple* $\mathcal{M} = (S, R, AP, L, s_0)$ *where:*

- *S is a non-empty finite set of states,*
- *$R : S \times S \to \mathbb{R}_{\geq 0}$ is a rate function,*
- *AP is a finite set of atomic propositions,*
- *$L : S \to 2^{AP}$ is a labeling function,*
- *$s_0 \in S$ is the initial state.*

The *exit rate* $E(s)$ for state $s \in S$ is defined by $E(s) = \sum_{s' \in S} R(s, s')$. A state s is called *absorbing* iff $E(s) = 0$. The semantics of a CTMC is defined as follows. The probability of moving from s to s' in a single step is defined by $P(s, s') = \frac{R(s,s')}{E(s)}$, if s is non-absorbing and $P(s, s') = 0$ otherwise. The probability to exit state s within t time units is given by $1 - e^{-E(s) \cdot t}$. The probability to move from a non-absorbing state s to s' within t time units equals $P(s, s') \cdot (1 - e^{-E(s) \cdot t})$.

Definition 2 (CTMC timed paths). *Let $\mathcal{M} = (S, R, AP, L, s_0)$ be a CTMC. An infinite* path *π in $\mathcal{M}$ is an alternating sequence of states $s_i \in S$ and time instants $t_i \in \mathbb{R}_{>0}$, i.e., $s_0 \xrightarrow{t_0} s_1 \xrightarrow{t_1} s_2 \cdots s_{n-1} \xrightarrow{t_{n-1}} s_n \cdots$ s.t. $R(s_i, s_{i+1}) > 0$ for all $i \in \mathbb{N}$. A finite path π is an alternating sequence of states $s_i \in S$ and time instants $t_i \in \mathbb{R}_{>0}$, i.e., $s_0 \xrightarrow{t_0} s_1 \xrightarrow{t_1} s_2 \cdots s_{n-1} \xrightarrow{t_{n-1}} s_n$ s.t. $R(s_i, s_{i+1}) > 0$ for all $i < n$.*

Let $Paths^{\mathcal{M}} = Paths^{\mathcal{M}}_{fin} \cup Paths^{\mathcal{M}}_{\omega}$ denote the set of all paths in $\mathcal{M}$, where $Paths^{\mathcal{M}}_{fin} = \bigcup_{n \in \mathbb{N}} Paths^{\mathcal{M}}_n$ is the set of all finite paths in $\mathcal{M}$ and $Paths^{\mathcal{M}}_{\omega}$ is the set of all infinite paths in $\mathcal{M}$. For infinite path $\pi = s_0 \xrightarrow{t_0} s_1 \xrightarrow{t_1} s_2 \cdots s_{n-1} \xrightarrow{t_{n-1}} s_n \cdots$ and any $i \in \mathbb{N}$, let $\pi[i] = s_i$, the $(i+1)$st state of π. Let $\delta(\pi, i) = t_i$ be the time spent in state s_i. For any $t \in \mathbb{R}_{\geq 0}$ and i, the smallest index s.t. $t \leq \sum_{j=0}^{i} t_j$, let $\pi@t = \pi[i]$, the state occupied at time t. For finite path $s_0 \xrightarrow{t_0} s_1 \xrightarrow{t_1} s_2 \cdots s_{n-1} \xrightarrow{t_{n-1}} s_n$, which is either a finite prefix of an infinite path or s_n is absorbing, $\pi[i]$, $\delta(\pi, i)$ are only defined for $i \leq n$, and for $i < n$ defined as in the case of infinite paths. For all $t > \sum_{j=0}^{n-1} t_j$, let $\pi@t = s_n$; otherwise $\pi@t$ is defined as in the case of infinite paths. Let $\delta(\pi, n) = \infty$. Let $\alpha : S \to [0, 1]$, be the intial probability distribution s.t. $\sum_{s \in S} \alpha(s) = 1$. Since $\mathcal{M}$ has a single initial state s_0, $\alpha(s_0) = 1$, and $\forall s \in S$ s.t. $s \neq s_0$, $\alpha(s) = 0$. Let $Paths(s_0)$ denote the set of all paths that start in s_0.

Example 1. Consider the CTMC $\mathcal{M}$ in Fig. 2(a), where $S = \{s_0, s_1, s_2, s_3, s_4, s_5, s_6, s_7\}$, $AP = \{a, b\}$ and s_0 is the initial state. The transition rates are associated with the transitions. An example timed path π is $s_0 \xrightarrow{1.3} s_1 \xrightarrow{1.5} s_3 \xrightarrow{2} s_6$. Here $\pi[3] = s_6$ and $\pi@3 = s_3$.

Definition 3 (Cylinder set). *Let $s_0, \ldots, s_k \in S$ with $P(s_i, s_{i+1}) > 0$ for $0 \leq i < k$ and $I_0, \ldots, I_{k-1}$ be nonempty intervals in $\mathbb{R}_{\geq 0}$. $Cyl(s_0, I_0, \ldots, I_{k-1}, s_k)$ denotes the* cylinder set *consisting of all paths $\pi \in Paths(s_0)$ s.t. $\pi[i] = s_i$ for $i \leq k$, and $\delta(\pi, i) \in I_i$ for $(i < k)$.*

The definition of a Borel space on paths of a CTMC follows [2]. Let $\mathcal{F}(Paths(s_0))$ be the smallest σ-algebra on $Paths(s_0)$ which contains all sets $Cyl(s_0, I_0, \ldots, I_{k-1} s_k)$ s.t. $s_0, \ldots, s_k$ is a state sequence with $P(s_i, s_{i+1}) > 0$ $(0 \leq i < k)$ and $I_0, \ldots, I_{k-1}$ ranges over all sequences of nonempty intervals in $\mathbb{R}_{\geq 0}$.

Definition 4. *The probability measure* $\Pr_\alpha$ *on* $\mathcal{F}(Path(s_0))$ *is the unique measure defined by induction on* k *in the following way. Let* $\Pr_\alpha(Cyl(s_0)) = \alpha(s_0)$ *and for* $k > 0$*:*

$$\Pr_\alpha(Cyl(s_0, I_0, \ldots, s_k, I', s')) = \Pr_\alpha(Cyl(s_0, I_0, \ldots, s_k)) \cdot P(s_k, s', I')$$

where $P(s_k, s', I') = P(s_k, s') \cdot \left(e^{E(s_k)\cdot a} - e^{E(s_k)\cdot b}\right)$ *with* $a = \inf I'$ *and* $b = \sup I'$.

Assumptions. Throughout this paper we assume that every state of CTMC $\mathcal{M}$ has at least one predecessor, i.e., $pred(s) = \{s' \in S \mid P(s', s) > 0\} \neq \varnothing$ for any $s \in S$. This is not a restriction, as any CTMC (S, R, AP, L, s_0) can be transformed into an equivalent CTMC (S', R', AP', L', s_0') which fulfills this condition. This is done by adding a new state $\hat{s}$ to S equipped with a self-loop and which has a transition to each state in S without predecessors. The transition rates for $\hat{s}$ are set to some arbitrary value, e.g., $R(\hat{s}, \hat{s}) = 1$ and $R(\hat{s}, s) = 1$ if $pred(s) = \varnothing$ and 0 otherwise. To distinguish this state from the others we set $L'(\hat{s}) = \bot$ with $\bot \notin$ AP. (All other labels, states and transitions remain unaffected.) Let $s_0' = s_0$. It follows that all states in $S' = S \cup \{\hat{s}\}$ have at least one predecessor. Moreover, the reachable state space of both CTMCs coincides. We also assume that the initial state s_0 of a CTMC is distinguished from all other states by a unique label, say \$. This assumption implies that for any equivalence that groups equally labeled states, $\{s_0\}$ constitutes a seperate equivalence class. Both assumptions do not affect the basic properties of the CTMC such as transient or steady-state distributions. For convenience, we neither show the state $\hat{s}$ nor the label \$ in figures.

3 Weighted Lumpability

Before defining weighted lumpability, we first define two auxiliary concepts. All definitions in this section are relative to a CTMC $\mathcal{M} = (S, R, AP, L, s_0)$. For $C \subseteq S$ and $s \in S$, let $P(s, C) = \sum_{s' \in C} P(s, s')$ be the cumulative probability to directly move from state s to some state in $C \subseteq S$.

Definition 5. *For* $s, s' \in S$ *and* $C \subseteq S$*, the function* $P : S \times S \times 2^S \to \mathbb{R}_{\geq 0}$ *is defined by:*

$$P(s, s', C) = \begin{cases} \frac{P(s,s')}{P(s,C)} & \text{if } s' \in C \text{ and } P(s, C) > 0 \\ 0 & \text{otherwise.} \end{cases}$$

Intuitively, $P(s, s', C)$ is the probability to move from state s to s' under the condition that s moves to some state in C.

Example 2. Consider the example in Fig. 2(a). Let $C = \{s_3, s_4, s_5\}$. Then $P(s_1, s_3, C) = 1/4$, $P(s_1, s_4, C) = 3/4$, $P(s_2, s_4, C) = 3/4$, and $P(s_2, s_5, C) = 1/4$.

Definition 6 (Weighted rate). *For $s \in S$, and $C, D \subseteq S$, the function $wr : S \times 2^S \times 2^S \to \mathbb{R}_{\geq 0}$ is defined by:*

$$wr(s, C, D) = \sum_{s' \in C} P(s, s', C) \cdot R(s', D)$$

where $R(s', D) = \sum_{s'' \in D} R(s', s'')$.

Intuitively, $wr(s, C, D)$ is the (weighted) rate to move from s to some state in D in two steps via any state $s' \in C$. Since $P(s', D) = \frac{R(s',D)}{E(s')}$, $wr(s, C, D)$ equals $\sum_{s' \in C} P(s, s', C) \cdot P(s', D) \cdot E(s')$.

Example 3. Consider the example in Fig. 2(a). Let $D = \{s_6\}$. Then $wr(s_1, C, D) = P(s_1, s_3, C) \cdot R(s_3, D) + P(s_1, s_4, C) \cdot R(s_4, D) = \frac{1}{2}$, $wr(s_2, C, D) = P(s_2, s_4, C) \cdot R(s_4, D) + P(s_2, s_5, C) \cdot R(s_5, D) = \frac{1}{2}$. Similarly, for $D = \{s_7\}$, we get $wr(s_1, C, D) = P(s_1, s_3, C) \cdot R(s_3, D) + P(s_1, s_4, C) \cdot R(s_4, D) = \frac{3}{2}$, $wr(s_2, C, D) = P(s_2, s_4, C) \cdot R(s_4, D) + P(s_2, s_5, C) \cdot R(s_5, D) = \frac{3}{2}$.

The above ingredients allow for the following definition of weighted lumpability, the central notion in this paper. For $C \subseteq S$, let $pred(C) = \bigcup_{s \in C} pred(s)$.

Definition 7 (WL). *Equivalence $\mathcal{R}$ on S is a* weighted lumping *(WL) if we have:*

1. *$\forall (s_1, s_2) \in \mathcal{R}$ it holds: $L(s_1) = L(s_2)$ and $E(s_1) = E(s_2)$, and*
2. *$\forall C, D \in S/_{\mathcal{R}}$ and $\forall s', s'' \in pred(C)$ it holds: $wr(s', C, D) = wr(s'', C, D)$.*

States s_1, s_2 are weighted lumpable*, denoted by $s_1 \cong s_2$, if $(s_1, s_2) \in \mathcal{R}$ for some WL $\mathcal{R}$.*

The first condition asserts that s_1 and s_2 are equally labeled and have identical exit rates. The second condition requires that for any two equivalence classes $C, D \in S/_{\mathcal{R}}$, where $S/_{\mathcal{R}}$ denotes the set consisting of all $\mathcal{R}$-equivalence classes, the weighted rate of going from any two predecessors of C to D via any state in C must be equal. Note that, by definition, any WL is an equivalence relation. Weighted lumpability coincides with Bernardo's notion of T-lumpability [6,7] that is defined in an axiomatic manner for action-labeled CTMCs. Roughly speaking, two states are T-lumpable if their expected delays w.r.t. to any test process, put in parallel to the CTMC, coincide for both the states.

Example 4. For the CTMC in Fig. 2(a), the equivalence relation induced by the partitioning $\{\{s_0\}, \{s_1\}, \{s_2\}, \{s_3, s_4, s_5\}, \{s_6\}, \{s_7\}\}$ is a WL relation.

Definition 8 (Quotient CTMC). *For WL relation $\mathcal{R}$ on $\mathcal{M}$, the quotient CTMC $\mathcal{M}/_{\mathcal{R}}$ is defined by $\mathcal{M}/_{\mathcal{R}} = (S/_{\mathcal{R}}, R', AP, L', s'_0)$ where:*

- *$S/_{\mathcal{R}}$ is the set of all equivalence classes under $\mathcal{R}$,*
- *$R'(C, D) = wr(s', C, D)$ where $C, D \in S/_{\mathcal{R}}$ and $s' \in pred(C)$,*
- *$L'(C) = L(s)$, where $s \in C$ and*
- *$s'_0 = C$ where $s_0 \in C$.*

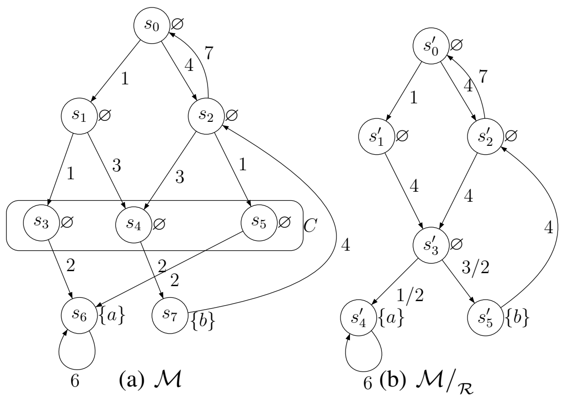

Fig. 2. (a) A CTMC and (b) its quotient under weighted lumpability

Note that $R'(C, D)$ is well-defined as for any predecessors s', s'' of C it follows $wr(s', C, D) = wr(s'', C, D)$. Similarly, L' is well-defined as states in any equivalence class C are equally labeled.

Example 5. The quotient CTMC for the Fig. 2(a) under the WL relation with partition $\{\{s_0\}, \{s_1\}, \{s_2\}, \{s_3, s_4, s_5\}, \{s_6\}, \{s_7\}\}$ is shown in Fig. 2(b).

Next, we show that any CTMC $\mathcal{M}$ and its quotient under WL relation are $\cong$-equivalent.

Definition 9. *For WL $\mathcal{R}$ on $\mathcal{M}$, let $\mathcal{M} \cong \mathcal{M}/_\mathcal{R}$ iff $\forall C \in S/_\mathcal{R}, s \in C$ it holds $s \cong C$.*

Theorem 1. *Let $\mathcal{M}$ be a CTMC and $\mathcal{R}$ be a WL on $\mathcal{M}$. Then $\mathcal{M} \cong \mathcal{M}/_\mathcal{R}$.*

Remark 1. The notion of WL-equivalent states cannot be lifted to WL equivalent time-abstract paths. That is to say, in general, $s \cong s'$ does not imply that for every path π_1 of s, there exists a statewise $\cong$-equivalent path π_2 in s', i.e., $\pi_1 \cong \pi_2$ iff $s_{i,1} \cong s_{i,2}$ for all $i \geq 0$. Consider, e.g., Fig. 2(a). Here $s_3 \cong s_4$, but the time-abstract paths from s_3, s_4 are not $\cong$-equivalent, as s_3 can move to s_6 but there is no direct successor s of s_4 with $s \cong s_6$. (Note that $L(s_6) \neq L(s_7)$.) As a consequence, $\cong$ is not finer than probabilistic trace equivalence [20].

To conclude this section, we investigate the relationship of WL to bisimulation, i.e., ordinary lumping [4,11]. This relationship is not novel; it is also given for T-lumpability in [6], but its proof is now quite simple thanks to the simplicity of the definition of WL.

Definition 10 (Bisimulation [4]). *Equivalence $\mathcal{R}$ on S is a bisimulation on $\mathcal{M}$ if for any $(s_1, s_2) \in \mathcal{R}$ we have: $L(s_1) = L(s_2)$, and $R(s_1, C) = R(s_2, C)$ for all C in $S/_\mathcal{R}$. s_1 and s_2 are bisimilar, denoted $s_1 \sim s_2$, if $(s_1, s_2) \in \mathcal{R}$ for some bisimulation $\mathcal{R}$.*

These conditions require that any two bisimilar states are equally labeled and have identical cumulative rates to move to any equivalence class C. Note that as $R(s, C) = P(s, C) \cdot E(s)$, the condition on the cumulative rates can be reformulated as $P(s_1, C) = P(s_2, C)$ for all $C \in S/_\mathcal{R}$ and $E(s_1) = E(s_2)$.

Lemma 1. *$\sim$ is strictly finer than $\cong$.*

The proof that $\sim$ is finer than $\cong$ is in the appendix. It follows from Fig. 2 that $s_1 \cong s_2 \not\Rightarrow s_1 \sim s_2$. Consider the equivalence class $C = \{s_3, s_4, s_5\}$ in Fig. 2(a). Here $s_3 \not\sim s_4$ since s_3 can reach an a-state while s_4 cannot.

4 Preservation of DTA Specifications

Bisimulation equivalence coincides with the logical equivalence of the branching-time logic CSL [4], a probabilistic real-time variant of CTL [2]. This implies that bisimilar states satisfy the same CSL formulas, a property that—thanks to efficient minimisation algorithms [13]— is exploited by model checkers to minimise the state space prior to verification. In order to investigate the kind of real-time properties for CTMCs that are preserved by WL, we study in this section *linear* real-time objectives that are given by Deterministic Timed Automata (DTA) [1]. These include, e.g., properties of the form: what is the probability to reach a given target state within the deadline, while avoiding "forbidden" states and not staying too long in any of the "dangerous" states on the way. Such properties can neither be expressed in CSL nor in dialects thereof [14]. A model-checking algorithm that verifies a CTMC against a DTA specification has recently been developed [12]; first experimental results are provided in [5]. The key issue is to compute the probability of all CTMC paths that are accepted by a DTA. In this section, we will deal with finite acceptance conditions, i.e., a DTA accepts the timed path if one of its final locations is reached. The results, however, also carry over to Muller acceptance conditions.

Deterministic timed automata. A DTA is a finite-state automaton equipped with a finite set of real-valued variables, called *clocks*. Clocks increase implicitly, all at the same pace, they can be inspected (in guards) and can be reset to the value zero. Let $\mathcal{X}$ be a finite set of clocks ranged over by x and y. A *clock constraint* g over set $\mathcal{X}$ is either of the form $x \bowtie c$ with $c \in \mathbb{N}$ and $\bowtie \in \{<, \leq, >, \geq\}$, or of the form $x - y \bowtie c$, or a conjunction of clock constraints. Let $\mathcal{CC}(\mathcal{X})$ denote the set of clock constraints over $\mathcal{X}$.

Definition 11 (DTA). *A* deterministic timed automaton *(DTA) is a tuple $\mathcal{A} = (\Sigma, \mathcal{X}, Q, q_0, F, \rightarrow)$ where:*

- *Σ is a finite alphabet,*
- *$\mathcal{X}$ is a finite set of clocks,*
- *Q is a nonempty finite set of locations with the initial location $q_0 \in Q$,*
- *$F \subseteq Q$ is a set of accepting (or final) locations,*
- *$\rightarrow \subseteq Q \times \Sigma \times \mathcal{CC}(\mathcal{X}) \times 2^{\mathcal{X}} \times Q$ is the edge relation satisfying:*

$$\left(q \xrightarrow{a,g,X} q' \text{ and } q \xrightarrow{a,g',X'} q'' \text{ with } g \neq g'\right) \quad \text{implies} \quad g \cap g' = \varnothing.$$

Intuitively, the edge $q \xrightarrow{a,g,X} q'$ asserts that the DTA $\mathcal{A}$ can move from location q to q' when the input symbol is a and the guard g holds, while the clocks in X should be reset when entering q' (all other clocks keep their value). DTA are deterministic as they have a single initial location, and outgoing edges of a location labeled with the same input symbol are required to have disjoint guards. In this way, the next location is uniquely determined for a given location and a given set of clock values. In case no guard is satisfied in a location for a given clock valuation, time can progress. If the advance of time will never reach a situation in which a guard holds, the DTA will stay in that location *ad infinitum*. Note that DTA do not have location invariants.

The semantics of a DTA is given by an infinite-state transition system. We do not provide the full semantics, cf. [1], but we define the notion of paths, i.e., runs or executions of a DTA. This is done using some auxiliary notions. A *clock valuation* η for a set $\mathcal{X}$ of clocks is a function $\eta : \mathcal{X} \to \mathbb{R}_{\geq 0}$, assigning to each clock $x \in \mathcal{X}$ its current value $\eta(x)$. The clock valuation η over $\mathcal{X}$ satisfies the clock constraint g, denoted $\eta \models g$, iff the values of the clocks under η fulfill g. For instance, $\eta \models x - y > c$ iff $\eta(x) - \eta(y) > c$. Other cases are defined analogously. For $d \in \mathbb{R}_{\geq 0}$, $\eta{+}d$ denotes the clock valuation where all clocks of η are increased by d. That is, $(\eta{+}d)(x) = \eta(x){+}d$ for all clocks $x \in \mathcal{X}$. Clock *reset* for a subset $X \subseteq \mathcal{X}$, denoted by $\eta[X := 0]$, is the valuation η' defined by: $\forall x \in X.\eta'(x) := 0$ and $\forall x \notin X.\eta'(x) := \eta(x)$. We denote the valuation that assigns 0 to all the clocks by $\mathbf{0}$. An (infinite) path of DTA $\mathcal{A}$ has the form $\rho = q_0 \xrightarrow{a_0,t_0} q_1 \xrightarrow{a_1,t_1} \ldots$ such that $\eta_0 = \mathbf{0}$, and for all $j \geqslant 0$, it holds $t_j > 0$, $\eta_j{+}t_j \models g_j$, $\eta_{j+1} = (\eta_j{+}t_j)[X_j := 0]$, where η_j is the clock evaluation on entering q_j. Here, g_j is the guard of the j-th edge taken in the DTA and X_j the set of clock to be reset on that edge. A path ρ is accepted by $\mathcal{A}$ if $q_i \in F$ for some $i \geqslant 0$. Since the DTA is deterministic, the successor location is uniquely determined; for convenience we write $q' = \mathit{succ}(q, a, g)$. A path in a CTMC $\mathcal{M}$ can be "matched" by a path through DTA $\mathcal{A}$ by regarding sets of atomic propositions in $\mathcal{M}$ as input symbols of $\mathcal{A}$. Such path is accepted, if at some point an accepting location in the DTA is reached:

Definition 12 (CTMC paths accepted by a DTA). *Let CTMC $\mathcal{M} = (S, R, AP, L, s_0)$ and DTA $\mathcal{A} = (2^{\mathrm{AP}}, \mathcal{X}, Q, q_0, F, \to)$. The CTMC path $\pi = s_0 \xrightarrow{t_0} s_1 \xrightarrow{t_1} s_2 \cdots$ is accepted by $\mathcal{A}$ if there exists a corresponding DTA path*

$$q_0 \xrightarrow{L(s_0),t_0} \underbrace{\mathit{succ}\big(q_0, L(s_0), g_0\big)}_{=q_1} \xrightarrow{L(s_1),t_1} \underbrace{\mathit{succ}\big(q_1, L(s_1), g_1\big)}_{=q_2} \cdots$$

such that $q_j \in F$ for some $j \geqslant 0$. Here, $\eta_0 = \mathbf{0}$, g_i is the (unique) guard in q_i such that $\eta_i{+}t_i \models g_i$ and $\eta_{i+1} = (\eta_i{+}t_i)[X_i := 0]$, and η_i is the clock evaluation on entering q_i, for $i \geqslant 0$. Let $\mathit{Paths}^{\mathcal{M}}(\mathcal{A}) = \{\pi \in \mathit{Paths}^{\mathcal{M}} \mid \pi$ is accepted by DTA $\mathcal{A}\}$.

Theorem 2 ([12]). *For any CTMC $\mathcal{M}$ and DTA $\mathcal{A}$, the set $\mathit{Paths}^{\mathcal{M}}(\mathcal{A})$ is measurable.*

The main result of this theorem is that $\mathit{Paths}^{\mathcal{M}}(\mathcal{A})$ can be rewritten as the combination of cylinder sets of the form $\mathit{Cyl} = (s_0, I_0, \ldots, I_{n-1}, s_n)$ (Cyl for short) which are all accepted by DTA $\mathcal{A}$. A cylinder set (Cyl) is accepted by DTA $\mathcal{A}$ if all its paths are accepted by $\mathcal{A}$. That is

$$\mathit{Paths}^{\mathcal{M}}(\mathcal{A}) = \bigcup_{n \in \mathbb{N}} \bigcup_{\pi \in \mathit{Paths}_n^{\mathcal{M}}(\mathcal{A})} \mathit{Cyl}_\pi, \qquad (1)$$

where $Paths_n^{\mathcal{M}}(\mathcal{A})$ is the set of accepting paths by $\mathcal{A}$ of length n and Cyl_π is the cylinder set that contains π.

Definition 13 (WL equivalent cylinder sets). *Cylinder sets $Cyl = (s_0, I_0, \ldots, I_{n-1}, s_n)$ and $Cyl' = (s'_0, I_0, \ldots, I_{n-1}, s'_n)$ are WL equivalent, denoted $Cyl \cong Cyl'$, if they are statewise WL equivalent: $Cyl \cong Cyl'$ iff $s_i \cong s'_i$ for all $0 \leqslant i \leqslant n$.*

Definition 14 (WL-closed set). *The set Π of cylinder sets is WL-closed if $\forall Cyl \in \Pi$, and Cyl' with $Cyl' \cong Cyl$ implies $Cyl' \in \Pi$.*

A finite path π in the CTMC $\mathcal{M}$ is compatible with Π if the cylinder set for this path $Cyl_\pi \in \Pi$. Since the cylinder sets contained in Π are disjoint, we have $\Pr_s(\Pi) = \Pr_s(\bigcup_{Cyl \in \Pi} Cyl) = \sum_{Cyl \in \Pi} \Pr_s(Cyl)$, where $\Pr_s(\Pi)$ is the probability of all the paths starting in s which are compatible with Π. For paths compatible with Π but not starting from s, the probability equals 0. We denote WL-closed set of cylinder sets of length n by Π_n. If $n = 0$, Π_n is the set of states and $\Pr_s(\Pi_n) = \alpha(s)$ if $s \in \Pi_n$, 0 otherwise, where $\alpha(s)$ is the probability of s being the initial state of CTMC $\mathcal{M}$.

Example 6. Consider the example given in Fig. 3, where we have the CTMC $\mathcal{M}$ (left) and its quotient $\mathcal{M}/_{\mathcal{R}}$ (right). If $\Pi = \{Cyl(s_0, I_0, s_1, I_1, s_3), Cyl(s'_0, I_0, s'_1, I_1, s'_2)\}$ is a WL closed set of cylinder sets in $\mathcal{M}$, and $\mathcal{M}/_{\mathcal{R}}$ that are accepted by DTA $\mathcal{A}$, then:

$$\begin{aligned}\Pr_{s_0}(\Pi) &= \Pr_{s_0}(Cyl(s_0, I_0, s_1, I_1, s_3)) + \Pr_{s_0}(Cyl(s'_0, I_0, s'_1, I_1, s'_2)) \\ &= 1/2 \cdot (e^{-E(s_0)\cdot \inf I_0} - e^{-E(s_0)\cdot \sup I_0}) \cdot (e^{-E(s_1)\cdot \inf I_1} - e^{-E(s_1)\cdot \sup I_1}) + 0.\end{aligned}$$

The second term is 0 as the cylinder set does not start from s_0. Similarly,

$$\begin{aligned}\Pr_{s'_0}(\Pi) &= \Pr_{s'_0}(Cyl(s_0, I_0, s_1, I_1, s_3)) + \Pr_{s'_0}(Cyl(s'_0, I_0, s'_1, I_1, s'_2)) \\ &= 0 + (e^{-E(s'_0)\cdot \inf I_0} - e^{-E(s'_0)\cdot \sup I_0}) \cdot 1/2 \cdot (e^{-E(s'_1)\cdot \inf I_1} - e^{-E(s'_1)\cdot \sup I_1}).\end{aligned}$$

Definition 15. *For CTMC $\mathcal{M}$ and DTA $\mathcal{A}$, let $\Pr(\mathcal{M} \models \mathcal{A}) = \Pr\left(Paths^{\mathcal{M}}(\mathcal{A})\right)$.*

Stated in words, $\Pr(\mathcal{M} \models \mathcal{A})$ denotes the probability of all the paths in CTMC $\mathcal{M}$ that are accepted by DTA $\mathcal{A}$. Note that we slightly abuse notation, since Pr on the right-hand side is the probability measure on the Borel space of infinite paths in the CTMC. This brings us to one of the main results of this paper:

Theorem 3 (Preservation of DTA specifications). *For any CTMC $\mathcal{M}$, a WL $\mathcal{R}$ on $\mathcal{M}$ and DTA $\mathcal{A}$:*

$$\Pr(\mathcal{M} \models \mathcal{A}) \; = \; \Pr(\mathcal{M}/_{\mathcal{R}} \models \mathcal{A}).$$

The detailed proof is in the appendix and consists of two main steps:

1. We prove that for any cylinder set Cyl in the quotient CTMC $\mathcal{M}/_{\mathcal{R}}$ which is accepted by the DTA $\mathcal{A}$, there is a corresponding *set* of cylinder sets in the CTMC $\mathcal{M}$ that are acccepted by the DTA $\mathcal{A}$ and that jointly have the same probability as Cyl, cf. Lemma 2 below.

2. We show that the sum of probabilities of all the cylinder sets in $\mathcal{M}/_{\mathcal{R}}$ that are accepted by DTA $\mathcal{A}$ equals the sum of probabilities of all the corresponding sets of cylinder sets in $\mathcal{M}$.

Lemma 2. *Let $\mathcal{M} = (S, R, AP, L, s_0)$ be a CTMC and $\mathcal{R}$ be a WL on $\mathcal{M}$. If Π is a WL-closed set of cylinder sets which are accepted by DTA $\mathcal{A}$, then for any $D \in S/_{\mathcal{R}}$ and $s_0' \in pred(D)$:*

$$\sum_{s_1 \in D} P(s_0', s_1, D) \cdot \Pr_{s_1}(\Pi) = \Pr_{D}(\Pi).$$

From Lemma 2 we conclude

$$\sum_{D \in S/_{\mathcal{R}}} \sum_{s_1 \in D} P(s_0', s_1, D) \cdot \Pr_{s_1}(\Pi) = \sum_{D \in S/_{\mathcal{R}}} \Pr_{D}(\Pi). \tag{2}$$

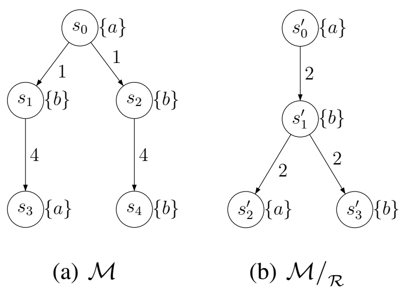

Fig. 3. WL equivalent cylinder sets

Corollary 1. *WL preserves transient state probabilities.*

5 Preservation of MTL Specifications

In this section we show that the quotient CTMC obtained under WL can be used for verifying Metric Temporal Logic (MTL) formulae [16,18,10]. It is interesting to note that the expressive power of MTL is different from that of DTA. Temporal properties like $(\Diamond\Box a)$ cannot be expressed using deterministic timed automata, since nondeterminism is needed to compensate for the non causality [17]. On the other hand, DTA expressible languages that involve counting [3], e.g., a should only occur at even positions, cannot be expressed using MTL. We now recall the syntax and semantics of Metric Temporal Logic [18,10].

Definition 16 (Syntax of MTL). *Let AP be a set of atomic propositions, then the formulas of MTL are built from AP using Boolean connectives, and time-constrained versions of the* until *operator* U *as follows:*

$$\varphi ::= tt \mid a \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \, \mathrm{U}^I \, \varphi$$

where $I \subseteq \mathbb{R}_{\geq 0}$ is a nonempty interval with rational bounds, and $a \in AP$.

Whereas, typically, the semantics of MTL is defined over timed paths of timed automata, we take a similar approach by interpreting MTL formulas over CTMC paths.

Definition 17 (Semantics of MTL formulas). *The meaning of MTL formulas is defined by means of a satisfaction relation, denoted by* $\models$*, between a CTMC* $\mathcal{M}$*, one of its paths* π*, MTL formula* φ*, and time* $t \in \mathbb{R}_{\geq 0}$*. Let* $\pi = s_0 \xrightarrow{t_0} s_1 \cdots s_{n-1} \xrightarrow{t_{n-1}} s_n \cdots$ *be a finite or infinite path of* $\mathcal{M}$*, then* $(\pi, t) \models \varphi$ *is defined inductively by:*

$$
\begin{array}{ll}
(\pi, t) \models tt & \\
(\pi, t) \models a & \textit{iff } a \in L(\pi@t) \\
(\pi, t) \models \neg\varphi & \textit{iff not } (\pi, t) \models \varphi \\
(\pi, t) \models \varphi_1 \wedge \varphi_2 & \textit{iff } (\pi, t) \models \varphi_1 \textit{ and } (\pi, t) \models \varphi_2 \\
(\pi, t) \models \varphi_1 \, \mathrm{U}^I \, \varphi_2 & \textit{iff } \exists t' \in t{+}I. \, ((\pi, t') \models \varphi_2 \wedge \forall t \leq t'' < t'. \, (\pi, t'') \models \varphi_1) \, .
\end{array}
$$

The semantics for the propositional fragment is straightforward. Recall that $\pi@t$ denotes the state occupied along path π at time t. Path π at time t satisfies $\varphi_1 \, \mathrm{U}^I \, \varphi_2$ whenever for some time point t' in the interval $I{+}t$, defined as $[a, b]{+}t = [a{+}t, b{+}t]$ (and similarly for open intervals), φ_2 holds, and at all time points between t and t', path π satisfies φ_1. Let $\pi \models \varphi$ if and only if $(\pi, 0) \models \varphi$. The standard temporal operators like $\Diamond$ ("eventually") and its timed variant $\Diamond^I$ are derived in the following way: $\Diamond^I \varphi = tt \, \mathrm{U}^I \, \varphi$ and$\Diamond \varphi = tt \, \mathrm{U} \, \varphi$. Similarly, $\Box$ ("globally") and its timed variant are derived as follows:

$$\Box^I \varphi = \neg(\Diamond^I \neg\varphi) \text{ and } \Box\varphi = \neg(\Diamond\neg\varphi).$$

Example 7. Using MTL, various interesting properties can be specified such as:

- $\Box(down \rightarrow \Diamond^{[0,5]} up)$, which asserts that whenever the system is down, it should be up again within 5 time units.
- $\Box(down \rightarrow alarm \, \mathrm{U}^{[0,10]} \, up)$, which states that whenever the system is down, an alarm should ring until it is up again within 10 time units.

More complex properties can be specified by nesting of until path formulas.

Theorem 4 ([2]). *The probability measure of the set of converging paths is zero.*

As a next result, we address the measurability of a set of CTMC paths satisfying an MTL formula φ.

Theorem 5. *For each MTL formula* φ *and state* s *of CTMC* $\mathcal{M}$*, the set* $\{\pi \in Paths(s) \mid \pi \models \varphi\}$ *is measurable.*

Definition 18 (Probability of MTL formulas). *The probability that state* s *satisfies MTL formula* φ *refers to the probability for the sets of paths for which that formula holds as follows:*

$$\Pr(s \models \varphi) \;=\; \Pr_s(\pi \in Paths(s) \mid \pi \models \varphi).$$

Since $\mathcal{M}$ has a single initial state, i.e., s_0, the probability of all the paths in $\mathcal{M}$ that satisfy MTL formula φ is given by $\Pr(\mathcal{M} \models \varphi) = \Pr(s_0 \models \varphi)$.

Theorem 6. *Let* $\mathcal{M}$ *be a CTMC and* $\mathcal{R}$ *be a WL on* $\mathcal{M}$*. Then for any MTL formula* φ*:*

$$\Pr(\mathcal{M} \models \varphi) = \Pr(\mathcal{M}/_{\mathcal{R}} \models \varphi).$$

6 Conclusions

This paper considered weighted lumpability (WL), a structural notion that coincides with Bernardo's T-lumpability [6] defined in a process-algebraic setting. Whereas Bernardo defines T-lumpability in an axiomatic manner, our starting point is a structural definition using first CTMC principles. The main contribution of this paper is the preservation of DTA and MTL specifications under WL quotienting. We note that this implies the preservation of transient probabilities as well as timed reachability probabilities. Future work is to develop an efficient quotienting algorithm for WL; we hope that our structural definition facilitates a reduction algorithm along the partition-refinement paradigm.

Acknowledgements. We thank Marco Bernardo for his constructive comments on a draft version of this paper.

References

1. Alur, R., Dill, D.L.: A theory of timed automata. Theor. Comput. Sci. 126(2), 183–235 (1994)
2. Baier, C., Haverkort, B.R., Hermanns, H., Katoen, J.-P.: Model-checking algorithms for continuous-time Markov chains. IEEE Trans. Software Eng. 29(6), 524–541 (2003)
3. Baier, C., Katoen, J.-P.: Principles of Model Checking. MIT Press (2008)
4. Baier, C., Katoen, J.-P., Hermanns, H., Wolf, V.: Comparative branching-time semantics for Markov chains. Inf. Comput. 200(2), 149–214 (2005)
5. Barbot, B., Chen, T., Han, T., Katoen, J.-P., Mereacre, A.: Efficient CTMC Model Checking of Linear Real-Time Objectives. In: Abdulla, P.A., Leino, K.R.M. (eds.) TACAS 2011. LNCS, vol. 6605, pp. 128–142. Springer, Heidelberg (2011)
6. Bernardo, M.: Non-bisimulation-based Markovian behavioral equivalences. J. Log. Algebr. Program. 72(1), 3–49 (2007)
7. Bernardo, M.: Towards State Space Reduction Based on T-Lumpability-Consistent Relations. In: Thomas, N., Juiz, C. (eds.) EPEW 2008. LNCS, vol. 5261, pp. 64–78. Springer, Heidelberg (2008)
8. Bernardo, M.: Uniform logical characterizations of testing equivalences for nondeterministic, probabilistic and Markovian processes. ENTCS 253(3), 3–23 (2009)
9. Bernardo, M., Botta, S.: A survey of modal logics characterising behavioural equivalences for non-deterministic and stochastic systems. Math. Structures in Comp. Sci. 18(1), 29–55 (2008)
10. Bouyer, P.: From Qualitative to Quantitative Analysis of Timed Systems. Mémoire d'habilitation, Université Paris 7, Paris, France (January 2009)
11. Buchholz, P.: Exact and ordinary lumpability in finite Markov chains. J. of Appl. Prob. 31(1), 59–75 (1994)
12. Chen, T., Han, T., Katoen, J.-P., Mereacre, A.: Quantitative model checking of continuous-time Markov chains against timed automata specifications. In: LICS, pp. 309–318. IEEE Computer Society (2009)
13. Derisavi, S., Hermanns, H., Sanders, W.H.: Optimal state-space lumping in Markov chains. Inf. Process. Lett. 87(6), 309–315 (2003)
14. Donatelli, S., Haddad, S., Sproston, J.: Model checking timed and stochastic properties with CSL^{TA}. IEEE Trans. Software Eng. 35(2), 224–240 (2009)

15. Katoen, J.-P., Kemna, T., Zapreev, I.S., Jansen, D.N.: Bisimulation Minimisation Mostly Speeds Up Probabilistic Model Checking. In: Grumberg, O., Huth, M. (eds.) TACAS 2007. LNCS, vol. 4424, pp. 87–101. Springer, Heidelberg (2007)
16. Koymans, R.: Specifying real-time properties with metric temporal logic. Real-Time Systems 2(4), 255–299 (1990)
17. Maler, O., Nickovic, D., Pnueli, A.: Checking Temporal Properties of Discrete, Timed and Continuous Behaviors. In: Avron, A., Dershowitz, N., Rabinovich, A. (eds.) Trakhtenbrot/Festschrift. LNCS, vol. 4800, pp. 475–505. Springer, Heidelberg (2008)
18. Ouaknine, J., Worrell, J.: Some Recent Results in Metric Temporal Logic. In: Cassez, F., Jard, C. (eds.) FORMATS 2008. LNCS, vol. 5215, pp. 1–13. Springer, Heidelberg (2008)
19. Valmari, A., Franceschinis, G.: Simple $O(m \log n)$ Time Markov Chain Lumping. In: Esparza, J., Majumdar, R. (eds.) TACAS 2010. LNCS, vol. 6015, pp. 38–52. Springer, Heidelberg (2010)
20. Wolf, V., Baier, C., Majster-Cederbaum, M.E.: Trace machines for observing continuous-time Markov chains. ENTCS 153(2), 259–277 (2006)

Appendix

Proof of Theorem 1

Proof. Let $\mathcal{M} = (S, R, AP, L, s_0)$ be a CTMC and $\mathcal{M}/_{\mathcal{R}} = (S/_{\mathcal{R}}, R', AP, L', s'_0)$ be its quotient under WL. Since we have defined the WL relation on a single state space, to prove this theorem we take the disjoint union $S \cup S/_{\mathcal{R}}$. Let us define an equivalence relation $\mathcal{R}^* \subseteq (S \cup S/_{\mathcal{R}}) \times (S \cup S/_{\mathcal{R}})$ with $\{(s, C) | s \in C\} \subseteq \mathcal{R}^*$. The exit rate $E'(C)$ for $C \in S/_{\mathcal{R}}$ is defined by $\sum\limits_{x \in (S \cup S/_{\mathcal{R}})} R'(C, x)$.

Now we prove that $\mathcal{R}^*$ is a WL relation. This is done by checking both conditions of Def. 7. Let $(s, C) \in \mathcal{R}^*$. The proofs for pairs (s, s'), (C, s), and (C, C) are similar and omitted.

1. $L'(C) = L(s)$ by definition of $\mathcal{M}/_{\mathcal{R}}$. We prove that $E'(C) = E(s)$ as follows:

$$\begin{aligned}
E'(C) &= \sum_{x \in (S \cup S/_{\mathcal{R}})} R'(C, x) = \sum_{D \in S/_{\mathcal{R}}} R'(C, D) \\
&= \sum_{D \in S/_{\mathcal{R}}} wr(s'_0, C, D) \text{ for some } s'_0 \in pred(C) \\
&= \sum_{D \in S/_{\mathcal{R}}} \sum_{s \in C} P(s'_0, s, C) \cdot R(s, D) \\
&= \sum_{s \in C} \left(P(s'_0, s, C) \cdot \sum_{D \in S/_{\mathcal{R}}} R(s, D) \right) \\
&= \sum_{s \in C} \left(P(s'_0, s, C) \cdot \sum_{D \in S/_{\mathcal{R}}} \sum_{s' \in D} R(s, s') \right) \\
&= \sum_{s \in C} \left(P(s'_0, s, C) \cdot \sum_{s' \in S} R(s, s') \right)
\end{aligned}$$

$$= \sum_{s \in C} (P(s_0', s, C) \cdot E(s))$$
$$= \left(\sum_{s \in C} P(s_0', s, C) \right) \cdot E(s), \text{ since for all } s' \in C, E(s') = E(s)$$
$$= E(s).$$

2. Finally we prove that $\forall E, F \in (S \cup S/_{\mathcal{R}})/_{\mathcal{R}^*}$ and $\forall x_0', x_0'' \in pred(E)$ it holds $wr(x_0', E, F) = wr(x_0'', E, F)$. Let $x_0', x_0'' \in pred(E)$. Consider the following three cases based on the successors of x_0', x_0'' such that these successors are in E.
   a) The successors of both x_0', x_0'' belong to S. Since we know that $\mathcal{R}$ is a WL, it follows $wr(x_0', E, F) = wr(x_0'', E, F)$.
   b) The successors of both x_0', x_0'' belong to $S/_{\mathcal{R}}$. In this case, $wr(x_0', E, F)$ $= wr(x_0', \{E_1\}, F)$ where $E_1 \in E \cap S/_{\mathcal{R}}$, which equals

$$\sum_{x' \in \{E_1\}} \frac{P(x_0', x')}{P(x_0', E_1)} \cdot R'(x', F) = R'(E_1, F).$$

   Similarily $wr(x_0'', E, F) = wr(x_0'', \{E_1\}, F) = R'(E_1, F)$.
   c) The successors of x_0', x_0'' belong to S and $S/_{\mathcal{R}}$ respectively. In this case we get, $wr(x_0'', E, F) = wr(x_0'', \{E_1\}, F) = R'(E_1, F)$.
   We know that the successors of $E_1 \in S/_{\mathcal{R}}$, hence using Def. 8 we conclude:

$$R'(E_1, F) = wr(x_0', E_1, F) = wr(x_0', E, F).$$

Since all the conditions of Def. 7 are satisfied by the relation $\mathcal{R}^*$, it is a WL relation. □

Proof of Lemma 1

Proof. Let $s_1 \sim s_2$. We prove that both conditions for $\cong$ are satisfied.

- $L(s_1) = L(s_2)$, follows directly from $s_1 \sim s_2$.
- $E(s_1) = E(s_2)$, since we know that

$$E(s_1) = \sum_{s \in S} R(s_1, s) = \sum_{C \in S/_\sim} \sum_{s \in C} R(s_1, s) = \sum_{C \in S/_\sim} R(s_1, C).$$

  If $s_1 \sim s_2$, then $R(s_1, C) = R(s_2, C)$. Therefore:

$$E(s_1) = \sum_{C \in S/_\sim} R(s_1, C) = \sum_{C \in S/_\sim} R(s_2, C) = E(s_2).$$

– Let $C, D \in S/_\sim$ and $s_0', s_0'' \in \mathit{pred}(C)$. Since $R(s_1, D) = R(s_2, D)$ for all $s_1, s_2 \in C$, then for all $s^* \in C$:

$$\begin{aligned} wr(s_0', C, D) &= \sum_{s \in C} \frac{P(s_0', s)}{P(s_0', C)} \cdot R(s, D) \\ &= R(s^*, D) \cdot \sum_{s \in C} \frac{P(s_0', s)}{P(s_0', C)} \\ &= R(s^*, D) \\ &= R(s^*, D) \cdot \sum_{s \in C} \frac{P(s_0'', s)}{P(s_0'', C)} \\ &= \sum_{s \in C} \frac{P(s_0'', s)}{P(s_0'', C)} \cdot R(s, D) \\ &= wr(s_0'', C, D). \end{aligned}$$

Thus $s_1 \cong s_2$. □

Proof of Lemma 2

Proof. We will prove this lemma by induction over the length of the cylinder set $Cyl \in \mathit{\Pi}$. That is, we will prove for any $n \in \mathbb{N}$:

$$\sum_{s_1 \in D} P(s_0', s_1, D) \cdot \Pr_{s_1}(\mathit{\Pi}_n) = \Pr_D(\mathit{\Pi}_n).$$

– **Base Case**: In this case, $n = 0$ and

$$\sum_{s_1 \in D} P(s_0', s_1, D) \cdot \Pr_{s_1}(\mathit{\Pi}_0) = 1 = \Pr_D(\mathit{\Pi}_0),$$

if $s_0 \in D, \mathit{\Pi}_0$, and 0, otherwise.

– **Induction Hypothesis**: Assume that for cylinder sets of length $n \in \mathbb{N}$, it holds:

$$\sum_{s_1 \in D} P(s_0', s_1, D) \cdot \Pr_{s_1}(\mathit{\Pi}_n) = \Pr_D(\mathit{\Pi}_n).$$

– **Induction Step**: Consider the case $n + 1$:

$$\begin{aligned} &\sum_{s_1 \in D} P(s_0', s_1, D) \cdot \Pr_{s_1}(\mathit{\Pi}_{n+1}) \\ &= \sum_{s_1 \in D} P(s_0', s_1, D) \cdot \sum_{s_2 \in S} P(s_1, s_2) \cdot (e^{-E(s_1) \cdot \inf I_0} - e^{-E(s_1) \cdot \sup I_0}) \cdot \Pr_{s_2}(\mathit{\Pi}_n) \end{aligned}$$

Let $(e^{-E(s_1)\cdot \inf I_0} - e^{-E(s_1)\cdot \sup I_0}) = \delta(s_1, I_0)$, then the above expression is equal to:

$$\begin{aligned}
&\sum_{s_1\in D} P(s_0', s_1, D)\cdot \sum_{s_2\in S} P(s_1, s_2)\cdot \delta(s_1, I_0)\cdot \Pr_{s_2}(\Pi_n)\\
&\quad = \sum_{s_1\in D} P(s_0', s_1, D)\cdot \sum_{C\in S/_{\mathcal{R}}} \sum_{s_2\in C} P(s_1, s_2)\cdot \delta(s_1, I_0)\cdot \Pr_{s_2}(\Pi_n)\\
&\quad = \sum_{C\in S/_{\mathcal{R}}} \sum_{s_1\in D} P(s_0', s_1, D)\cdot \sum_{s_2\in C} P(s_1, s_2)\cdot \delta(s_1, I_0)\cdot \Pr_{s_2}(\Pi_n).
\end{aligned}$$

Multiplying the above expression by $\dfrac{R(s_1, C)}{R(s_1, C)}$ and using $P(s_1, s_2) = \dfrac{R(s_1, s_2)}{E(s_1)}$ yields:

$$\sum_{C\in S/_{\mathcal{R}}} \sum_{s_1\in D} P(s_0', s_1, D)\cdot \sum_{s_2\in D} \frac{R(s_1, C)}{R(s_1, C)}\cdot \frac{R(s_1, s_2)}{E(s_1)}\cdot \delta(s_1, I_0)\cdot \Pr_{s_2}(\Pi_n).$$

Since $\forall s_1, s_1' \in D$, $E(s_1) = E(s_1')$, we have $\delta(s_1, I_0) = \delta(s_1', I_0)$. We get:

$$\frac{\delta(s_1, I_0)}{E(s_1)} \sum_{C\in S/_{\mathcal{R}}} \sum_{s_1\in D} P(s_0', s_1, D)\cdot R(s_1, C)\cdot \sum_{s_2\in C} \frac{R(s_1, s_2)}{R(s_1, C)}\cdot \Pr_{s_2}(\Pi_n)$$

$$= \frac{\delta(s_1, I_0)}{E(s_1)} \sum_{C\in S/_{\mathcal{R}}} \sum_{s_1\in D} P(s_0', s_1, D)\cdot R(s_1, C)\cdot \sum_{s_2\in C} P(s_1, s_2, C)\cdot \Pr_{s_2}(\Pi_n).$$

We have already proved that $\forall s \in D$, $E(s) = E(D)$, cf. Theorem 1. From the induction hypothesis we have:

$$\sum_{s_2\in C} P(s_1, s_2, C)\cdot \Pr_{s_2}(\Pi_n) = \Pr_{C}(\Pi_n).$$

Also from Def. 6 and Def. 8 we know that:

$$\sum_{C\in S/_{\mathcal{R}}} \sum_{s_1\in D} P(s_0', s_1, D)\cdot R(s_1, C) = \sum_{C\in S/_{\mathcal{R}}} R'(D, C),$$

since $\sum_{s_1\in D} P(s_0', s_1, D)\cdot R(s_1, C) = wr(s_0', D, C) = R'(D, C)$. Therefore we get:

$$\frac{\delta(D, I_0)}{E(D)} \sum_{C\in S/_{\mathcal{R}}} R'(D, C)\cdot \Pr_{C}(\Pi_n) = \Pr_{D}(\Pi_{n+1}).$$

□

Proof of Theorem 3

Proof. Let C_n be the set of all the cylinder sets in $\mathcal{M}$, and $\mathcal{M}/_{\mathcal{R}}$ of length n that are accepted by DTA $\mathcal{A}$ and $C_{n/\Pi}$ be the set of subsets of C_n grouped according to WL-closed set of cylinder sets. Let Cyl_π be the cylinder set that contains π. Since the cylinder sets in Eq. (1) are disjoint, we have:

$$\begin{aligned}\Pr(\mathcal{M} \models \mathcal{A}) &= \Pr\left(\bigcup_{n\in\mathbb{N}} \bigcup_{\pi \in Paths_n^{\mathcal{M}}(\mathcal{A})} Cyl_\pi\right)\\ &= \sum_{n\in\mathbb{N}} \sum_{Cyl \in C_n} \Pr(Cyl)\\ &= \sum_{n\in\mathbb{N}} \sum_{\Pi \in C_{n/\Pi}} \sum_{D\in S/_R} \sum_{s_1 \in D} P(s_0', s_1, D) \cdot \Pr_{s_1}(\Pi).\end{aligned}$$

Then we get using Eq. (2):

$$\begin{aligned}\Pr(\mathcal{M} \models \mathcal{A}) &= \sum_{n\in\mathbb{N}} \sum_{\Pi \in C_{n/\Pi}} \sum_{D\in S/_{\mathcal{R}}} \sum_{s_1 \in D} P(s_0', s_1, D) \cdot \Pr_{s_1}(\Pi)\\ &= \sum_{n\in\mathbb{N}} \sum_{\Pi \in C_{n/\Pi}} \sum_{D\in S/_{\mathcal{R}}} \Pr_D(\Pi)\\ &= \Pr(\mathcal{M}/_{\mathcal{R}} \models \mathcal{A}).\end{aligned}$$

□

Proof of Theorem 5

Proof. We prove the measurability by showing that for any path $\pi = s_0 \xrightarrow{t_0} s_1 \xrightarrow{t_1} s_2 \cdots s_{n-1} \xrightarrow{t_{n-1}} s_n \in Paths_n^{\mathcal{M}}(s_0 \models \varphi)$ where $Paths_n^{\mathcal{M}}(s_0 \models \varphi)$ is the set of paths of length n starting from s_0 that satisfy φ, there exists a cylinder set $Cyl(s_0, I_0, ..., I_{n-1}, s_n$ (Cyl for short) s.t. $\pi \in Cyl$ and $Cyl \subseteq Paths_n^{\mathcal{M}}(s_0 \models \varphi)$. Since the only interesting case is time-bounded "until", we consider $\varphi = \varphi_1 \mathrm{U}^{[a,b]} \varphi_2$, where a, $b \in \mathbb{Q}$. Let $\sum_{i=0}^{n-1} t_i - \Delta > a$ and $\sum_{i=0}^{n-2} t_i + \Delta < b$, where $\Delta = \dfrac{2n}{10^k}$, and k is large enough. We construct Cyl by considering intervals I_i with rational bounds that are based on t_i. Let $I_i = [t_i^-, t_i^+]$ s.t. $t_i^- = t_i = t_i^+$ if $t_i \in \mathbb{Q}$, and otherwise:

$$t_i^- < t_i < t_i^+,\ t_i^- > t_i - \frac{\Delta}{2n} \text{ and } t_i^+ < t_i + \frac{\Delta}{2n}.$$

We have to show for $t_i \notin \mathbb{Q}$, Eq. (3) and Eq. (4) hold:

$$\sum_{i=0}^{n-1} t_i^- > a. \tag{3}$$

Proof. We know that $\displaystyle\sum_{i=0}^{n-1} t_i - \Delta > a \implies \sum_{i=0}^{n-1} t_i^- + n \cdot \frac{\Delta}{2n} - \Delta > a$

$$\implies \sum_{i=0}^{n-1} t_i^- + \frac{\Delta}{2} - \Delta > a \implies \sum_{i=0}^{n-1} t_i^- - \frac{\Delta}{2} > a \implies \sum_{i=0}^{n-1} t_i^- > a.$$

$$\sum_{i=0}^{n-2} t_i^+ < b. \tag{4}$$

Proof. We know that $\sum_{i=0}^{n-2} t_i + \Delta < b \implies \sum_{i=0}^{n-2} t_i^+ - (n-1) \cdot \dfrac{\Delta}{2n} + \Delta < b$
$\implies \sum_{i=0}^{n-2} t_i^+ + \dfrac{(n+1) \cdot \Delta}{2n} < b \implies \sum_{i=0}^{n-2} t_i^+ < b.$

One way is to pick t_i^-, t_i^+ as follows:

$$t_i^- = \lfloor t_i \rfloor + \frac{\lfloor \{t_i\} \cdot 10^k \rfloor}{10^k},$$

$$t_i^+ = \lfloor t_i \rfloor + \frac{\lfloor \{t_i\} \cdot 10^k \rfloor + 1}{10^k},$$

where $\{t_i\}$ represents the fractional part of the irrational number t_i. It can be checked that picking t_i^-, t_i^+ this way satisfies the above mentioned constraints.

From this derivation we conclude that $\{\pi \in Paths(s_0) | \pi \models \varphi\}$ can be rewritten as the combination of cylinder sets of the form $Cyl = (s_0, I_0, ..., I_{n-1}, s_n)$. That is,

$$\{\pi \in Paths(s_0) | \pi \models \varphi\} = \bigcup_{n \in \mathbb{N}} \bigcup_{\pi \in Paths_n^{\mathcal{M}}(s_0 \models \varphi)} Cyl_\pi, \tag{5}$$

where $Paths_n^{\mathcal{M}}(s_0 \models \varphi)$ is the set of paths of length n starting from s_0 which satisfy φ. □

Proof of Theorem 6

Proof. The proof is similar to that of Theorem 3. We consider the WL-closed set of cylinder sets of length n in $\mathcal{M}$, $\mathcal{M}/_{\mathcal{R}}$ such that this set satisfies φ. The rest of the proof remains the same. □