

Bisimulation and Logical Preservation for Continuous-Time Markov Decision Processes

Martin R. Neuhäuser^{1,2} and Joost-Pieter Katoen^{1,2}

¹ Software Modeling and Verification Group
RWTH Aachen University, Germany

² Formal Methods and Tools Group
University of Twente, The Netherlands
{neuhaeusser,katoen}@cs.rwth-aachen.de

Abstract. This paper introduces strong bisimulation for continuous-time Markov decision processes (CTMDPs), a stochastic model which allows for a nondeterministic choice between exponential distributions, and shows that bisimulation preserves the validity of CSL. To that end, we interpret the semantics of CSL—a stochastic variant of CTL for continuous-time Markov chains—on CTMDPs and show its measure-theoretic soundness. The main challenge faced in this paper is the proof of logical preservation that is substantially based on measure theory.

1 Introduction

Discrete-time probabilistic models, in particular Markov decision processes (MDP) [20], are used in various application areas such as randomized distributed algorithms and security protocols. A plethora of results in the field of concurrency theory and verification are known for MDPs. Efficient model-checking algorithms exist for probabilistic variants of CTL [9,11], linear-time [30] and long-run properties [15], process algebraic formalisms for MDPs have been developed and bisimulation is used to minimize MDPs prior to analysis [18].

In contrast, CTMDPs [26], a continuous-time variant of MDPs, where state residence times are exponentially distributed, have received scant attention. Whereas in MDPs nondeterminism occurs between discrete probability distributions, in CTMDPs the choice between various exponential distributions is nondeterministic. In case all exponential delays are uniquely determined, a continuous-time Markov chain (CTMC) results, a widely studied model in performance and dependability analysis.

This paper proposes strong bisimulation on CTMDPs—this notion is a conservative extension of bisimulation on CTMCs [13]—and investigates which kind of logical properties this preserves. In particular, we show that bisimulation preserves the validity of CSL [3,5], a well-known logic for CTMCs. To that end, we provide a semantics of CSL on CTMDPs which is in fact obtained in a similar way as the semantics of PCTL on MDPs [9,11]. We show the semantic soundness of the logic using measure-theoretic arguments, and prove that bisimilar states

preserve full CSL. Although this result is perhaps not surprising, its proof is non-trivial and strongly relies on measure-theoretic aspects. It shows that reasoning about CTMDPs, as witnessed also by [31,7,10] is not straightforward. As for MDPs, CSL equivalence does not coincide with bisimulation as only maximal and minimal probabilities can be logically expressed.

Apart from the theoretical contribution, we believe that the results of this paper have wider applicability. CTMDPs are the semantic model of stochastic Petri nets [14] that exhibit confusion, stochastic activity networks [28] (where absence of nondeterminism is validated by a “well-specified” check), and is strongly related to interactive Markov chains which are used to provide compositional semantics to process algebras [19] and dynamic fault trees [12]. Besides, CTMDPs have practical applicability in areas such as stochastic scheduling [17,1] and dynamic power management [27]. Our interest in CTMDPs is furthermore stimulated by recent results on abstraction—where the introduction of nondeterminism is the key principle—of CTMCs [21] in the context of probabilistic model checking.

In our view, it is a challenge to study this continuous-time stochastic model in greater depth. This paper is a small, though important, step towards a better understanding of CTMDPs. More details and all proofs can be found in [25].

2 Continuous-time Markov decision processes

Continuous-time Markov decision processes extend continuous-time Markov chains by nondeterministic choices. Therefore each transition is labelled with an action referring to the nondeterministic choice and the rate of a negative exponential distribution which determines the transition’s delay:

Definition 1 (Continuous-time Markov decision process). *A tuple $\mathcal{C} = (\mathcal{S}, Act, \mathbf{R}, AP, L)$ is a labelled continuous-time Markov decision process if \mathcal{S} is a finite, nonempty set of states, Act a finite, nonempty set of actions and $\mathbf{R} : \mathcal{S} \times Act \times \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$ a three-dimensional rate matrix. Further, AP is a finite set of atomic propositions and $L : \mathcal{S} \rightarrow 2^{AP}$ is a state labelling function.*

The set of actions that are enabled in a state $s \in \mathcal{S}$ is denoted $Act(s) := \{\alpha \in Act \mid \exists s' \in \mathcal{S}. \mathbf{R}(s, \alpha, s') > 0\}$. A CTMDP is *well-formed* if $Act(s) \neq \emptyset$ for all $s \in \mathcal{S}$, that is, if every state has at least one outgoing transition. Note that this can easily be established for any CTMDP by adding self-loops.

Example 1. When entering state s_1 of the CTMDP in Fig. 1 (without state labels) one action from the set of enabled actions $Act(s_1) = \{\alpha, \beta\}$ is chosen nondeterministically, say α . Next, the rate of the α -transition determines its exponentially distributed delay. Hence for a single transition, the probability to go from s_1 to s_3 within time t is $1 - e^{-\mathbf{R}(s_1, \alpha, s_3)t} = 1 - e^{-0.1t}$.




Fig. 1. Example of a CTMDP.

If multiple outgoing transitions exist for the chosen action, they compete according to their exponentially distributed delays: In Fig. 1 such a *race condition* occurs if action β is chosen in state s_1 . In this situation, two β -transitions (to s_2 and s_3) with rates $\mathbf{R}(s_1, \beta, s_2) = 15$ and $\mathbf{R}(s_1, \beta, s_3) = 5$ become available and state s_1 is left as soon as the first transition's delay expires. Hence the sojourn time in state s_1 is distributed according to the minimum of both exponential distributions, i.e. with rate $\mathbf{R}(s_1, \beta, s_2) + \mathbf{R}(s_1, \beta, s_3) = 20$. In general, $E(s, \alpha) := \sum_{s' \in \mathcal{S}} \mathbf{R}(s, \alpha, s')$ is the *exit rate* of state s under action α . Then $\mathbf{R}(s_1, \beta, s_2)/E(s_1, \beta) = 0.75$ is the probability to move with β from s_1 to s_2 , i.e. the probability that the delay of the β -transition to s_2 expires first. Formally, the *discrete branching probability* is $\mathbf{P}(s, \alpha, s') := \frac{\mathbf{R}(s, \alpha, s')}{E(s, \alpha)}$ if $E(s, \alpha) > 0$ and 0 otherwise. By $\mathbf{R}(s, \alpha, Q) := \sum_{s' \in Q} \mathbf{R}(s, \alpha, s')$ we denote the total rate to states in $Q \subseteq \mathcal{S}$.

Definition 2 (Path). Let $\mathcal{C} = (\mathcal{S}, Act, \mathbf{R}, AP, L)$ be a CTMDP. $Paths^n(\mathcal{C}) := \mathcal{S} \times (Act \times \mathbb{R}_{\geq 0} \times \mathcal{S})^n$ is the set of paths of length n in \mathcal{C} ; the set of finite paths in \mathcal{C} is defined by $Paths^*(\mathcal{C}) = \bigcup_{n \in \mathbb{N}} Paths^n$ and $Paths^\omega(\mathcal{C}) := (\mathcal{S} \times Act \times \mathbb{R}_{\geq 0})^\omega$ is the set of infinite paths in \mathcal{C} . $Paths(\mathcal{C}) := Paths^*(\mathcal{C}) \cup Paths^\omega(\mathcal{C})$ denotes the set of all paths in \mathcal{C} .

We write *Paths* instead of $Paths(\mathcal{C})$ whenever \mathcal{C} is clear from the context. Paths are denoted $\pi = s_0 \xrightarrow{\alpha_0, t_0} s_1 \xrightarrow{\alpha_1, t_1} \dots \xrightarrow{\alpha_{n-1}, t_{n-1}} s_n$ where $|\pi|$ is the length of π . Given a finite path $\pi \in Paths^n$, $\pi \downarrow$ is the last state of π . For $n < |\pi|$, $\pi[n] := s_n$ is the n -th state of π and $\delta(\pi, n) := t_n$ is the time spent in state s_n . Further, $\pi[i..j]$ is the path-infix $s_i \xrightarrow{\alpha_i, t_i} s_{i+1} \xrightarrow{\alpha_{i+1}, t_{i+1}} \dots \xrightarrow{\alpha_{j-1}, t_{j-1}} s_j$ of π for $i < j \leq |\pi|$. We write $\xrightarrow{\alpha, t} s'$ for a transition with action α at time point t to a successor state s' . The extension of a path π by a transition m is denoted $\pi \circ m$. Finally, $\pi @ t$ is the state occupied in π at time point $t \in \mathbb{R}_{\geq 0}$, i.e. $\pi @ t := \pi[n]$ where n is the smallest index such that $\sum_{i=0}^n t_i > t$.

Note that Def. 2 does not impose any semantic restrictions on paths, i.e. the set *Paths* usually contains paths which do not exist in the underlying CTMDP. However, the following definition of the probability measure (Def. 4) justifies this as it assigns probability zero to those sets of paths.

2.1 The probability space

In probability theory (see [2]), a *field* of sets $\mathfrak{F} \subseteq 2^\Omega$ is a family of subsets of a set Ω which contains the empty set and is closed under complement and finite union. A field \mathfrak{F} is a σ -*field*³ if it is also closed under countable union, i.e. if for all countable families $\{A_i\}_{i \in I}$ of sets $A_i \in \mathfrak{F}$ it holds $\bigcup_{i \in I} A_i \in \mathfrak{F}$. Any subset A of Ω which is in \mathfrak{F} is called *measurable*. To measure the probability of sets of paths, we define a σ -field of sets of *combined transitions* which we later use to define σ -fields of sets of finite and infinite paths: For CTMDP $\mathcal{C} = (\mathcal{S}, Act, \mathbf{R}, AP, L)$,

³ In the literature [22], σ -fields are also called σ -algebras.

the set of combined transitions is $\Omega = Act \times \mathbb{R}_{\geq 0} \times \mathcal{S}$. As \mathcal{S} and Act are finite, the corresponding σ -fields are $\mathfrak{F}_{Act} := 2^{Act}$ and $\mathfrak{F}_{\mathcal{S}} := 2^{\mathcal{S}}$; further, $Distr(Act)$ and $Distr(\mathcal{S})$ denote the sets of probability distributions on \mathfrak{F}_{Act} and $\mathfrak{F}_{\mathcal{S}}$. Any combined transition occurs at some time point $t \in \mathbb{R}_{\geq 0}$ so that we can use the Borel σ -field $\mathfrak{B}(\mathbb{R}_{\geq 0})$ to measure the corresponding subsets of $\mathbb{R}_{\geq 0}$.

A Cartesian product is a *measurable rectangle* if its constituent sets are elements of their respective σ -fields, i.e. the set $A \times T \times S$ is a measurable rectangle if $A \in \mathfrak{F}_{Act}$, $T \in \mathfrak{B}(\mathbb{R}_{\geq 0})$ and $S \in \mathfrak{F}_{\mathcal{S}}$. We use $\mathfrak{F}_{Act} \times \mathfrak{B}(\mathbb{R}_{\geq 0}) \times \mathfrak{F}_{\mathcal{S}}$ to denote the set of all measurable rectangles⁴. It generates the desired σ -field \mathfrak{F} of sets of combined transitions, i.e. $\mathfrak{F} := \sigma(\mathfrak{F}_{Act} \times \mathfrak{B}(\mathbb{R}_{\geq 0}) \times \mathfrak{F}_{\mathcal{S}})$.

Now \mathfrak{F} may be used to infer the σ -fields \mathfrak{F}_{Paths^n} of sets of paths of length n : \mathfrak{F}_{Paths^n} is generated by the set of measurable (path) rectangles, i.e. $\mathfrak{F}_{Paths^n} := \sigma(\{S_0 \times M_0 \times \dots \times M_n \mid S_0 \in \mathfrak{F}_{\mathcal{S}}, M_i \in \mathfrak{F}, 0 \leq i \leq n\})$. The σ -field of sets of infinite paths is obtained using the cylinder-set construction [2]: A set C^n of paths of length n is called a *cylinder base*; it induces the infinite *cylinder* $C_n = \{\pi \in Paths^\omega \mid \pi[0..n] \in C^n\}$. A cylinder C_n is *measurable* if $C^n \in \mathfrak{F}_{Paths^n}$; C_n is a *rectangle* if $C^n = S_0 \times A_0 \times T_0 \times \dots \times A_{n-1} \times T_{n-1} \times S_n$ and $S_i \subseteq \mathcal{S}$, $A_i \subseteq Act$ and $T_i \subseteq \mathbb{R}_{\geq 0}$. It is a *measurable rectangle*, if $S_i \in \mathfrak{F}_{\mathcal{S}}$, $A_i \in \mathfrak{F}_{Act}$ and $T_i \in \mathfrak{B}(\mathbb{R}_{\geq 0})$. Finally, the σ -field of sets of infinite paths is defined as $\mathfrak{F}_{Paths^\omega} := \sigma(\bigcup_{n=0}^{\infty} \{C_n \mid C^n \in \mathfrak{F}_{Paths^n}\})$.

2.2 The probability measure

To define a semantics for CTMDP we use schedulers⁵ to resolve the nondeterministic choices. Thereby we obtain probability measures on the probability spaces defined above. A scheduler quantifies the probability of the next action based on the history of the system: If state s is reached via finite path π , the scheduler yields a probability distribution over $Act(\pi \downarrow)$. The type of schedulers we use is the class of measurable timed history-dependent randomized schedulers [31]:

Definition 3 (Measurable scheduler). *Let \mathcal{C} be a CTMDP with action set Act . A mapping $\mathcal{D} : Paths^* \times \mathfrak{F}_{Act} \rightarrow [0, 1]$ is a measurable scheduler if $\mathcal{D}(\pi, \cdot) \in Distr(Act(\pi \downarrow))$ for all $\pi \in Paths^*$ and the functions $\mathcal{D}(\cdot, A) : Paths^* \rightarrow [0, 1]$ are measurable for all $A \in \mathfrak{F}_{Act}$. THR denotes the set of measurable schedulers.*

In Def. 3, the measurability condition states that for any $B \in \mathfrak{B}([0, 1])$ and $A \in \mathfrak{F}_{Act}$ the set $\{\pi \in Paths^* \mid \mathcal{D}(\pi, A) \in B\} \in \mathfrak{F}_{Paths^*}$, see [31]. In the following, note that $\mathcal{D}(\pi, \cdot)$ is a probability measure with support $\subseteq Act(\pi \downarrow)$; further $\mathbf{P}(s, \alpha, \cdot) \in Distr(\mathcal{S})$ if $\alpha \in Act(s)$. Let $\eta_{E(\pi \downarrow, \alpha)}(t) := E(\pi \downarrow, \alpha) \cdot e^{-E(\pi \downarrow, \alpha)t}$ denote the probability density function of the negative exponential distribution with parameter $E(\pi \downarrow, \alpha)$. To derive a probability measure on $\mathfrak{F}_{Paths^\omega}$, we first define a probability measure on (Ω, \mathfrak{F}) : For history $\pi \in Paths^*$, let $\mu_{\mathcal{D}}(\pi, \cdot) : \mathfrak{F} \rightarrow [0, 1]$ such that

$$\mu_{\mathcal{D}}(\pi, M) := \int_{Act} \mathcal{D}(\pi, d\alpha) \int_{\mathbb{R}_{\geq 0}} \eta_{E(\pi \downarrow, \alpha)}(dt) \int_{\mathcal{S}} \mathbf{I}_M(\alpha, t, s) \mathbf{P}(\pi \downarrow, \alpha, ds).$$

⁴ Despite notation, $\mathfrak{F}_{Act} \times \mathfrak{B}(\mathbb{R}_{\geq 0}) \times \mathfrak{F}_{\mathcal{S}}$ is not a Cartesian product.

⁵ Schedulers are also called policies or adversaries in the literature.

Then $\mu_{\mathcal{D}}(\pi, \cdot)$ defines a probability measure on \mathfrak{F} where the indicator function $\mathbf{I}_M(\alpha, t, s) := 1$ if the combined transition $(\alpha, t, s) \in M$ and 0 otherwise [31]. For a measurable rectangle $A \times T \times S' \in \mathfrak{F}$ we obtain

$$\mu_{\mathcal{D}}(\pi, A \times T \times S') = \sum_{\alpha \in A} \mathcal{D}(\pi, \{\alpha\}) \cdot \mathbf{P}(\pi \downarrow, \alpha, S') \cdot \int_T E(\pi \downarrow, \alpha) \cdot e^{-E(\pi \downarrow, \alpha)t} dt. \quad (1)$$

Intuitively, $\mu_{\mathcal{D}}(\pi, A \times T \times S')$ is the probability to leave $\pi \downarrow$ via some action in A within time interval T to a state in S' . To extend this to a probability measure on paths, we now assume an *initial distribution* $\nu \in \text{Distr}(\mathcal{S})$ for the probability to start in a certain state s ; instead of $\nu(\{s\})$ we also write $\nu(s)$.

Definition 4 (Probability measure [31]). For initial distribution $\nu \in \text{Distr}(\mathcal{S})$ the probability measure on $\mathfrak{F}_{\text{Paths}^n}$ is defined inductively:

$$\begin{aligned} Pr_{\nu, \mathcal{D}}^0 : \mathfrak{F}_{\text{Paths}^0} &\rightarrow [0, 1] : \Pi \mapsto \sum_{s \in \Pi} \nu(s) \quad \text{and for } n > 0 \\ Pr_{\nu, \mathcal{D}}^n : \mathfrak{F}_{\text{Paths}^n} &\rightarrow [0, 1] : \Pi \mapsto \int_{\text{Paths}^{n-1}} Pr_{\nu, \mathcal{D}}^{n-1}(d\pi) \int_{\Omega} \mathbf{I}_{\Pi}(\pi \circ m) \mu_{\mathcal{D}}(\pi, dm). \end{aligned}$$

By Def. 4 we obtain measures on all σ -fields $\mathfrak{F}_{\text{Paths}^n}$. This extends to a measure on $(\text{Paths}^{\omega}, \mathfrak{F}_{\text{Paths}^{\omega}})$ as follows: First, note that any measurable cylinder can be represented by a base of finite length, i.e. $C_n = \{\pi \in \text{Paths}^{\omega} \mid \pi[0..n] \in C^n\}$. Now the measures $Pr_{\nu, \mathcal{D}}^n$ on $\mathfrak{F}_{\text{Paths}^n}$ extend to a unique probability measure $Pr_{\nu, \mathcal{D}}^{\omega}$ on $\mathfrak{F}_{\text{Paths}^{\omega}}$ by defining $Pr_{\nu, \mathcal{D}}^{\omega}(C_n) = Pr_{\nu, \mathcal{D}}^n(C^n)$. Although any measurable rectangle with base C^m can equally be represented by a higher-dimensional base (more precisely, if $m < n$ and $C^n = C^m \times \Omega^{n-m}$ then $C_n = C_m$), the Ionescu–Tulcea extension theorem [2] is applicable due to the inductive definition of the measures $Pr_{\nu, \mathcal{D}}^n$ and assures the extension to be well defined and unique.

Definition 4 inductively *appends* transition triples to the path prefixes of length n to obtain a measure on sets of paths of length $n + 1$. In the proof of Theorem 3, we use an equivalent characterization that constructs paths reversely, i.e. paths of length $n + 1$ are obtained from paths of length n by concatenating an *initial triple* from the set $\mathcal{S} \times \text{Act} \times \mathbb{R}_{\geq 0}$ to the suffix of length n :

Definition 5 (Initial triples). Let $\mathcal{C} = (\mathcal{S}, \text{Act}, \mathbf{R}, AP, L)$ be a CTMDP, $\nu \in \text{Distr}(\mathcal{S})$ and \mathcal{D} a scheduler. Then the measure $\mu_{\nu, \mathcal{D}} : \mathfrak{F}_{\mathcal{S} \times \text{Act} \times \mathbb{R}_{\geq 0}} \rightarrow [0, 1]$ on sets I of initial triples (s, α, t) is defined as

$$\mu_{\nu, \mathcal{D}}(I) = \int_{\mathcal{S}} \nu(ds) \int_{\text{Act}} \mathcal{D}(s, d\alpha) \int_{\mathbb{R}_{\geq 0}} \mathbf{I}_I(s, \alpha, t) \eta_{E(s, \alpha)}(dt).$$

This allows to decompose a path $\pi = s_0 \xrightarrow{\alpha_0, t_0} \dots \xrightarrow{\alpha_{n-1}, t_{n-1}} s_n$ into an initial triple $i = (s_0, \alpha_0, t_0)$ and the path suffix $\pi[1..n]$. For this to be measure preserving, a new $\nu_i \in \text{Distr}(\mathcal{S})$ is defined based on the original initial distribution ν of $Pr_{\nu, \mathcal{D}}^n$ on $\mathfrak{F}_{\text{Paths}^n}$ which reflects the fact that state s_0 has already been left with action α_0 at time t_0 . Hence ν_i is the initial distribution for the suffix-measure on $\mathfrak{F}_{\text{Paths}^{n-1}}$. Similarly, a scheduler \mathcal{D}_i is defined which reproduces the decisions of the original scheduler \mathcal{D} given that the first i -step is already taken. Hence $Pr_{\nu_i, \mathcal{D}_i}^{n-1}$ is the adjusted probability measure on $\mathfrak{F}_{\text{Paths}^{n-1}}$ given ν_i and \mathcal{D}_i .

Lemma 1. For $n \geq 1$ let $I \times \Pi \in \mathfrak{F}_{Paths^n}$ be a measurable rectangle, where $I \in \mathfrak{F}_{\mathcal{S}} \times \mathfrak{F}_{Act} \times \mathfrak{B}(\mathbb{R}_{\geq 0})$. For $i = (s, \alpha, t) \in I$, let $\nu_i := \mathbf{P}(s, \alpha, \cdot)$ and $\mathcal{D}_i(\pi) := \mathcal{D}(i \circ \pi)$. Then $Pr_{\nu, \mathcal{D}}^n(I \times \Pi) = \int_I Pr_{\nu_i, \mathcal{D}_i}^{n-1}(\Pi) \mu_{\nu, \mathcal{D}}(di)$.

Proof. By induction on n :

– induction start ($n = 1$): Let $\Pi \in \mathfrak{F}_{Paths^0}$, i.e. $\Pi \subseteq \mathcal{S}$.

$$\begin{aligned}
Pr_{\nu, \mathcal{D}}^1(I \times \Pi) &= \int_{Paths^0} Pr_{\nu, \mathcal{D}}^0(d\pi) \int_{\Omega} \mathbf{I}_{I \times \Pi}(\pi \circ m) \mu_{\mathcal{D}}(\pi, dm) && (* \text{ Definition 4 } *) \\
&= \int_{\mathcal{S}} \nu(ds_0) \int_{\Omega} \mathbf{I}_{I \times \Pi}(s_0 \circ m) \mu_{\mathcal{D}}(s_0, dm) && (* Paths^0 = \mathcal{S} *) \\
&= \int_{\mathcal{S}} \nu(ds_0) \int_{Act} \mathcal{D}(s_0, d\alpha_0) \int_{\mathbb{R}_{\geq 0}} \eta_{E(s_0, \alpha_0)}(dt_0) \int_{\mathcal{S}} \mathbf{I}_{I \times \Pi}(s_0 \xrightarrow{\alpha_0, t_0} s_1) \mathbf{P}(s_0, \alpha_0, ds_1) \\
&= \int_I \mu_{\nu, \mathcal{D}}(ds_0, d\alpha_0, dt_0) \int_{\mathcal{S}} \mathbf{I}_{\Pi}(s_1) \mathbf{P}(s_0, \alpha_0, ds_1) && (* \text{ definition of } \mu_{\nu, \mathcal{D}} *) \\
&= \int_I \mu_{\nu, \mathcal{D}}(di) \int_{\mathcal{S}} \mathbf{I}_{\Pi}(s_1) \nu_i(ds_1) && (* i = (s_0, \alpha_0, t_0) *) \\
&= \int_I Pr_{\nu_i, \mathcal{D}_i}^0(\Pi) \mu_{\nu, \mathcal{D}}(di). && (* \text{ Definition 4 } *)
\end{aligned}$$

– induction step ($n > 1$): Let $I \times \Pi \times M$ be a measurable rectangle in $\mathfrak{F}_{Paths^{n+1}}$ such that $I \in \mathfrak{F}_{\mathcal{S}} \times \mathfrak{F}_{Act} \times \mathfrak{B}(\mathbb{R}_{\geq 0})$ is a set of initial triples, $\Pi \in \mathfrak{F}_{Paths^{n-1}}$ and $M \in \mathfrak{F}$ is a set of combined transitions. Using the induction hypothesis $Pr_{\nu, \mathcal{D}}^n(I \times \Pi) = \int_I Pr_{\nu_i, \mathcal{D}_i}^{n-1}(\Pi) \mu_{\nu, \mathcal{D}}(di)$ we derive:

$$\begin{aligned}
Pr_{\nu, \mathcal{D}}^{n+1}(I \times \Pi \times M) &= \int_{I \times \Pi} \mu_{\mathcal{D}}(\pi, M) Pr_{\nu, \mathcal{D}}^n(d\pi) && (* \text{ Definition 4 } *) \\
&= \int_{I \times \Pi} \mu_{\mathcal{D}}(i \circ \pi', M) Pr_{\nu, \mathcal{D}}^n(d(i \circ \pi')) && (* \pi \simeq i \circ \pi' *) \\
&= \int_I \int_{\Pi} \mu_{\mathcal{D}}(i \circ \pi', M) Pr_{\nu_i, \mathcal{D}_i}^{n-1}(d\pi') \mu_{\nu, \mathcal{D}}(di) && (* \text{ ind. hypothesis } *) \\
&= \int_I \int_{\Pi} \mu_{\mathcal{D}_i}(\pi', M) Pr_{\nu_i, \mathcal{D}_i}^{n-1}(d\pi') \mu_{\nu, \mathcal{D}}(di) && (* \text{ definition of } \mathcal{D}_i *) \\
&= \int_I Pr_{\nu_i, \mathcal{D}_i}^n(\Pi \times M) \mu_{\nu, \mathcal{D}}(di). && (* \text{ Definition 4 } *)
\end{aligned}$$

□

A class of pathological paths that are not ruled out by Def. 2 are infinite paths whose duration converges to some real constant, i.e. paths that visit infinitely many states in a finite amount of time. For $n = 0, 1, 2, \dots$, an increasing sequence $r_n \in \mathbb{R}_{>0}$ is *Zeno* if it converges to a positive real number. For example, $r_n := \sum_{i=1}^n \frac{1}{2^i}$ converges to 1, hence is Zeno. The following theorem justifies to rule out such Zeno behaviour:

Theorem 1 (Converging paths theorem). *The probability measure of the set of converging paths is zero.*

Proof. Let $ConvPaths := \{s_0 \xrightarrow{\alpha_0, t_0} s_1 \xrightarrow{\alpha_1, t_1} \dots \mid \sum_{i=0}^n t_i \text{ converges}\}$. Then for $\pi \in ConvPaths$ the sequence t_i converges to 0. Thus there exists $k \in \mathbb{N}$ such that

$t_i \leq 1$ for all $i \geq k$. Hence $ConvPaths \subseteq \bigcup_{k \in \mathbb{N}} \mathcal{S} \times \Omega^k \times (Act \times [0, 1] \times \mathcal{S})^\omega$. Similar to [5, Prop. 1], it can be shown that $Pr_{\nu, \mathcal{D}}^\omega(\mathcal{S} \times \Omega^k \times (Act \times [0, 1] \times \mathcal{S})^\omega) = 0$ for all $k \in \mathbb{N}$. Thus also $Pr_{\nu, \mathcal{D}}^\omega(\bigcup_{k \in \mathbb{N}} \mathcal{S} \times \Omega^k \times (Act \times [0, 1] \times \mathcal{S})^\omega) = 0$. $ConvPaths$ is a subset of a set of measure zero; hence, on $\mathfrak{F}_{Paths^\omega}$ completed⁶ w.r.t. $Pr_{\nu, \mathcal{D}}^\omega$ we obtain $Pr_{\nu, \mathcal{D}}^\omega(ConvPaths) = 0$. \square

3 Strong bisimulation

Strong bisimulation [8,23] is an equivalence on the set of states of a CTMDP which relates two states if they are equally labelled and exhibit the same stepwise behaviour. As shown in Theorem 4, strong bisimilarity allows one to aggregate the state space while preserving transient and long run measures.

In the following we denote the equivalence class of s under equivalence $\mathcal{R} \subseteq \mathcal{S} \times \mathcal{S}$ by $[s]_{\mathcal{R}} = \{s' \in \mathcal{S} \mid (s, s') \in \mathcal{R}\}$; if \mathcal{R} is clear from the context we also write $[s]$. Further, $\mathcal{S}_{\mathcal{R}} := \{[s]_{\mathcal{R}} \mid s \in \mathcal{S}\}$ is the quotient space of \mathcal{S} under \mathcal{R} .

Definition 6 (Strong bisimulation relation). *Let $\mathcal{C} = (\mathcal{S}, Act, \mathbf{R}, AP, L)$ be a CTMDP. An equivalence $\mathcal{R} \subseteq \mathcal{S} \times \mathcal{S}$ is a strong bisimulation relation if $L(u) = L(v)$ for all $(u, v) \in \mathcal{R}$ and $\mathbf{R}(u, \alpha, C) = \mathbf{R}(v, \alpha, C)$ for all $\alpha \in Act$ and all $C \in \mathcal{S}_{\mathcal{R}}$.*

Two states u and v are strongly bisimilar ($u \sim v$) if there exists a strong bisimulation relation \mathcal{R} such that $(u, v) \in \mathcal{R}$. Strong bisimilarity is the union of all strong bisimulation relations.

Formally, $\sim = \{(u, v) \in \mathcal{S} \times \mathcal{S} \mid \exists \text{ str. bisimulation rel. } \mathcal{R} \text{ with } (u, v) \in \mathcal{R}\}$ defines strong bisimilarity which itself is (the largest) strong bisimulation relation.

Definition 7 (Quotient). *Let $\mathcal{C} = (\mathcal{S}, Act, \mathbf{R}, AP, L)$ be a CTMDP. Then $\tilde{\mathcal{C}} := (\tilde{\mathcal{S}}, Act, \tilde{\mathbf{R}}, AP, \tilde{L})$ where $\tilde{\mathcal{S}} := \mathcal{S}_{\sim}$, $\tilde{\mathbf{R}}([s], \alpha, C) := \mathbf{R}(s, \alpha, C)$ and $\tilde{L}([s]) := L(s)$ for all $s \in \mathcal{S}$, $\alpha \in Act$ and $C \in \tilde{\mathcal{S}}$ is the quotient of \mathcal{C} under strong bisimilarity.*

To distinguish between a CTMDP \mathcal{C} and its quotient, let $\tilde{\mathbf{P}}$ denote the quotient's discrete branching probabilities and \tilde{E} its exit rates. Note however, that exit rates and branching probabilities are preserved by strong bisimilarity, i.e. $E(s, \alpha) = \tilde{E}([s], \alpha)$ and $\tilde{\mathbf{P}}([s], \alpha, [t]) = \sum_{t' \in [t]} \mathbf{P}(s, \alpha, t')$ for $\alpha \in Act$ and $s, t \in \mathcal{S}$.

Example 2. Consider the CTMDP over the set $AP = \{a\}$ of atomic propositions in Fig. 2(a). Its quotient under strong bisimilarity is outlined in Fig. 2(b).

4 Continuous Stochastic Logic

Continuous stochastic logic [3,5] is a state-based logic to reason about continuous-time Markov chains. In this context, its formulas characterize strong bisimilarity [16] as defined in [5]; moreover, strongly bisimilar states satisfy the same CSL formulas [5]. In this paper, we extend CSL to CTMDPs along the lines of [6] and further introduce a long-run average operator [15]. Our semantics is based on ideas from [9,11] where variants of PCTL are extended to (discrete time) MDPs.

⁶ We may assume $\mathfrak{F}_{Paths^\omega}$ to be complete, see [2, p. 18ff].




Fig. 2. Quotient under strong bisimilarity.

4.1 Syntax and Semantics

Definition 8 (CSL syntax). For $a \in AP$, $p \in [0, 1]$, $I \subseteq \mathbb{R}_{\geq 0}$ a nonempty interval and $\sqsubseteq \in \{<, \leq, \geq, >\}$, CSL state and CSL path formulas are defined by

$$\Phi ::= a \mid \neg\Phi \mid \Phi \wedge \Psi \mid \forall^{\sqsubseteq p} \varphi \mid \mathbb{L}^{\sqsubseteq p} \Phi \quad \text{and} \quad \varphi ::= X^I \Phi \mid \Phi U^I \Psi.$$

The Boolean connectives \vee and \rightarrow are defined as usual; further we extend the syntax by deriving the timed modal operators “eventually” and “always” using the equalities $\diamond^I \Phi \equiv \text{tt} U^I \Phi$ and $\square^I \Phi \equiv \neg \diamond^I \neg \Phi$ where $\text{tt} := a \vee \neg a$ for some $a \in AP$. Similarly, the equality $\exists^{\sqsubseteq p} \varphi \equiv \neg \forall^{\sqsupset p} \varphi$ defines an existentially quantified transient state operator.

Example 3. Reconsider the CTMDP from Fig. 2(a). The *transient state formula* $\forall^{>0.1} \diamond^{[0,1]} a$ states that the probability to reach an a -labelled state within at most one time unit exceeds 0.1 no matter how the nondeterministic choices in the current state are resolved. Further, the *long-run average formula* $\mathbb{L}^{<0.25} \neg a$ states that for all scheduling decisions, the system spends less than 25% of its execution time in non- a states, on average.

Formally the long-run average is derived as follows: For $B \subseteq \mathcal{S}$, let \mathbf{I}_B denote an indicator with $\mathbf{I}_B(s) = 1$ if $s \in B$ and 0 otherwise. Following the ideas of [15,24], we compute the fraction of time spent in states from the set B on an infinite path π up to time bound $t \in \mathbb{R}_{\geq 0}$ and define $\text{avg}_{B,t}(\pi) = \frac{1}{t} \int_0^t \mathbf{I}_B(\pi @ t') dt'$. As $\text{avg}_{B,t}$ is a random variable, its expectation can be derived given an initial distribution $\nu \in \text{Distr}(\mathcal{S})$ and a measurable scheduler $\mathcal{D} \in \text{THR}$, i.e. $E(\text{avg}_{B,t}) = \int_{\text{Paths}^\omega} \text{avg}_{B,t}(\pi) Pr_{\nu, \mathcal{D}}^\omega(d\pi)$. Having the expectation for fixed time bound t , we now let $t \rightarrow \infty$ and obtain the long-run average as $\lim_{t \rightarrow \infty} E(\text{avg}_{B,t})$.

Definition 9 (CSL semantics). Let $\mathcal{C} = (\mathcal{S}, \text{Act}, \mathbf{R}, AP, L)$ be a CTMDP, $s, t \in \mathcal{S}$, $a \in AP$, $\sqsubseteq \in \{<, \leq, \geq, >\}$ and $\pi \in \text{Paths}^\omega$. Further let $\nu_s(t) := 1$ if $s = t$ and 0 otherwise. The semantics of state formulas is defined by

$$\begin{aligned} s \models a &\iff a \in L(s) \\ s \models \neg\Phi &\iff \text{not } s \models \Phi \\ s \models \Phi \wedge \Psi &\iff s \models \Phi \text{ and } s \models \Psi \\ s \models \forall^{\sqsubseteq p} \varphi &\iff \forall \mathcal{D} \in \text{THR}. Pr_{\nu_s, \mathcal{D}}^\omega \{ \pi \in \text{Paths}^\omega \mid \pi \models \varphi \} \sqsubseteq p \\ s \models \mathbb{L}^{\sqsubseteq p} \Phi &\iff \forall \mathcal{D} \in \text{THR}. \lim_{t \rightarrow \infty} \int_{\text{Paths}^\omega} \text{avg}_{\text{Sat}(\Phi), t}(\pi) Pr_{\nu_s, \mathcal{D}}^\omega(d\pi) \sqsubseteq p. \end{aligned}$$

Path formulas are defined by

$$\begin{aligned}\pi &\models \mathbf{X}^I \Phi \iff \pi[1] \models \Phi \wedge \delta(\pi, 0) \in I \\ \pi &\models \Phi \mathbf{U}^I \Psi \iff \exists t \in I. (\pi @ t \models \Psi \wedge (\forall t' \in [0, t). \pi @ t' \models \Phi))\end{aligned}$$

where $Sat(\Phi) := \{s \in \mathcal{S} \mid s \models \Phi\}$ and $\delta(\pi, n)$ is the time spent in state $\pi[n]$.

In Def. 9 the transient-state operator $\forall^{\overline{I}^p} \varphi$ is based on the measure of the set of paths that satisfy φ . For this to be well defined we must show that the set $\{\pi \in Paths^\omega \mid \pi \models \varphi\}$ is measurable:

Theorem 2 (Measurability of path formulas). *For any CSL path formula φ the set $\{\pi \in Paths^\omega \mid \pi \models \varphi\}$ is measurable.*

Proof. For next formulas, the proof is straightforward. For until formulas, let $\pi = s_0 \xrightarrow{\alpha_0, t_0} s_1 \xrightarrow{\alpha_1, t_1} \dots \in Paths^\omega$ and assume $\pi \models \Phi \mathbf{U}^I \Psi$. By Def. 9 it holds $\pi \models \Phi \mathbf{U}^I \Psi$ iff $\exists t \in I. (\pi @ t \models \Psi \wedge \forall t' \in [0, t). \pi @ t' \models \Phi)$. As we may exclude Zeno behaviour by Theorem 1, there exists $n \in \mathbb{N}$ with $\pi @ t = \pi[n] = s_n$ such that I and the period of time $[\sum_{i=0}^{n-1} t_i, \sum_{i=0}^n t_i)$ spent in state s_n overlap; further $s_n \models \Psi$ and $s_i \models \Phi$ for $i = 0, \dots, n-1$. Note however, that s_n must also satisfy Φ except for the case of *instantaneous arrival* where $\sum_{i=0}^{n-1} t_i \in I$. Accordingly, the set $\{\pi \in Paths^\omega \mid \pi \models \Phi \mathbf{U}^I \Psi\}$ can be represented by the union

$$\bigcup_{n=0}^{\infty} \left\{ \pi \in Paths^\omega \mid \sum_{i=0}^{n-1} t_i \in I \wedge \pi[n] \models \Psi \wedge \forall m < n. \pi[m] \models \Phi \right\} \quad (2)$$

$$\cup \bigcup_{n=0}^{\infty} \left\{ \pi \in Paths^\omega \mid \left(\sum_{i=0}^{n-1} t_i, \sum_{i=0}^n t_i \right) \cap I \neq \emptyset \wedge \pi[n] \models \Psi \wedge \forall m \leq n. \pi[m] \models \Phi \right\}. \quad (3)$$

It suffices to show that the subsets of (2) and (3) induced by any $n \in \mathbb{N}$ are measurable cylinders. In the following, we exhibit the proof for (3) and closed intervals $I = [a, b]$ as the other cases are similar. For fixed $n \geq 0$ we show that the corresponding cylinder base is measurable using a discretization argument:

$$\begin{aligned}& \left\{ \pi \in Paths^{n+1} \mid \left(\sum_{i=0}^{n-1} t_i, \sum_{i=0}^n t_i \right) \cap [a, b] \neq \emptyset \wedge \pi[n] \models \Psi \wedge \forall m \leq n. \pi[m] \models \Phi \right\} \\ &= \bigcup_{k=1}^{\infty} \bigcup_{\substack{c_0 + \dots + c_n \geq ak \\ d_0 + \dots + d_{n-1} \leq bk \\ c_i < d_i}} \prod_{i=0}^{n-1} \left[Sat(\Phi) \times Act \times \left(\frac{c_i}{k}, \frac{d_i}{k} \right) \right] \times Sat(\Phi \wedge \Psi) \times Act \times \left(\frac{c_n}{k}, \infty \right) \times \mathcal{S} \quad (4)\end{aligned}$$

where $c_i, d_j \in \mathbb{N}$. To shorten notation, let $c := \sum_{i=0}^{n-1} t_i$ and $d := \sum_{i=0}^n t_i$.

\subseteq : Let $\pi = s_0 \xrightarrow{\alpha_0, t_0} s_1 \xrightarrow{\alpha_1, t_1} \dots \xrightarrow{\alpha_n, t_n} s_{n+1}$ be in the set on the left-hand side of equation (4). The intervals (c, d) and $[a, b]$ overlap, hence $c < b$ and $d > a$ (see top of Fig. 3). Further $\pi[i] \models \Phi$ for $i = 0, \dots, n$ and $\pi[n] \models \Psi$. To show that π is in the set on the right-hand side, let $c_i = \lceil t_i \cdot k - 1 \rceil$ and $d_i = \lfloor t_i \cdot k + 1 \rfloor$ for $k > 0$. Then $\frac{c_i}{k} < t_i < \frac{d_i}{k}$ approximates the sojourn times t_i as depicted in




Fig. 3. Discretization of intervals with $n = 4$ and $I = (a, b)$.

Fig. 3. Further let $\varepsilon = \sum_{i=0}^n t_i - a$ and choose k_0 such that $\frac{n+1}{k_0} \leq \varepsilon$ to obtain

$$a = \sum_{i=0}^n t_i - \varepsilon \leq \sum_{i=0}^n t_i - \frac{n+1}{k_0} \leq \sum_{i=0}^n \frac{c_i + 1}{k_0} - \frac{n+1}{k_0} = \sum_{i=0}^n \frac{c_i}{k_0}.$$

Thus $ak \leq \sum_{i=0}^n c_i$ for all $k \geq k_0$. Similarly, we obtain $k'_0 \in \mathbb{N}$ s.t. $\sum_{i=0}^{n-1} d_i \leq bk$ for all $k \geq k'_0$. Hence for large k , π is in the set on the right-hand side.

\supseteq : Let π be in the set on the right-hand side of equation (4) with corresponding values for c_i, d_i and k . Then $t_i \in (\frac{c_i}{k}, \frac{d_i}{k})$. Hence $a \leq \sum_{i=0}^n \frac{c_i}{k} < \sum_{i=0}^n t_i = c$ and $b \geq \sum_{i=0}^{n-1} \frac{d_i}{k} > \sum_{i=0}^{n-1} t_i = c$ so that the time-interval (c, d) of state s_n and the time interval $I = [a, b]$ of the formula overlap. Further, $\pi[m] \models \Phi$ for $m \leq n$ and $\pi[n] \models \Psi$; thus π is in the set on the left-hand side of equation (4).

The right-hand side of equation (4) is measurable, hence also the cylinder base. This extends to its cylinder and the countable union in equation (3). \square

4.2 Strong bisimilarity preserves CSL

We now prepare the main result of our paper. To prove that strong bisimilarity preserves CSL formulas we establish a correspondence between certain sets of paths of a CTMDP and its quotient which is measure-preserving:

Definition 10 (Simple bisimulation closed). *Let $\mathcal{C} = (\mathcal{S}, Act, \mathbf{R}, AP, L)$ be a CTMDP. A measurable rectangle $\Pi = S_0 \times A_0 \times T_0 \times \cdots \times A_{n-1} \times T_{n-1} \times S_n$ is simple bisimulation closed if $S_i \in (\tilde{\mathcal{S}} \cup \{\emptyset\})$ for $i = 0, \dots, n$. Further, let $\tilde{\Pi} = \{S_0\} \times A_0 \times T_0 \times \cdots \times A_{n-1} \times T_{n-1} \times \{S_n\}$ be the corresponding rectangle in the quotient $\tilde{\mathcal{C}}$.*

An essential step in our proof strategy is to obtain a scheduler on the quotient. The following example illustrates the intuition for such a scheduler.

Example 4. Let \mathcal{C} be the CTMDP in Fig. 4(a) where $\nu(s_0) = \frac{1}{4}$, $\nu(s_1) = \frac{2}{3}$ and $\nu(s_2) = \frac{1}{12}$. Assume a scheduler \mathcal{D} where $\mathcal{D}(s_0, \{\alpha\}) = \frac{2}{3}$, $\mathcal{D}(s_0, \{\beta\}) = \frac{1}{3}$, $\mathcal{D}(s_1, \{\alpha\}) = \frac{1}{4}$ and $\mathcal{D}(s_1, \{\beta\}) = \frac{3}{4}$. Intuitively, a scheduler \mathcal{D}' that mimics \mathcal{D} 's




Fig. 4. Derivation of the quotient scheduler.

behaviour on the quotient $\tilde{\mathcal{C}}$ in Fig. 4(b) can be defined by

$$\mathcal{D}_{\sim}^{\nu}([s_0], \{\alpha\}) = \frac{\sum_{s \in [s_0]} \nu(s) \cdot \mathcal{D}(s, \{\alpha\})}{\sum_{s \in [s_0]} \nu(s)} = \frac{\frac{1}{4} \cdot \frac{2}{3} + \frac{2}{3} \cdot \frac{1}{4}}{\frac{1}{4} + \frac{2}{3}} = \frac{4}{11} \quad \text{and}$$

$$\mathcal{D}_{\sim}^{\nu}([s_0], \{\beta\}) = \frac{\sum_{s \in [s_0]} \nu(s) \cdot \mathcal{D}(s, \{\beta\})}{\sum_{s \in [s_0]} \nu(s)} = \frac{\frac{1}{4} \cdot \frac{1}{3} + \frac{2}{3} \cdot \frac{3}{4}}{\frac{1}{4} + \frac{2}{3}} = \frac{7}{11}.$$

Even though s_0 and s_1 are bisimilar, the scheduler \mathcal{D} decides differently for the histories $\pi_0 = s_0$ and $\pi_1 = s_1$. As π_0 and π_1 collapse into $\tilde{\pi} = [s_0]$ on the quotient, \mathcal{D}_{\sim}^{ν} can no longer distinguish between π_0 and π_1 . Therefore \mathcal{D} 's decision for any history $\pi \in \tilde{\pi}$ is weighed w.r.t. the total probability of $\tilde{\pi}$.

Definition 11 (Quotient scheduler). Let $\mathcal{C} = (\mathcal{S}, Act, \mathbf{R}, AP, L)$ be a CTMDP, $\nu \in \text{Distr}(\mathcal{S})$ and $\mathcal{D} \in \text{THR}$. First, define the history weight of finite paths of length n inductively as follows:

$$hw_0(\nu, \mathcal{D}, s_0) := \nu(s_0) \text{ and}$$

$$hw_{n+1}(\nu, \mathcal{D}, \pi \xrightarrow{\alpha_n, t_n} s_{n+1}) := hw_n(\nu, \mathcal{D}, \pi) \cdot \mathcal{D}(\pi, \{\alpha_n\}) \cdot \mathbf{P}(\pi \downarrow, \alpha_n, s_{n+1}).$$

Let $\tilde{\pi} = [s_0] \xrightarrow{\alpha_0, t_0} \dots \xrightarrow{\alpha_{n-1}, t_{n-1}} [s_n]$ be a timed history of $\tilde{\mathcal{C}}$ and $\Pi = [s_0] \times \{\alpha_0\} \times \{t_0\} \times \dots \times \{\alpha_{n-1}\} \times \{t_{n-1}\} \times [s_n]$ be the corresponding set of paths in \mathcal{C} . The quotient scheduler \mathcal{D}_{\sim}^{ν} on $\tilde{\mathcal{C}}$ is then defined as follows:

$$\mathcal{D}_{\sim}^{\nu}(\tilde{\pi}, \alpha_n) := \frac{\sum_{\pi \in \Pi} hw_n(\nu, \mathcal{D}, \pi) \cdot \mathcal{D}(\pi, \{\alpha_n\})}{\sum_{\pi \in \Pi} hw_n(\nu, \mathcal{D}, \pi)}.$$

Further, let $\tilde{\nu}([s]) := \sum_{s' \in [s]} \nu(s')$ be the initial distribution on $\tilde{\mathcal{C}}$.

A history $\tilde{\pi}$ of $\tilde{\mathcal{C}}$ corresponds to a set of paths Π in \mathcal{C} ; given $\tilde{\pi}$, the quotient scheduler decides by multiplying \mathcal{D} 's decision on each path in Π with its corresponding weight and normalizing with the weight of Π afterwards. Now we obtain a first intermediate result: For CTMDP \mathcal{C} , if Π is a simple bisimulation closed set of paths, ν an initial distribution and $\mathcal{D} \in \text{THR}$, the measure of Π in \mathcal{C} coincides with the measure of $\tilde{\Pi}$ in $\tilde{\mathcal{C}}$ which is induced by $\tilde{\nu}$ and \mathcal{D}_{\sim}^{ν} :

Theorem 3. *Let \mathcal{C} be a CTMDP with set of states \mathcal{S} and $\nu \in \text{Distr}(\mathcal{S})$. Then $Pr_{\nu, \mathcal{D}}^\omega(\Pi) = Pr_{\tilde{\nu}, \mathcal{D}^\nu}^\omega(\tilde{\Pi})$ where $\mathcal{D} \in \text{THR}$ and Π simple bisimulation closed.*

Proof. By induction on the length n of cylinder bases. The induction base holds for $Pr_{\nu, \mathcal{D}}^0([s]) = \sum_{s' \in [s]} \nu(s') = \tilde{\nu}([s]) = Pr_{\tilde{\nu}, \mathcal{D}^\nu}^0(\{[s]\})$. With the induction hypothesis that $Pr_{\nu, \mathcal{D}}^n(\Pi) = Pr_{\tilde{\nu}, \mathcal{D}^\nu}^n(\tilde{\Pi})$ for all $\nu \in \text{Distr}(\mathcal{S})$, $\mathcal{D} \in \text{THR}$ and bisimulation closed $\Pi \subseteq \text{Paths}^n$ we obtain the induction step:

$$\begin{aligned}
Pr_{\nu, \mathcal{D}}^{n+1}([s_0] \times A_0 \times T_0 \times \Pi) &= \int_{[s_0] \times A_0 \times T_0} Pr_{\mathbf{P}(s, \alpha, \cdot), \mathcal{D}(s \xrightarrow{\alpha, t} \cdot)}^n(\Pi) \mu_{\nu, \mathcal{D}}(ds, d\alpha, dt) \\
&= \int_{s \in [s_0]} \nu(ds) \int_{\alpha \in A_0} \mathcal{D}(s, d\alpha) \int_{T_0} Pr_{\mathbf{P}(s, \alpha, \cdot), \mathcal{D}(s \xrightarrow{\alpha, t} \cdot)}^n(\Pi) \eta_{E(s, \alpha)}(dt) \\
&= \sum_{s \in [s_0]} \nu(s) \sum_{\alpha \in A_0} \mathcal{D}(s, \{\alpha\}) \int_{T_0} Pr_{\mathbf{P}(s, \alpha, \cdot), \mathcal{D}(s \xrightarrow{\alpha, t} \cdot)}^n(\Pi) \eta_{\tilde{E}([s_0], \alpha)}(dt) \\
&\stackrel{\text{i.h.}}{=} \sum_{s \in [s_0]} \sum_{\alpha \in A_0} \int_{T_0} Pr_{\tilde{\mathbf{P}}([s_0], \alpha, \cdot), \mathcal{D}^\nu([s_0] \xrightarrow{\alpha, t} \cdot)}^n(\tilde{\Pi}) \cdot \nu(s) \cdot \mathcal{D}(s, \{\alpha\}) \eta_{\tilde{E}([s_0], \alpha)}(dt) \\
&= \sum_{\alpha \in A_0} \int_{T_0} Pr_{\tilde{\mathbf{P}}([s_0], \alpha, \cdot), \mathcal{D}^\nu([s_0] \xrightarrow{\alpha, t} \cdot)}^n(\tilde{\Pi}) \cdot \sum_{s \in [s_0]} (\nu(s) \cdot \mathcal{D}(s, \{\alpha\})) \eta_{\tilde{E}([s_0], \alpha)}(dt) \\
&= \sum_{\alpha \in A_0} \int_{T_0} Pr_{\tilde{\mathbf{P}}([s_0], \alpha, \cdot), \mathcal{D}^\nu([s_0] \xrightarrow{\alpha, t} \cdot)}^n(\tilde{\Pi}) \cdot \tilde{\nu}([s_0]) \cdot \mathcal{D}^\nu([s_0], \{\alpha\}) \eta_{\tilde{E}([s_0], \alpha)}(dt) \\
&= \int_{\{[s_0]\}} \tilde{\nu}(d[s]) \int_{A_0} \mathcal{D}^\nu([s], d\alpha) \int_{T_0} Pr_{\tilde{\mathbf{P}}([s], \alpha, \cdot), \mathcal{D}^\nu([s] \xrightarrow{\alpha, t} \cdot)}^n(\tilde{\Pi}) \eta_{\tilde{E}([s], \alpha)}(dt) \\
&= \int_{\{[s_0]\} \times A_0 \times T_0} Pr_{\tilde{\mathbf{P}}([s], \alpha, \cdot), \mathcal{D}^\nu([s] \xrightarrow{\alpha, t} \cdot)}^n(\tilde{\Pi}) \tilde{\mu}_{\tilde{\nu}, \mathcal{D}^\nu}(d[s], d\alpha, dt) \\
&= Pr_{\tilde{\nu}, \mathcal{D}^\nu}^{n+1}(\{[s_0]\} \times A_0 \times T_0 \times \tilde{\Pi})
\end{aligned}$$

where $\tilde{\mu}_{\tilde{\nu}, \mathcal{D}^\nu}$ is the extension of $\mu_{\nu, \mathcal{D}}$ (Def. 5) to sets of initial triples in $\tilde{\mathcal{C}}$:

$$\tilde{\mu}_{\tilde{\nu}, \mathcal{D}^\nu}: \mathfrak{F}_{\tilde{\mathcal{S}} \times \text{Act} \times \mathbb{R}_{\geq 0}} \rightarrow [0, 1] : I \mapsto \int_{\tilde{\mathcal{S}}} \tilde{\nu}(d[s]) \int_{\text{Act}} \mathcal{D}^\nu([s], d\alpha) \int_{\mathbb{R}_{\geq 0}} \mathbf{1}_I([s], \alpha, t) \eta_{\tilde{E}([s], \alpha)}(dt). \quad \square$$

According to Theorem 3, the quotient scheduler preserves the measure for *simple* bisimulation closed sets of paths, i.e. for paths, whose state components are equivalence classes under \sim . To generalize this to sets of paths that satisfy a CSL path formula, we introduce *general* bisimulation closed sets of paths:

Definition 12 (Bisimulation closed). *Let $\mathcal{C} = (\mathcal{S}, \text{Act}, \mathbf{R}, AP, L)$ be a CTMDP and $\tilde{\mathcal{C}}$ its quotient under strong bisimilarity. A measurable rectangle $\Pi = S_0 \times A_0 \times T_0 \times \cdots \times A_{n-1} \times T_{n-1} \times S_n$ is bisimulation closed if $S_i = \bigsqcup_{j=0}^{k_i} [s_{i,j}]$ for $k_i \in \mathbb{N}$ and $0 \leq i \leq n$. Let $\tilde{\Pi} = \bigcup_{j=0}^{k_0} \{[s_{0,j}]\} \times A_0 \times T_0 \times \cdots \times A_{n-1} \times T_{n-1} \times \bigcup_{j=0}^{k_n} \{[s_{n,j}]\}$ be the corresponding rectangle in the quotient $\tilde{\mathcal{C}}$.*

Lemma 2. *Any bisimulation closed set of paths Π can be represented as a finite disjoint union of simple bisimulation closed sets of paths.*

Proof. Direct consequence of Def. 12. \square

Corollary 1. *Let \mathcal{C} be a CTMDP with set of states \mathcal{S} and $\nu \in \text{Distr}(\mathcal{S})$ an initial distribution. Then $Pr_{\nu, \mathcal{D}}^\omega(\Pi) = Pr_{\tilde{\nu}, \mathcal{D}^\sim}^\omega(\tilde{\Pi})$ for any $\mathcal{D} \in \text{THR}$ and any bisimulation closed set of paths Π .*

Proof. Follows directly from Lemma 2 and Theorem 3. \square

Using these extensions we can now prove our main result:

Theorem 4. *Let \mathcal{C} be a CTMDP with set of states \mathcal{S} and $u, v \in \mathcal{S}$. Then $u \sim v$ implies $u \models \Phi$ iff $v \models \Phi$ for all CSL state formulas Φ .*

Proof. By structural induction on Φ . If $\Phi = a$ and $a \in AP$ the induction base follows as $L(u) = L(v)$. In the induction step, conjunction and negation are obvious.

Let $\Phi = \forall^{\square p} \varphi$ and $\Pi = \{\pi \in \text{Paths}^\omega \mid \pi \models \varphi\}$. To show $u \models \forall^{\square p} \varphi$ implies $v \models \forall^{\square p} \varphi$ it suffices to show that for any $\mathcal{V} \in \text{THR}$ there exists $\mathcal{U} \in \text{THR}$ with $Pr_{\nu_u, \mathcal{U}}^\omega(\Pi) = Pr_{\nu_v, \mathcal{V}}^\omega(\Pi)$. By Theorem 2 the set Π is measurable, hence $\Pi = \bigsqcup_{i=0}^\infty \Pi_i$ for disjoint $\Pi_i \in \mathfrak{F}_{\text{Paths}^\omega}$. By *induction hypothesis* for path formulas $X^I \Phi$ and $\Phi U^I \Psi$ the sets $\text{Sat}(\Phi)$ and $\text{Sat}(\Psi)$ are disjoint unions of \sim -equivalence classes. The same holds for any Boolean combination of Φ and Ψ . Hence $\Pi = \bigsqcup_{i=0}^\infty \Pi_i$ where the Π_i are bisimulation closed. For all $\mathcal{V} \in \text{THR}$ and $\pi = s_0 \xrightarrow{\alpha_0, t_0} \dots \xrightarrow{\alpha_{n-1}, t_{n-1}} s_n$ let $\mathcal{U}(\pi) := \mathcal{V}^{\nu_v}([s_0] \xrightarrow{\alpha_0, t_0} \dots \xrightarrow{\alpha_{n-1}, t_{n-1}} [s_n])$. Thus \mathcal{U} mimics on π the decision of \mathcal{V}^{ν_v} on $\tilde{\pi}$. In fact $\mathcal{U}^{\nu_u} = \mathcal{V}^{\nu_v}$ since

$$\mathcal{U}^{\nu_u}(\tilde{\pi}, \alpha_n) = \frac{\sum_{\pi \in \Pi} hw_n(\nu_u, \mathcal{U}, \pi) \cdot \mathcal{V}^{\nu_v}(\tilde{\pi}, \alpha_n)}{\sum_{\pi \in \Pi} hw_n(\nu_u, \mathcal{U}, \pi)}$$

and $\mathcal{V}^{\nu_v}(\tilde{\pi}, \alpha_n)$ is independent of π . With $\tilde{\nu}_u = \tilde{\nu}_v$ and by Corollary 1 we obtain $Pr_{\nu_u, \mathcal{U}}^\omega(\Pi_i) = Pr_{\tilde{\nu}_u, \mathcal{U}^{\nu_u}}^\omega(\tilde{\Pi}_i) = Pr_{\tilde{\nu}_v, \mathcal{V}^{\nu_v}}^\omega(\tilde{\Pi}_i) = Pr_{\nu_v, \mathcal{V}}^\omega(\Pi_i)$ which carries over to Π for Π is a countable union of disjoint sets Π_i .

Let $\Phi = \mathbb{L}^{\square p} \Psi$. Since $u \sim v$, it suffices to show that for all $s \in \mathcal{S}$ it holds $s \models \mathbb{L}^{\square p} \Psi$ iff $[s] \models \mathbb{L}^{\square p} \Psi$. The expectation of $avg_{\text{Sat}(\Psi), t}$ for $t \in \mathbb{R}_{\geq 0}$ can be expressed as follows:

$$\int_{\text{Paths}^\omega} \left(\frac{1}{t} \int_0^t \mathbf{1}_{\text{Sat}(\Psi)}(\pi @ t') dt' \right) Pr_{\nu_s, \mathcal{D}}^\omega(d\pi) = \frac{1}{t} \int_0^t Pr_{\nu_s, \mathcal{D}}^\omega\{\pi \in \text{Paths}^\omega \mid \pi @ t' \models \Psi\} dt'.$$

Further, the sets $\{\pi \in \text{Paths}^\omega \mid \pi @ t' \models \Psi\}$ and $\{\pi \in \text{Paths}^\omega \mid \pi \models \diamond^{[t', t']}\Psi\}$ have the same measure and the *induction hypothesis* applies to Ψ . Applying the previous reasoning for the until case to the formula $\text{tt } U^{[t', t']}\Psi$ once, we obtain

$$Pr_{\nu_s, \mathcal{D}}^\omega\{\pi \in \text{Paths}^\omega(\mathcal{C}) \mid \pi \models \diamond^{[t', t']}\Psi\} = Pr_{\tilde{\nu}_s, \mathcal{D}^{\nu_s}}^\omega\{\tilde{\pi} \in \text{Paths}^\omega(\tilde{\mathcal{C}}) \mid \tilde{\pi} \models \diamond^{[t', t']}\Psi\}$$

for all $t' \in \mathbb{R}_{\geq 0}$. Thus the expectations of $avg_{\text{Sat}(\Psi), t}$ on \mathcal{C} and $\tilde{\mathcal{C}}$ are equal for all $t \in \mathbb{R}_{\geq 0}$ and the same holds for their limits if $t \rightarrow \infty$. This completes the proof as for $u \sim v$ we obtain $u \models \mathbb{L}^{\square p} \Psi$ iff $[u] \models \mathbb{L}^{\square p} \Psi$ iff $[v] \models \mathbb{L}^{\square p} \Psi$ iff $v \models \mathbb{L}^{\square p} \Psi$. \square

This theorem shows that bisimilar states satisfy the same CSL formulas. The reverse direction, however, does not hold in general. One reason is obvious: In this paper we use a purely state-based logic whereas our definition of strong bisimulation also accounts for action names. Therefore it comes to no surprise that CSL cannot characterize strong bisimulation. However, there is another more profound reason which is analogous to the discrete-time setting where extensions of PCTL to Markov decision processes [29,4] also cannot express strong bisimilarity: CSL and PCTL only allow to specify infima and suprema as probability bounds under a denumerable class of randomized schedulers; therefore intuitively, CSL cannot characterize exponential distributions which neither contribute to the supremum nor to the infimum of the probability measures of a given set of paths. Thus the counterexample from [4, Fig 9.5] interpreted as a CTMDP applies verbatim to our case.

5 Conclusion

In this paper we define strong bisimulation on CTMDPs and propose a nondeterministic extension of CSL to CTMDP that allows to express a wide class of performance and dependability measures. Using a measure-theoretic argument we prove our logic to be well-defined. Our main contribution is the proof that strong bisimilarity preserves the validity of CSL formulas. However, our logic is not capable of characterizing strong bisimilarity. To this end, action-based logics provide a natural starting point.

Acknowledgements This research has been performed as part of the QUPES project that is financed by the Netherlands Organization for Scientific Research (NWO). Daniel Klink and David N. Jansen are kindly acknowledged for many fruitful discussions.

References

1. Abdeddaïm, Y., Asarin, E., Maler, O.: On optimal scheduling under uncertainty. In: *TACAS. LNCS*, Vol. 2619. Springer (2003) 240–253
2. Ash, R. B., Doléans-Dade, C. A.: *Probability & Measure Theory*. 2nd edn. Academic Press (2000)
3. Aziz, A., Sanwal, K., Singhal, V., Brayton, R. K.: Model-checking continuous-time Markov chains. *ACM Trans. Comput. Log.* **1** (2000) 162–170
4. Baier, C.: On Algorithmic Verification Methods for Probabilistic Systems. Habilitation Thesis (1998) University of Mannheim.
5. Baier, C., Haverkort, B. R., Hermanns, H., Katoen, J.-P.: Model-checking algorithms for continuous-time Markov chains. *IEEE TSE* **29** (2003) 524–541
6. Baier, C., Haverkort, B. R., Hermanns, H., Katoen, J.-P.: Nonuniform CTMDPs. unpublished manuscript (2004)
7. Baier, C., Hermanns, H., Katoen, J.-P., Haverkort, B. R.: Efficient computation of time-bounded reachability probabilities in uniform continuous-time Markov decision processes. *Theor. Comp. Sci.* **345** (2005) 2–26
8. Baier, C., Katoen, J.-P., Hermanns, H., Wolf, V.: Comparative branching-time semantics for Markov chains. *Information and Computation* **200** (2005) 149–214

9. Baier, C., Kwiatkowska, M. Z.: Model checking for a probabilistic branching time logic with fairness. *Distr. Comp.* **11** (1998) 125–155
10. Beutler, F. J., Ross, K. W.: Optimal policies for controlled Markov chains with a constraint. *Journal of Mathematical Analysis and Appl.* **112** (1985) 236–252
11. Bianco, A., de Alfaro, L.: Model checking of probabilistic and nondeterministic systems. In: *FSTTCS. LNCS*, Vol. 1026. Springer (1995) 499–513
12. Boudali, H., Crouzen, P., Stoelinga, M. I. A.: Dynamic fault tree analysis using input/output interactive Markov chains. In: *Dependable Systems and Networks*. IEEE (2007)
13. Buchholz, P.: Exact and ordinary lumpability in finite Markov chains. *Journal of Applied Probability* **31** (1994) 59–75
14. Chiola, G., Marsan, M. A., Balbo, G., Conte, G.: Generalized stochastic Petri nets: A definition at the net level and its implications. *IEEE TSE* **19** (1993) 89–107
15. de Alfaro, L.: *Formal Verification of Probabilistic Systems*. PhD thesis, Stanford University (1997)
16. Desharnais, J., Panangaden, P.: Continuous stochastic logic characterizes bisimulation of continuous-time Markov processes. *Journal of Logic and Algebraic Programming* **56** (2003) 99–115
17. Feinberg, E. A.: Continuous time discounted jump Markov decision processes: A discrete-event approach. *Mathematics of Operations Research* **29** (2004) 492–524
18. Givan, R., Dean, T., Greig, M.: Equivalence notions and model minimization in Markov decision processes. *Artificial Intelligence* **147** (2003) 163–223
19. Hermanns, H.: *Interactive Markov Chains: The Quest for Quantified Quality*. LNCS, Vol. 2428. Springer (2002)
20. Howard, R. A.: *Dynamic Probabilistic Systems*. John Wiley and Sons (1971)
21. Katoen, J.-P., Klink, D., Leucker, M., Wolf, V.: Three-valued abstraction for continuous-time Markov chains. In: *CAV*. LNCS. Springer (2007)
22. Kemeny, J. G., Snell, J. L., Knapp, A. W.: *Denumerable Markov Chains*. 2nd edn. Springer (1976)
23. Larsen, K. G., Skou, A.: Bisimulation through probabilistic testing. *Information and Computation* **94** (1991) 1–28
24. López, G. G. I., Hermanns, H., Katoen, J.-P.: Beyond memoryless distributions: Model checking semi-Markov chains. In: *PAPM-PROBMIV. LNCS*, Vol. 2165. Springer (2001) 57–70
25. Neuhäüßer, M. R., Katoen, J.-P.: Bisimulation and logical preservation for continuous-time markov decision processes. Technical Report 10, RWTH Aachen (2007)
26. Puterman, M. L.: *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley and Sons (1994)
27. Qiu, Q., Pedram, M.: Dynamic power management based on continuous-time Markov decision processes. In: *DAC*. ACM Press (1999) 555–561
28. Sanders, W. H., Meyer, J. F.: Stochastic activity networks: Formal definitions and concepts. In: *FMPA. LNCS*, Vol. 2090. Springer (2000) 315–343
29. Segala, R., Lynch, N.: Probabilistic simulations for probabilistic processes. *Nordic Journal of Computing* **2** (1995) 250–273
30. Vardi, M. Y.: Automatic verification of probabilistic concurrent finite-state programs. In: *FOCS*. IEEE (1985) 327–338
31. Wolovick, N., Johr, S.: A characterization of meaningful schedulers for continuous-time Markov decision processes. In: *FORMATS. LNCS*, Vol. 4202. Springer (2006) 352–367