

Bisimulation and Logical Preservation for Continuous-Time Markov Decision Processes

Martin Neuhäuser and Joost-Pieter Katoen

The publications of the Department of Computer Science of *RWTH Aachen University* are in general accessible through the World Wide Web.

<http://aib.informatik.rwth-aachen.de/>

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 [29] 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 [25], a continuous-time variant of MDPs, where state residence times are exponentially distributed, have received scant attention. Unlike in MDPs, where 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 [30,7,10] 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 [27] (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 [26]. 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.

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 α -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 $i < j \leq |\pi|$. 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. The set $Paths$ in general contains paths which do not comply with the rate matrix of the underlying CTMDP. However, the following definition of the probability measure (Def. 4) justifies this as it assigns probability zero to such 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 -field³ 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 first define a σ -field of sets of *combined transitions* which we later use to define σ -fields of sets of finite and infinite paths. Here, a combined transition is a tuple (α, t, s) which links the decision for action α (which is given by a scheduler, see Def. 3) with the exponentially distributed time-point t to move to a successor state s of the underlying CTMDP. Formally, for a CTMDP $\mathcal{C} = (\mathcal{S}, Act, \mathbf{R}, AP, L)$, the set of combined transitions is $\Omega = Act \times \mathbb{R}_{\geq 0} \times \mathcal{S}$. To define a probability space on Ω , note that \mathcal{S} and Act are finite; hence, the corresponding σ -fields are defined as $\mathfrak{F}_{Act} := 2^{Act}$ and $\mathfrak{F}_{\mathcal{S}} := 2^{\mathcal{S}}$. 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 $\mathbb{R}_{\geq 0}$. In the following, we denote the sets of probability distributions on \mathfrak{F}_{Act} and $\mathfrak{F}_{\mathcal{S}}$ by $Distr(Act)$ and $Distr(\mathcal{S})$, respectively. Note that any path $\pi = s_0 \xrightarrow{\alpha_0, t_0} s_1 \xrightarrow{\alpha_1, t_1} \dots \xrightarrow{\alpha_{n-1}, t_{n-1}} s_n$ of length n can be extended by a combined transition $m = (\alpha_n, t_n, s_{n+1})$ to a path of length $n + 1$, denoted $\pi \circ m$.

Generally, a Cartesian product is a *measurable rectangle* if its constituent sets are elements of their respective σ -fields. For example, in our case 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 $S \in \mathfrak{F}_{\mathcal{S}}$.

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

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. $\mathfrak{F}_{Paths^n} := \sigma(\{S_0 \times M_0 \times \dots \times M_n \mid S_0 \in \mathfrak{F}_{\mathcal{S}}, M_i \in \mathfrak{F}, 0 \leq i \leq n\})$. Intuitively, \mathfrak{F}_{Paths^n} consists of all possible (even countable infinite) unions and intersections of measurable path rectangles.

Example 2. For the CTMDP in Fig. 1, the event “from s_1 we directly reach state s_3 within 0.5 time units” and the event “if action α is chosen in state s_1 , we remain in s_1 for less than 0.2 or more than 1 time units” are described by the Cartesian products $\Pi_1 = \{s_1\} \times Act \times [0, 0.5] \times \{s_3\}$ and $\Pi_2 = \{s_1\} \times \{\alpha\} \times ([0, 0.2] \cup (1, \infty)) \times \mathcal{S}$. Π_1 and Π_2 are measurable rectangles whereas their union $\Pi_1 \cup \Pi_2$ is an element of the σ -field \mathfrak{F}_{Paths^2} .

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 called an *infinite 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 infinite rectangle*, if $S_i \in \mathfrak{F}_{\mathcal{S}}$, $A_i \in \mathfrak{F}_{Act}$ and $T_i \in \mathfrak{B}(\mathbb{R}_{\geq 0})$. We obtain the desired σ -field of sets of infinite paths as the minimal σ -field generated by the set of measurable cylinders; formally: $\mathfrak{F}_{Paths^\omega} := \sigma(\bigcup_{n=0}^{\infty} \{C_n \mid C^n \in \mathfrak{F}_{Paths^n}\})$.

Finally, the σ -field \mathfrak{F}_{Paths^*} over finite and infinite paths is the smallest σ -field generated by the disjoint union $\bigcup_{n=0}^{\infty} \mathfrak{F}_{Paths^n} \cup \mathfrak{F}_{Paths^\omega}$.

2.2 The probability measure

To define a semantics for CTMDPs we use schedulers⁵ to resolve the non-deterministic 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 [30]:

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 [30]. 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 $E(\pi \downarrow, \alpha)$.

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

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

To derive a probability measure on $\mathfrak{F}_{Paths^\omega}$, we first define a probability measure on combined transitions, i.e. on the measurable space (Ω, \mathfrak{F}) : For history $\pi \in Paths^*$, let $\mu_{\mathcal{D}}(\pi, \cdot) : \mathfrak{F} \rightarrow [0, 1]$ such that

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

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 [30]. Intuitively, for a given finite path π and a set M of combined transitions, $\mu_{\mathcal{D}}(\pi, M)$ is the probability to continue from $\pi \downarrow$ by one of the combined transitions in M . For a measurable rectangle $A \times T \times S' \in \mathfrak{F}$ and time interval T , we obtain

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

which is the probability to leave $\pi \downarrow$ via some action in A within time interval T to a state in S' .

Lemma 1. *For any $\pi \in Paths^*$, the function $\mu_{\mathcal{D}}(\pi, \cdot) : \mathfrak{F} \rightarrow [0, 1]$ is a probability measure on (Ω, \mathfrak{F}) .*

Proof. This follows from [2, Theorem 2.6.7], for $\mathcal{D}(\pi, \cdot)$ is a probability measure and all $\eta_{E(\pi \downarrow, \alpha)}$ as well as $\mathbf{P}(\pi \downarrow, \alpha, \cdot)$ are probability measures for $\alpha \in Act(\pi \downarrow)$. \square

To extend this to a probability measure on \mathfrak{F}_{Paths^n} , we assume an *initial distribution* $\nu \in Distr(\mathcal{S})$ for the probability to start in a certain state s and inductively append sets of combined transitions. To ease notation, we write $\nu(s)$ instead of $\nu(\{s\})$ where appropriate.

As the probability measures in Def. 4 (see below) depend on the Lebesgue integral of a function involving the measure $\mu_{\mathcal{D}}$, we have to show that $\mu_{\mathcal{D}} : Paths^* \times \mathfrak{F} \rightarrow [0, 1]$ is measurable in its first argument, i.e. that for all $M \in \mathfrak{F}$ and $B \in \mathfrak{B}([0, 1])$ it is the case that $\mu_{\mathcal{D}}(\cdot, M)^{-1}(B) \in \mathfrak{F}_{Paths^*}$. The following theorem stems from Wolovick and Johr in [30] and is restated here only for the sake of completeness:

Theorem 1 (Combined transition measurability [30, Theorem 1]). *Let \mathcal{C} be a CTMDP with set Act of actions and \mathcal{D} a scheduler. For all $A \in \mathfrak{F}_{Act}$, it holds: $\mathcal{D}(\cdot, A) : Paths^* \rightarrow [0, 1]$ is measurable iff $\forall M \in \mathfrak{F}, \mu_{\mathcal{D}}(\cdot, M) : Paths^* \rightarrow [0, 1]$ is measurable.*

Hence $\mu_{\mathcal{D}} : Paths^* \times \mathfrak{F} \rightarrow [0, 1]$ is measurable in its first argument whenever \mathcal{D} is a measurable scheduler as defined in Def. 3. Note also, that the restriction $\mu_{\mathcal{D}} : Paths^n \times \mathfrak{F} \rightarrow [0, 1]$ is measurable w.r.t. \mathfrak{F}_{Paths^n} .

With this precondition satisfied, we can define the probability measure on sets of finite paths as follows:

Definition 4 (Probability measure [30]). For initial distribution $\nu \in \text{Distr}(\mathcal{S})$ the probability measure on $\mathfrak{F}_{\text{Paths}^n}$ is defined inductively:

$$Pr_{\nu, \mathcal{D}}^0 : \mathfrak{F}_{\text{Paths}^0} \rightarrow [0, 1] : \Pi \mapsto \sum_{s \in \Pi} \nu(s) \quad \text{and for } n > 0$$

$$Pr_{\nu, \mathcal{D}}^n : \mathfrak{F}_{\text{Paths}^n} \rightarrow [0, 1] : \Pi \mapsto \int_{\text{Paths}^{n-1}} Pr_{\nu, \mathcal{D}}^{n-1}(d\pi) \int_{\Omega} \mathbf{I}_{\Pi}(\pi \circ m) \mu_{\mathcal{D}}(\pi, dm).$$

One further remark might be in order: For $n > 0$, the Lebesgue integral in Def. 4 is well defined as the functions

$$f_{\Pi} : \text{Paths}^{n-1} \rightarrow [0, 1] : \pi \mapsto \int_{\Omega} \mathbf{I}_{\Pi}(\pi \circ m) \mu_{\mathcal{D}}(\pi, dm)$$

are measurable for all $\Pi \in \mathfrak{F}_{\text{Paths}^n}$. First, $\{m \in \Omega \mid \pi \circ m \in \Pi\} \in \mathfrak{F}$ for all $\pi \in \text{Paths}^{n-1}$: If $\Pi = S_0 \times M_0 \times \cdots \times M_{n-1}$ is a measurable rectangle such that $M_i \in \mathfrak{F}$ for $0 \leq i < n$, we obtain

$$\{m \in \Omega \mid \pi \circ m \in \Pi\} = \begin{cases} M_{n-1} & \text{if } \pi \in S_0 \times M_0 \times \cdots \times M_{n-2} \\ \emptyset & \text{otherwise.} \end{cases}$$

Hence, for measurable rectangle Π , the set $\{m \in \Omega \mid \pi \circ m \in \Pi\}$ is measurable. Now, let $\Pi = \Pi_1 \cup \Pi_2$ and $M_i = \{m \in \Omega \mid \pi \circ m \in \Pi_i\}$ for $i = 1, 2$. By induction hypothesis, $M_i \in \mathfrak{F}$; further, $\{m \in \Omega \mid \pi \circ m \in \Pi\} = M_1 \cup M_2$. As \mathfrak{F} is closed under countable union, $M_1 \cup M_2 \in \mathfrak{F}$. For the complement Π^c , define $M = \{m \in \Omega \mid \pi \circ m \in \Pi\}$. By induction hypothesis, $M \in \mathfrak{F}$. Further observe that $\{m \in \Omega \mid \pi \circ m \in \Pi^c\} = \{m \in \Omega \mid \pi \circ m \notin \Pi\} = \{m \in \Omega \mid \pi \circ m \in \Pi\}^c = M^c$. Then $M^c \in \mathfrak{F}$ follows since $M \in \mathfrak{F}$ and \mathfrak{F} is closed under complement. Now the functions f_{Π} can be restated as follows:

$$f_{\Pi} : \text{Paths}^{n-1} \rightarrow [0, 1] : \pi \mapsto \mu_{\mathcal{D}}(\pi, \{m \in \Omega \mid \pi \circ m \in \Pi\})$$

which is measurable w.r.t. $\mathfrak{F}_{\text{Paths}^{n-1}}$ by Theorem 1, where $\mu_{\mathcal{D}}$ is restricted to Paths^{n-1} .

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.

Lemma 2. $Pr_{\nu, \mathcal{D}}^n$ is a probability measure on $(\text{Paths}^n, \mathfrak{F}_{\text{Paths}^n})$ for all $n \in \mathbb{N}$.

Proof. By induction on n . ν is a probability measure on $(\mathcal{S}, \mathfrak{F}_{\mathcal{S}})$ and so is $Pr_{\nu, \mathcal{D}}^0$. For $n > 0$,

$$Pr_{\nu, \mathcal{D}}^n(\Pi) = \int_{\text{Paths}^{n-1}} Pr_{\nu, \mathcal{D}}^{n-1}(d\pi) \int_{\Omega} \mathbf{I}_{\Pi}(\pi \circ m) \mu_{\mathcal{D}}(\pi, dm).$$

By the induction hypothesis, $Pr_{\nu, \mathcal{D}}^{n-1}$ is a probability measure; the same holds for $\mu_{\mathcal{D}}(\pi, \cdot)$ by Lemma 1. The induction step then follows by [2, 2.6.2]. \square

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 5, 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 3. For $n \geq 1$ let $I \times \Pi \in \mathfrak{F}_{\text{Paths}^n}$ be a measurable rectangle, where $I \in \mathfrak{F}_{\mathcal{S}} \times \mathfrak{F}_{\text{Act}} \times \mathfrak{B}(\mathbb{R}_{\geq 0})$. 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 :

For the induction start ($n = 1$), let $\Pi \in \mathfrak{F}_{\text{Paths}^0}$, i.e. $\Pi \subseteq \mathcal{S}$. Then:

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

For the induction step ($n > 1$), let $I \times \Pi \times M$ be a measurable rectangle in $\mathfrak{F}_{\text{Paths}^{n+1}}$ such that $I \in \mathfrak{F}_{\mathcal{S}} \times \mathfrak{F}_{\text{Act}} \times \mathfrak{B}(\mathbb{R}_{\geq 0})$ is a set of initial triples, $\Pi \in \mathfrak{F}_{\text{Paths}^{n-1}}$ 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:

$$\begin{aligned} Pr_{\nu, \mathcal{D}}^{n+1}(I \times \Pi \times M) &= \int_{I \times \Pi} \mu_{\mathcal{D}}(\pi, M) Pr_{\nu, \mathcal{D}}^n(d\pi) && (* \text{ Definition 4 } *) \\ &= \int_{I \times \Pi} \mu_{\mathcal{D}}(i \circ \pi', M) Pr_{\nu, \mathcal{D}}^n(d(i \circ \pi')) && (* \pi \simeq i \circ \pi' *) \end{aligned}$$

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

□

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 $r_n \in \mathbb{R}_{\geq 0}$ 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.

Lemma 4. *Let $k \in \mathbb{N}$ and $B = \mathcal{S} \times \Omega^k \times (\text{Act} \times [0, 1] \times \mathcal{S})^\omega$; then $Pr_{\nu, \mathcal{D}}^\omega(B) = 0$.*

Proof. The proof goes along the lines of [5, Prop. 1]:

As \mathcal{S} is finite, we can define $\Lambda := \max \{E(s, \alpha) \mid s \in \mathcal{S}, \alpha \in \text{Act}\}$. For $n \geq 0$, let $B^n := \mathcal{S} \times \Omega^k \times (\text{Act} \times [0, 1] \times \mathcal{S})^n$ be a measurable base and B_n the induced infinite measurable rectangle. By induction on n , we show that $Pr_{\nu, \mathcal{D}}^\omega(B_n) \leq (1 - e^{-\Lambda})^n$:

- Let $n = 0$. Then $Pr_{\nu, \mathcal{D}}^\omega(B_0) = Pr_{\nu, \mathcal{D}}^k(\mathcal{S} \times \Omega^k) = 1$.
- As induction hypothesis let $Pr_{\nu, \mathcal{D}}^\omega(B_n) \leq (1 - e^{-\Lambda})^n$. For B_{n+1} we obtain:

$$\begin{aligned}
Pr_{\nu, \mathcal{D}}^\omega(B_{n+1}) &= Pr_{\nu, \mathcal{D}}^{n+k+1}(B^n \times \text{Act} \times [0, 1] \times \mathcal{S}) \\
&= \int_{B^n} \mu_{\mathcal{D}}(\pi, \text{Act} \times [0, 1] \times \mathcal{S}) Pr_{\nu, \mathcal{D}}^{n+k}(d\pi) \\
&= \int_{B^n} \left(\sum_{\alpha \in \text{Act}} \mathcal{D}(\pi, \{\alpha\}) \cdot P(\pi \downarrow, \alpha, \mathcal{S}) \cdot \int_{[0,1]} E(\pi \downarrow, \alpha) e^{-E(\pi \downarrow, \alpha)t} dt \right) Pr_{\nu, \mathcal{D}}^{n+k}(d\pi) \\
&= \int_{B^n} \sum_{\alpha \in \text{Act}} \mathcal{D}(\pi, \{\alpha\}) \cdot P(\pi \downarrow, \alpha, \mathcal{S}) \cdot \left(1 - e^{-E(\pi \downarrow, \alpha)}\right) Pr_{\nu, \mathcal{D}}^{n+k}(d\pi) \\
&\leq (1 - e^{-\Lambda}) \cdot \underbrace{\int_{B^n} \sum_{\alpha \in \text{Act}} \mathcal{D}(\pi, \{\alpha\}) \cdot P(\pi \downarrow, \alpha, \mathcal{S}) Pr_{\nu, \mathcal{D}}^{n+k}(d\pi)}_{\leq 1} \\
&\leq (1 - e^{-\Lambda}) \cdot \int_{B^n} Pr_{\nu, \mathcal{D}}^{n+k}(d\pi) = (1 - e^{-\Lambda}) \cdot Pr_{\nu, \mathcal{D}}^{n+k}(B^n) \\
&= (1 - e^{-\Lambda}) \cdot Pr_{\nu, \mathcal{D}}^\omega(B_n) \leq (1 - e^{-\Lambda})^{n+1}.
\end{aligned}$$

Now $B_0 \supseteq B_1 \supseteq \dots$ and the B_n converge to B , i.e. $B_n \downarrow B$; hence $Pr_{\nu, \mathcal{D}}^\omega(B_n) \rightarrow Pr_{\nu, \mathcal{D}}^\omega(B)$ by [2, 1.2.7]. Further $\lim_{n \rightarrow \infty} Pr_{\nu, \mathcal{D}}^\omega(B_n) = 0$ for $Pr_{\nu, \mathcal{D}}^\omega$ is a measure (i.e. nonnegative) and $\lim_{n \rightarrow \infty} (1 - e^{-\Lambda})^n = 0$. Thus $Pr_{\nu, \mathcal{D}}^\omega(B) = 0$. □

With this result we can prove the following theorem which justifies to generally rule out Zeno behaviour:

Theorem 2 (Converging paths theorem). *The probability measure of the set of converging paths is zero.*

Proof. Let $ConvPaths := \{s_0 \xrightarrow{\alpha_0, t_0} s_1 \xrightarrow{\alpha_1, t_1} \dots \mid \sum_{i=0}^n t_i \text{ converges}\}$. For $\pi \in ConvPaths$, the sequence $\sum_{i=0}^{\infty} t_i$ converges; thus t_i converges to 0 and there exists $k \in \mathbb{N}$ such that $t_i \leq 1$ for all $i \geq k$. Hence $ConvPaths \subseteq \bigcup_{k=0}^{\infty} \mathcal{S} \times \Omega^k \times (Act \times [0, 1] \times \mathcal{S})^\omega$. By Lemma 4, $Pr_{\nu, \mathcal{D}}^\omega(\mathcal{S} \times \Omega^k \times (Act \times [0, 1] \times \mathcal{S})^\omega) = 0$ for all $k \in \mathbb{N}$. Thus we obtain

$$\begin{aligned} Pr_{\nu, \mathcal{D}}^\omega\left(\bigcup_{k=0}^{\infty} \mathcal{S} \times \Omega^k \times (Act \times [0, 1] \times \mathcal{S})^\omega\right) \\ \leq \sum_{k=0}^{\infty} Pr_{\nu, \mathcal{D}}^\omega(\mathcal{S} \times \Omega^k \times (Act \times [0, 1] \times \mathcal{S})^\omega) = 0. \end{aligned}$$

But then $ConvPaths$ 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 bisimilarity

Strong bisimilarity [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 6, 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.

Theorem 3 (Strong bisimilarity). Strong bisimilarity is

1. an equivalence,
2. a strong bisimulation relation and
3. the largest strong bisimulation relation.

Proof. Let $\sim = \bigcup\{\mathcal{R} \mid \mathcal{R} \text{ is a strong bisimulation relation on } \mathcal{S}\}$ denote strong bisimilarity.

1. \sim is an equivalence:

Reflexivity and symmetry follow directly from the definition.

We show transitivity: $(u, v) \in \sim$ and $(v, w) \in \sim \implies (u, w) \in \sim$.

$(u, v) \in \sim \implies$ ex. strong bisimulation relation $\mathcal{R}_1 \subseteq \sim$ s.t. $(u, v) \in \mathcal{R}_1$

$(v, w) \in \sim \implies$ ex. strong bisimulation relation $\mathcal{R}_2 \subseteq \sim$ s.t. $(v, w) \in \mathcal{R}_2$

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




Fig. 2. Example partitioning of an equivalence class $C \in \mathcal{S}_{\mathcal{R}}$.

Let \mathcal{R} denote the transitive closure of $\mathcal{R}_1 \cup \mathcal{R}_2$. Then $(u, w) \in \mathcal{R}$. Therefore it suffices to show that \mathcal{R} is a strong bisimulation relation. As \mathcal{R} obviously is an equivalence, it remains to show that for all $(u, v) \in \mathcal{R}$, $\alpha \in Act$ and $C \in \mathcal{S}_{\mathcal{R}}$ it holds $L(u) = L(v)$ and

$$\mathbf{R}(u, \alpha, C) = \mathbf{R}(v, \alpha, C). \quad (2)$$

The first condition, $L(u) = L(v)$ follows directly from the transitivity of the identity relation on 2^{AP} . For condition (2), let $C = \{s_1, \dots, s_n\}$. We have $C = \bigcup_{i=1}^n [s_i]_{\mathcal{R}_k}$ for $k \in \{1, 2\}$:

\subseteq : Let $s \in C$. Then $s \in [s_i]_{\mathcal{R}_k}$ for some $i \in \{1, \dots, n\}$. Hence $s \in \bigcup_{i=1}^n [s_i]_{\mathcal{R}_k}$.
 \supseteq : Let $i \in \{1, \dots, n\}$. Then it holds:

$$\begin{aligned} s \in [s_i]_{\mathcal{R}_k} &\iff (s, s_i) \in \mathcal{R}_k && (* \text{ by definition } *) \\ &\implies (s, s_i) \in \mathcal{R} && (* \mathcal{R}_k \subseteq \mathcal{R} *) \\ &\iff s \in [s_i]_{\mathcal{R}} && (* \mathcal{R} \text{ is an equivalence relation } *) \\ &\iff s \in C && (* [s_i]_{\mathcal{R}} = C *) \end{aligned}$$

Hence we can decompose C into equivalence classes w.r.t. \mathcal{R}_1 and \mathcal{R}_2 (see Fig. 2). As \mathcal{R}_1 is an equivalence relation, it induces a partitioning of C :

$$C = \bigsqcup \{[s_{i_1}]_{\mathcal{R}_1}, [s_{i_2}]_{\mathcal{R}_1}, \dots, [s_{i_m}]_{\mathcal{R}_1}\} \text{ where } m \leq n. \quad (3)$$

Note that the same applies to \mathcal{R}_2 for a different set of indices $i'_1, \dots, i'_{m'}$. Now we are able to prove property (2) by induction on the structure of \mathcal{R} . Therefore we provide an inductive definition of \mathcal{R} as follows:

$$\begin{aligned} \mathcal{R}^0 &= \mathcal{R}_1 \cup \mathcal{R}_2 \quad \text{and} \\ \mathcal{R}^{i+1} &= \{(u, w) \mid \exists v \in \mathcal{S}. (u, v) \in \mathcal{R}^i \wedge (v, w) \in \mathcal{R}^i\} \quad \text{for } i \geq 0. \end{aligned}$$

By construction, the subset-ordering on \mathcal{R}^i is bounded from above by $\mathcal{S} \times \mathcal{S}$. Further, \mathcal{S} is finite, so that $\mathcal{R}^0 \subseteq \mathcal{R}^1 \subseteq \dots$ is an ascending chain, that is, the transitive closure is reached after a finite number z of iterations such that $\mathcal{R}^{z+1} = \mathcal{R}^z$. Obviously, we have $\mathcal{R} = \mathcal{R}^z$.

By induction on i , we prove that if $(u, v) \in \mathcal{R}^i$, then $\mathbf{R}(u, \alpha, C) = \mathbf{R}(v, \alpha, C)$ for all $\alpha \in Act$ and $C \in \mathcal{S}_{\mathcal{R}}$:

– *induction base* ($i = 0$):

Distinguish two cases:

(a) *Case 1:* Let $(u, v) \in \mathcal{R}_1$:

$$\begin{aligned}
(u, v) \in \mathcal{R}_1 &\implies \forall C' \in \mathcal{S}_{\mathcal{R}_1}. \forall \alpha \in Act. \mathbf{R}(u, \alpha, C') = \mathbf{R}(v, \alpha, C') \\
&\implies \forall j \in \{1, \dots, m\}. \forall \alpha \in Act. \\
&\quad \mathbf{R}(u, \alpha, [s_{i_j}]_{\mathcal{R}_1}) = \mathbf{R}(v, \alpha, [s_{i_j}]_{\mathcal{R}_1}) \\
&\implies \forall \alpha \in Act. \sum_{j=1}^m \mathbf{R}(u, \alpha, [s_{i_j}]_{\mathcal{R}_1}) = \sum_{j=1}^m \mathbf{R}(v, \alpha, [s_{i_j}]_{\mathcal{R}_1}) \\
&\implies \forall \alpha \in Act. \mathbf{R}(u, \alpha, \bigoplus_{j=1}^m [s_{i_j}]_{\mathcal{R}_1}) = \mathbf{R}(v, \alpha, \bigoplus_{j=1}^m [s_{i_j}]_{\mathcal{R}_1}) \\
&\stackrel{(3)}{\implies} \forall \alpha \in Act. \mathbf{R}(u, \alpha, C) = \mathbf{R}(v, \alpha, C).
\end{aligned}$$

(b) *Case 2:* Let $(u, v) \in \mathcal{R}_2$:

The argument is completely analogue to the first case.

– *induction step* ($i \rightsquigarrow i + 1$):

Assume $(u, w) \in \mathcal{R}^{i+1}$. By construction, we have $(u, v) \in \mathcal{R}^i$ and $(v, w) \in \mathcal{R}^i$. Applying the *induction hypothesis* we have $\mathbf{R}(u, \alpha, C) = \mathbf{R}(v, \alpha, C)$ and $\mathbf{R}(v, \alpha, C) = \mathbf{R}(w, \alpha, C)$ for all actions $\alpha \in Act$ and all $C \in \mathcal{S}_{\mathcal{R}}$. Therefore $\mathbf{R}(u, \alpha, C) = \mathbf{R}(w, \alpha, C)$ directly follows from the transitivity of $=$ on $\mathbb{R}_{\geq 0}$.

Now we can conclude that \sim is indeed transitive: Given $(u, v) \in \mathcal{R}_1$ and $(v, w) \in \mathcal{R}_2$, there exists a strong bisimulation relation \mathcal{R} such that $(u, w) \in \mathcal{R}$. By definition, $\mathcal{R} \subseteq \sim$; whence $u \sim w$.

2. \sim is a strong bisimulation relation:

It remains to show for any $u \sim v$, that $\mathbf{R}(u, \alpha, C) = \mathbf{R}(v, \alpha, C)$ holds for all $\alpha \in Act$, $C \in \tilde{\mathcal{S}}$. Since $u \sim v$ implies the existence of a strong bisimulation relation $\mathcal{R} \subseteq \sim$ with $(u, v) \in \mathcal{R}$ we may follow the idea of (3) to express C as finite union of equivalence classes of $\mathcal{S}_{\mathcal{R}}$. Since \mathcal{R} is a strong bisimulation relation, the rates from u and v into those equivalence classes are equal and maintained by summation.

3. \sim is the largest (i.e. the coarsest) strong bisimulation relation:

Clear from the fact that \sim is the union of all strong bisimulation relations. \square

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.

For states $[s], [t] \in \tilde{\mathcal{S}}$ of the quotient $\tilde{\mathcal{C}}$, let $\tilde{E}([s], \alpha) := \sum_{s' \in [s]} E(s, \alpha)$ be the exit rate of $[s]$ under action α . Further, $\tilde{\mathbf{P}}([s], \alpha, [t]) := \frac{\tilde{\mathbf{R}}([s], \alpha, [t])}{\tilde{E}([s], \alpha)}$ is the discrete branching probability from state $[s]$ to state $[t]$ under action α .

Example 3. Consider the CTMDP over the set $AP = \{a\}$ of atomic propositions in Fig. 3(a). Its quotient under strong bisimilarity is outlined in Fig. 3(b).

In the quotient, exit rates and branching probabilities are preserved w.r.t. the underlying CTMDP as shown by the following two lemmas:




Fig. 3. Quotient under strong bisimilarity.

Lemma 5 (Preservation of exit rates). *Let $\mathcal{C} = (\mathcal{S}, Act, \mathbf{R}, AP, L)$ be a CTMDP and $\tilde{\mathcal{C}}$ its quotient under strong bisimilarity. Then $E(s, \alpha) = \tilde{E}([s], \alpha)$ for all $s \in \mathcal{S}$ and $\alpha \in Act$.*

Proof. Let $\mathcal{S} = \bigcup_{k=0}^n [s_{i_k}]$ such that $[s_{i_j}] \cap [s_{i_k}] = \emptyset$ for all $j \neq k$. For all states $s \in \mathcal{S}$ it holds:

$$\begin{aligned}
 E(s, \alpha) &= \sum_{s' \in \mathcal{S}} \mathbf{R}(s, \alpha, s') = \sum_{k=0}^n \sum_{s' \in [s_{i_k}]} \mathbf{R}(s, \alpha, s') = \sum_{k=0}^n \mathbf{R}(s, \alpha, [s_{i_k}]) \\
 &\stackrel{\text{Def. 7}}{=} \sum_{k=0}^n \tilde{\mathbf{R}}([s], \alpha, [s_{i_k}]) = \sum_{[s'] \in \tilde{\mathcal{S}}} \tilde{\mathbf{R}}([s], \alpha, [s']) = \tilde{E}([s], \alpha).
 \end{aligned}$$

□

With Lemma 5 it easily follows that the discrete transition probabilities are preserved under strong bisimulation:

Lemma 6 (Preservation of transition probabilities). *Let \mathcal{C} be as before and let $\tilde{\mathcal{C}}$ be its quotient under strong bisimilarity. For all states $s, t \in \mathcal{S}$ and all actions $\alpha \in Act$ it holds*

$$\tilde{\mathbf{P}}([s], \alpha, [t]) = \sum_{t' \in [t]} \mathbf{P}(s, \alpha, t').$$

Proof.

$$\begin{aligned}
 \tilde{\mathbf{P}}([s], \alpha, [t]) &= \frac{\tilde{\mathbf{R}}([s], \alpha, [t])}{\tilde{E}([s], \alpha)} \stackrel{\text{Def. 7}}{=} \frac{\mathbf{R}(s, \alpha, [t])}{\tilde{E}([s], \alpha)} \\
 &= \frac{\sum_{t' \in [t]} \mathbf{R}(s, \alpha, t')}{\tilde{E}([s], \alpha)} \stackrel{\text{Lemma 5}}{=} \frac{\sum_{t' \in [t]} \mathbf{R}(s, \alpha, t')}{E(s, \alpha)} = \sum_{t' \in [t]} \mathbf{P}(s, \alpha, t').
 \end{aligned}$$

□

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.

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

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

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 $\diamond^I \Phi \equiv \mathbb{tt} \mathbb{U}^I \Phi$ and $\square^I \Phi \equiv \neg \diamond^I \neg \Phi$ where $\mathbb{tt} := a \vee \neg a$ 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 4. Reconsider the CTMDP from Fig. 3(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 $avg_{B,t}(\pi) = \frac{1}{t} \int_0^t \mathbf{I}_B(\pi@t') dt'$. As $avg_{B,t}$ is a random variable, its expectation can be derived given an initial distribution $\nu \in Distr(\mathcal{S})$ and a measurable scheduler $\mathcal{D} \in THR$, i.e. $E(avg_{B,t}) = \int_{Paths^\omega} avg_{B,t}(\pi) Pr_{\nu, \mathcal{D}}^\omega(d\pi)$. Having the expectation for fixed time bound t , we now let $t \rightarrow \infty$ and obtain the long-run average as $\lim_{t \rightarrow \infty} E(avg_{B,t})$.

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

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

Path formulas are defined by

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

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 $\forall^{\sqsubseteq p} \varphi$ 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 4 (Measurability of path formulas). The set $\{\pi \in Paths^\omega \mid \pi \models \varphi\}$ is measurable for all CSL path formula φ .




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

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 $\pi \models \Phi U^I \Psi$. By Def. 9 it holds $\pi \models \Phi U^I \Psi$ 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 2, 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 $\{\pi \in Paths^\omega \mid \pi \models \Phi U^I \Psi\}$ can be represented by the union

$$\bigcup_{n=0}^{\infty} \left\{ \pi \in Paths^\omega \mid \sum_{i=0}^{n-1} t_i \in I \wedge \pi[n] \models \Psi \wedge \forall m < n. \pi[m] \models \Phi \right\} \quad (4)$$

$$\cup \bigcup_{n=0}^{\infty} \left\{ \pi \in Paths^\omega \mid \left(\sum_{i=0}^{n-1} t_i, \sum_{i=0}^n t_i \right) \cap I \neq \emptyset \wedge \pi[n] \models \Psi \wedge \forall m \leq n. \pi[m] \models \Phi \right\}. \quad (5)$$

It suffices to show that the subsets of (4) and (5) induced by any $n \in \mathbb{N}$ are measurable cylinders. In the following, we exhibit the proof for (5) 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:

$$\begin{aligned} & \left\{ \pi \in Paths^{n+1} \mid \left(\sum_{i=0}^{n-1} t_i, \sum_{i=0}^n t_i \right) \cap [a, b] \neq \emptyset \wedge \pi[n] \models \Psi \wedge \forall m \leq n. \pi[m] \models \Phi \right\} \\ &= \bigcup_{k=1}^{\infty} \bigcup_{\substack{c_0 + \dots + c_n \geq ak \\ d_0 + \dots + d_{n-1} \leq bk \\ c_i < d_i}} \prod_{i=0}^{n-1} \left[Sat(\Phi) \times Act \times \left(\frac{c_i}{k}, \frac{d_i}{k} \right) \right] \times Sat(\Phi \wedge \Psi) \times Act \times \left(\frac{c_n}{k}, \infty \right) \times \mathcal{S} \end{aligned} \quad (6)$$

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 (6). The intervals (c, d) and $[a, b]$ overlap, hence $c < b$ and $d > a$ (see top of Fig. 4). 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. 4. 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

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




Fig. 5. Derivation of the quotient scheduler.

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 (6) with corresponding values for c_i, d_i and k . Then $t_i \in (\frac{c_i}{k}, \frac{d_i}{k})$. Hence $a \leq \sum_{i=0}^n \frac{c_i}{k} < \sum_{i=0}^n t_i = d$ 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 (6).

The right-hand side of equation (6) is measurable, hence also the cylinder base. This extends to its cylinder and the countable union in equation (5). \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 $\Pi = S_0 \times A_0 \times T_0 \times \dots \times A_{n-1} \times T_{n-1} \times S_n$ is simple bisimulation closed if $S_i \in (\tilde{\mathcal{S}} \cup \{\emptyset\})$ for $i = 0, \dots, n$. Further, let $\tilde{\Pi} = \{S_0\} \times A_0 \times T_0 \times \dots \times A_{n-1} \times T_{n-1} \times \{S_n\}$ 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 5. Let \mathcal{C} be the CTMDP in Fig. 5(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}^ν that mimics \mathcal{D} 's behaviour on the quotient $\tilde{\mathcal{C}}$ in Fig. 5(b) can be defined by

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

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

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}^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, $\nu \in Distr(\mathcal{S})$ and $\mathcal{D} \in THR$. 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 $\mathcal{D} \in THR$, 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 5. Let \mathcal{C} be a CTMDP with set of states \mathcal{S} and $\nu \in Distr(\mathcal{S})$. Then $Pr_{\nu, \mathcal{D}}^{\omega}(\Pi) = Pr_{\tilde{\nu}, \mathcal{D}_{\sim}^{\nu}}^{\omega}(\tilde{\Pi})$ where $\mathcal{D} \in THR$ and Π simple bisimulation closed.

Proof. By induction on the length n of cylinder bases. The induction base holds for all $\nu \in Distr(\mathcal{S})$ since $Pr_{\nu, \mathcal{D}}^0([s]) = \sum_{s' \in [s]} \nu(s') = \tilde{\nu}([s]) = Pr_{\tilde{\nu}, \mathcal{D}_{\sim}^{\nu}}^0(\{[s]\})$. With the induction hypothesis that $Pr_{\nu, \mathcal{D}}^n(\Pi) = Pr_{\tilde{\nu}, \mathcal{D}_{\sim}^{\nu}}^n(\tilde{\Pi})$ for all $\nu \in Distr(\mathcal{S})$, $\mathcal{D} \in THR$ and bisimulation closed $\Pi \subseteq Paths^n$ 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) \quad (* \text{ by Lemma 5 } *) \\ &\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}_{\sim}^{\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}_{\sim}^{\nu}([s_0] \xrightarrow{\alpha, t} \cdot)}^n(\tilde{\Pi}) \cdot \sum_{s \in [s_0]} \left(\nu(s) \cdot \mathcal{D}(s, \{\alpha\}) \right) \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}_{\sim}^{\nu}([s_0] \xrightarrow{\alpha, t} \cdot)}^n(\tilde{\Pi}) \cdot \left(\sum_{s \in [s_0]} \nu(s) \right) \cdot \frac{\sum_{s \in [s_0]} \nu(s) \cdot \mathcal{D}(s, \{\alpha\})}{\sum_{s \in [s_0]} \nu(s)} \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}_{\sim}^{\nu}([s_0] \xrightarrow{\alpha, t} \cdot)}^n(\tilde{\Pi}) \cdot \tilde{\nu}([s_0]) \cdot \mathcal{D}_{\sim}^{\nu}([s_0], \{\alpha\}) \eta_{\tilde{E}([s_0], \alpha)}(dt) \\ &= \int_{\{[s_0]\}} \tilde{\nu}(d[s]) \int_{A_0} \mathcal{D}_{\sim}^{\nu}([s], d\alpha) \int_{T_0} Pr_{\tilde{\mathbf{P}}([s], \alpha, \cdot), \mathcal{D}_{\sim}^{\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}_{\sim}^{\nu}([s] \xrightarrow{\alpha, t} \cdot)}^n(\tilde{\Pi}) \tilde{\mu}_{\tilde{\nu}, \mathcal{D}_{\sim}^{\nu}}(d[s], d\alpha, dt) \end{aligned}$$

$$= Pr_{\tilde{\nu}, \mathcal{D}^{\nu}}^{n+1} (\{[s_0]\} \times A_0 \times T_0 \times \tilde{\Pi})$$

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

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

□

According to Theorem 5, 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 $\mathcal{C} = (\mathcal{S}, Act, \mathbf{R}, AP, L)$ be a CTMDP and $\tilde{\mathcal{C}}$ its quotient under strong bisimilarity. A measurable rectangle $\Pi = S_0 \times A_0 \times T_0 \times \dots \times A_{n-1} \times T_{n-1} \times S_n$ 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 $\tilde{\Pi} = \bigcup_{j=0}^{k_0} \{[s_{0,j}]\} \times A_0 \times T_0 \times \dots \times A_{n-1} \times T_{n-1} \times \bigcup_{j=0}^{k_n} \{[s_{n,j}]\}$ be the corresponding rectangle in the quotient $\tilde{\mathcal{C}}$.*

Lemma 7. *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. □

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 $Pr_{\nu, \mathcal{D}}^{\omega}(\Pi) = Pr_{\tilde{\nu}, \mathcal{D}^{\nu}}^{\omega}(\tilde{\Pi})$ for any $\mathcal{D} \in \text{THR}$ and any bisimulation closed set of paths Π .*

Proof. Follows directly from Lemma 7 and Theorem 5. □

Using these extensions we can now prove our main result:

Theorem 6. *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 4 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 $\mathcal{V}_{\sim}^{\nu_v}(\tilde{\pi}, \alpha_n)$ is independent of π . With $\tilde{v}_u = \tilde{v}_v$ and by Corollary 1 we obtain $Pr_{\tilde{v}_u, \mathcal{U}}^{\omega}(II_i) = Pr_{\tilde{v}_u, \mathcal{U}^{\nu_u}}^{\omega}(\tilde{II}_i) = Pr_{\tilde{v}_v, \mathcal{V}^{\nu_v}}^{\omega}(\tilde{II}_i) = Pr_{\tilde{v}_v, \mathcal{V}}^{\omega}(II_i)$ which carries over to Π for Π is a countable union of disjoint sets II_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 $avg_{Sat(\Psi), t}$ for $t \in \mathbb{R}_{\geq 0}$ can be expressed as follows:

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

Further, the sets $\{ \pi \in Paths^{\omega} \mid \pi @ t' \models \Psi \}$ and $\{ \pi \in Paths^{\omega} \mid \pi \models \diamond^{[t', t']} \Psi \}$ 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

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

for all $t' \in \mathbb{R}_{\geq 0}$. Thus the expectations of $avg_{Sat(\Psi), t}$ 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 [28,4] 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) 708–717
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*, Vol. 4590. Springer (2007) 316–329
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. Puterman, M. L.: *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley and Sons (1994)
26. Qiu, Q., Pedram, M.: Dynamic power management based on continuous-time Markov decision processes. In: *DAC*. ACM Press (1999) 555–561
27. Sanders, W. H., Meyer, J. F.: Stochastic activity networks: Formal definitions and concepts. In: *FMPA. LNCS*, Vol. 2090. Springer (2000) 315–343

28. Segala, R., Lynch, N.: Probabilistic simulations for probabilistic processes. *Nordic Journal of Computing* **2** (1995) 250–273
29. Vardi, M. Y.: Automatic verification of probabilistic concurrent finite-state programs. In: *FOCS*. IEEE (1985) 327–338
30. Wolovick, N., Johr, S.: A characterization of meaningful schedulers for continuous-time Markov decision processes. In: *FORMATS. LNCS*, Vol. 4202. Springer (2006) 352–367

Aachener Informatik-Berichte

This list contains all technical reports published during the past five years. A complete list of reports dating back to 1987 is available from <http://aib.informatik.rwth-aachen.de/>. To obtain copies consult the above URL or send your request to: Informatik-Bibliothek, RWTH Aachen, Ahornstr. 55, 52056 Aachen, Email: biblio@informatik.rwth-aachen.de

- 2001-01 * Jahresbericht 2000
- 2001-02 Benedikt Bollig, Martin Leucker: Deciding LTL over Mazurkiewicz Traces
- 2001-03 Thierry Cachat: The power of one-letter rational languages
- 2001-04 Benedikt Bollig, Martin Leucker, Michael Weber: Local Parallel Model Checking for the Alternation Free μ -Calculus
- 2001-05 Benedikt Bollig, Martin Leucker, Thomas Noll: Regular MSC Languages
- 2001-06 Achim Blumensath: Prefix-Recognisable Graphs and Monadic Second-Order Logic
- 2001-07 Martin Grohe, Stefan Wöhrle: An Existential Locality Theorem
- 2001-08 Mareike Schoop, James Taylor (eds.): Proceedings of the Sixth International Workshop on the Language-Action Perspective on Communication Modelling
- 2001-09 Thomas Arts, Jürgen Giesl: A collection of examples for termination of term rewriting using dependency pairs
- 2001-10 Achim Blumensath: Axiomatising Tree-interpretable Structures
- 2001-11 Klaus Indermark, Thomas Noll (eds.): Kolloquium Programmiersprachen und Grundlagen der Programmierung
- 2002-01 * Jahresbericht 2001
- 2002-02 Jürgen Giesl, Aart Middeldorp: Transformation Techniques for Context-Sensitive Rewrite Systems
- 2002-03 Benedikt Bollig, Martin Leucker, Thomas Noll: Generalised Regular MSC Languages
- 2002-04 Jürgen Giesl, Aart Middeldorp: Innermost Termination of Context-Sensitive Rewriting
- 2002-05 Horst Lichter, Thomas von der Maßen, Thomas Weiler: Modelling Requirements and Architectures for Software Product Lines
- 2002-06 Henry N. Adorna: 3-Party Message Complexity is Better than 2-Party Ones for Proving Lower Bounds on the Size of Minimal Nondeterministic Finite Automata
- 2002-07 Jörg Dahmen: Invariant Image Object Recognition using Gaussian Mixture Densities
- 2002-08 Markus Mohnen: An Open Framework for Data-Flow Analysis in Java
- 2002-09 Markus Mohnen: Interfaces with Default Implementations in Java
- 2002-10 Martin Leucker: Logics for Mazurkiewicz traces
- 2002-11 Jürgen Giesl, Hans Zantema: Liveness in Rewriting
- 2003-01 * Jahresbericht 2002
- 2003-02 Jürgen Giesl, René Thiemann: Size-Change Termination for Term Rewriting
- 2003-03 Jürgen Giesl, Deepak Kapur: Deciding Inductive Validity of Equations
- 2003-04 Jürgen Giesl, René Thiemann, Peter Schneider-Kamp, Stephan Falke: Improving Dependency Pairs
- 2003-05 Christof Löding, Philipp Rohde: Solving the Sabotage Game is PSPACE-hard
- 2003-06 Franz Josef Och: Statistical Machine Translation: From Single-Word Models to Alignment Templates
- 2003-07 Horst Lichter, Thomas von der Maßen, Alexander Nyßen, Thomas Weiler: Vergleich von Ansätzen zur Feature Modellierung bei der Softwareproduktlinienentwicklung
- 2003-08 Jürgen Giesl, René Thiemann, Peter Schneider-Kamp, Stephan Falke: Mechanizing Dependency Pairs

- 2004-01 * Fachgruppe Informatik: Jahresbericht 2003
- 2004-02 Benedikt Bollig, Martin Leucker: Message-Passing Automata are expressively equivalent to EMSO logic
- 2004-03 Delia Kesner, Femke van Raamsdonk, Joe Wells (eds.): HOR 2004 – 2nd International Workshop on Higher-Order Rewriting
- 2004-04 Slim Abdennadher, Christophe Ringeissen (eds.): RULE 04 – Fifth International Workshop on Rule-Based Programming
- 2004-05 Herbert Kuchen (ed.): WFLP 04 – 13th International Workshop on Functional and (Constraint) Logic Programming
- 2004-06 Sergio Antoy, Yoshihito Toyama (eds.): WRS 04 – 4th International Workshop on Reduction Strategies in Rewriting and Programming
- 2004-07 Michael Codish, Aart Middeldorp (eds.): WST 04 – 7th International Workshop on Termination
- 2004-08 Klaus Indermark, Thomas Noll: Algebraic Correctness Proofs for Compiling Recursive Function Definitions with Strictness Information
- 2004-09 Joachim Kneis, Daniel Mölle, Stefan Richter, Peter Rossmanith: Parameterized Power Domination Complexity
- 2004-10 Zinaida Benenson, Felix C. Gärtner, Dogan Kesdogan: Secure Multi-Party Computation with Security Modules
- 2005-01 * Fachgruppe Informatik: Jahresbericht 2004
- 2005-02 Maximilian Dornseif, Felix C. Gärtner, Thorsten Holz, Martin Mink: An Offensive Approach to Teaching Information Security: “Aachen Summer School Applied IT Security”
- 2005-03 Jürgen Giesl, René Thiemann, Peter Schneider-Kamp: Proving and Disproving Termination of Higher-Order Functions
- 2005-04 Daniel Mölle, Stefan Richter, Peter Rossmanith: A Faster Algorithm for the Steiner Tree Problem
- 2005-05 Fabien Pouget, Thorsten Holz: A Pointillist Approach for Comparing Honey Pots
- 2005-06 Simon Fischer, Berthold Vöcking: Adaptive Routing with Stale Information
- 2005-07 Felix C. Freiling, Thorsten Holz, Georg Wicherski: Botnet Tracking: Exploring a Root-Cause Methodology to Prevent Distributed Denial-of-Service Attacks
- 2005-08 Joachim Kneis, Peter Rossmanith: A New Satisfiability Algorithm With Applications To Max-Cut
- 2005-09 Klaus Kursawe, Felix C. Freiling: Byzantine Fault Tolerance on General Hybrid Adversary Structures
- 2005-10 Benedikt Bollig: Automata and Logics for Message Sequence Charts
- 2005-11 Simon Fischer, Berthold Vöcking: A Counterexample to the Fully Mixed Nash Equilibrium Conjecture
- 2005-12 Neeraj Mittal, Felix Freiling, S. Venkatesan, Lucia Draque Penso: Efficient Reductions for Wait-Free Termination Detection in Faulty Distributed Systems
- 2005-13 Carole Delporte-Gallet, Hugues Fauconnier, Felix C. Freiling: Revisiting Failure Detection and Consensus in Omission Failure Environments
- 2005-14 Felix C. Freiling, Sukumar Ghosh: Code Stabilization
- 2005-15 Uwe Naumann: The Complexity of Derivative Computation
- 2005-16 Uwe Naumann: Syntax-Directed Derivative Code (Part I: Tangent-Linear Code)
- 2005-17 Uwe Naumann: Syntax-directed Derivative Code (Part II: Intraprocedural Adjoint Code)
- 2005-18 Thomas von der Maßen, Klaus Müller, John MacGregor, Eva Geisberger, Jörg Dörr, Frank Houdek, Harbhajan Singh, Holger Wußmann, Hans-Veit Bacher, Barbara Paech: Einsatz von Features im Software-Entwicklungsprozess - Abschlußbericht des GI-Arbeitskreises “Features”
- 2005-19 Uwe Naumann, Andre Vehreschild: Tangent-Linear Code by Augmented LL-Parsers

- 2005-20 Felix C. Freiling, Martin Mink: Bericht über den Workshop zur Ausbildung im Bereich IT-Sicherheit Hochschulausbildung, berufliche Weiterbildung, Zertifizierung von Ausbildungsangeboten am 11. und 12. August 2005 in Köln organisiert von RWTH Aachen in Kooperation mit BITKOM, BSI, DLR und Gesellschaft fuer Informatik (GI) e.V.
- 2005-21 Thomas Noll, Stefan Rieger: Optimization of Straight-Line Code Revisited
- 2005-22 Felix Freiling, Maurice Herlihy, Lucia Draque Penso: Optimal Randomized Fair Exchange with Secret Shared Coins
- 2005-23 Heiner Ackermann, Alantha Newman, Heiko Röglin, Berthold Vöcking: Decision Making Based on Approximate and Smoothed Pareto Curves
- 2005-24 Alexander Becher, Zinaida Benenson, Maximillian Dornseif: Tampering with Motes: Real-World Physical Attacks on Wireless Sensor Networks
- 2006-01 * Fachgruppe Informatik: Jahresbericht 2005
- 2006-02 Michael Weber: Parallel Algorithms for Verification of Large Systems
- 2006-03 Michael Maier, Uwe Naumann: Intraprocedural Adjoint Code Generated by the Differentiation-Enabled NAGWare Fortran Compiler
- 2006-04 Ebadollah Varnik, Uwe Naumann, Andrew Lyons: Toward Low Static Memory Jacobian Accumulation
- 2006-05 Uwe Naumann, Jean Utke, Patrick Heimbach, Chris Hill, Derya Ozyurt, Carl Wunsch, Mike Fagan, Nathan Tallent, Michelle Strout: Adjoint Code by Source Transformation with OpenAD/F
- 2006-06 Joachim Kneis, Daniel Mölle, Stefan Richter, Peter Rossmann: Divide-and-Color
- 2006-07 Thomas Colcombet, Christof Löding: Transforming structures by set interpretations
- 2006-08 Uwe Naumann, Yuxiao Hu: Optimal Vertex Elimination in Single-Expression-Use Graphs
- 2006-09 Tingting Han, Joost-Pieter Katoen: Counterexamples in Probabilistic Model Checking
- 2006-10 Mesut Günes, Alexander Zimmermann, Martin Wenig, Jan Ritterfeld, Ulrich Meis: From Simulations to Testbeds - Architecture of the Hybrid MCG-Mesh Testbed
- 2006-11 Bastian Schlich, Michael Rohrbach, Michael Weber, Stefan Kowalewski: Model Checking Software for Microcontrollers
- 2006-12 Benedikt Bollig, Joost-Pieter Katoen, Carsten Kern, Martin Leucker: Replaying Play in and Play out: Synthesis of Design Models from Scenarios by Learning
- 2006-13 Wong Karianto, Christof Löding: Unranked Tree Automata with Sibling Equalities and Disequalities
- 2006-14 Danilo Beuche, Andreas Birk, Heinrich Dreier, Andreas Fleischmann, Heidi Galle, Gerald Heller, Dirk Janzen, Isabel John, Ramin Tavakoli Kolagari, Thomas von der Maßen, Andreas Wolfram: Report of the GI Work Group "Requirements Management Tools for Product Line Engineering"
- 2006-15 Sebastian Ullrich, Jakob T. Valvoda, Torsten Kuhlen: Utilizing optical sensors from mice for new input devices
- 2006-16 Rafael Ballagas, Jan Borchers: Selexels: a Conceptual Framework for Pointing Devices with Low Expressiveness
- 2006-17 Eric Lee, Henning Kiel, Jan Borchers: Scrolling Through Time: Improving Interfaces for Searching and Navigating Continuous Audio Timelines
- 2007-01 * Fachgruppe Informatik: Jahresbericht 2006
- 2007-02 Carsten Fuhs, Jürgen Giesl, Aart Middeldorp, Peter Schneider-Kamp, René Thiemann, and Harald Zankl: SAT Solving for Termination Analysis with Polynomial Interpretations
- 2007-03 Jürgen Giesl, René Thiemann, Stephan Swiderski, and Peter Schneider-Kamp: Proving Termination by Bounded Increase
- 2007-04 Jan Buchholz, Eric Lee, Jonathan Klein, and Jan Borchers: coJIVE: A System to Support Collaborative Jazz Improvisation

- 2007-05 Uwe Naumann: On Optimal DAG Reversal
- 2007-06 Joost-Pieter Katoen, Thomas Noll, and Stefan Rieger: Verifying Concurrent List-Manipulating Programs by LTL Model Checking
- 2007-07 Alexander Nyßen, Horst Lichter: MeDUSA - Method for UML2-based Design of Embedded Software Applications
- 2007-08 Falk Salewski and Stefan Kowalewski: Achieving Highly Reliable Embedded Software: An empirical evaluation of different approaches
- 2007-09 Tina Krauß, Heiko Mantel, and Henning Sudbrock: A Probabilistic Justification of the Combining Calculus under the Uniform Scheduler Assumption
- 2007-12 Uwe Naumann: An L-Attributed Grammar for Adjoint Code
- 2007-13 Uwe Naumann, Michael Maier, Jan Riehme, and Bruce Christianson: Second-Order Adjoints by Source Code Manipulation of Numerical Programs
- 2007-14 Jean Utke, Uwe Naumann, Mike Fagan, Nathan Tallent, Michelle Strout, Patrick Heimbach, Chris Hill, and Carl Wunsch: OpenAD/F: A Modular, Open-Source Tool for Automatic Differentiation of Fortran Codes

* These reports are only available as a printed version.

Please contact biblio@informatik.rwth-aachen.de to obtain copies.