Comparative Branching-Time Semantics for Markov Chains (extended abstract)

Christel Baier[a], Holger Hermanns[b,c], Joost-Pieter Katoen[b], and Verena Wolf[a]

[a] *Institut für Informatik I, University of Bonn Römerstraße 164, D-53117 Bonn, Germany*

[b] *Department of Computer Science, University of Twente P.O. Box 217, 7500 AE Enschede, The Netherlands*

[c] *Department of Computer Science Saarland University, D-66123 Saarbrücken, Germany*

Abstract. This paper presents various semantics in the branching-time spectrum of discrete-time and continuous-time Markov chains (DTMCs and CTMCs). Strong and weak bisimulation equivalence and simulation pre-orders are covered and are logically characterised in terms of the temporal logics PCTL and CSL. Apart from presenting various existing branching-time relations in a uniform manner, our contributions are: (i) weak simulation for DTMCs is defined, (ii) weak bisimulation equivalence is shown to coincide with weak simulation equivalence, (iii) logical characterisation of weak (bi)simulations are provided, and (iv) a classification of branching-time relations is presented, elucidating the semantics of DTMCs, CTMCs and their interrelation.

1 Introduction

Equivalences and pre-orders are important means to compare the behaviour of transition systems. Prominent branching-time relations are bisimulation and simulation. Bisimulations [36] are equivalences requiring related states to exhibit identical stepwise behaviour. Simulations [30] are preorders requiring state s' to mimic s in a stepwise manner, but not necessarily the reverse, i.e., s' may perform steps that cannot be matched by s. Typically, strong and weak relations are distinguished. Whereas in *strong* (bi)simulations, each individual step needs to be mimicked, in *weak* (bi)simulations this is only required for observable steps but not for internal computations. Weak relations thus allow for stuttering.

A plethora of strong and weak (bi)simulations for labelled transition systems has been defined in the literature, and their relationship has been studied by process algebraists, most notably by van Glabbeek [22, 23]. These “comparative” semantics have been extended with logical characterisations. Strong bisimulation, for instance, coincides with CTL-equivalence [13], whereas strong simulation agrees with a “preorder” on the universal (or existential) fragment of CTL [15]. Similar results hold for weak (bi)simulation where typically the next operator is omitted, which is not compatible with stuttering.

For probabilistic systems, a similar situation exists. Based on the seminal works of [31, 35], notions of (bi)simulation (see, e.g., [2, 7, 8, 11, 12, 24, 27, 28, 32, 38, 40, 41]) for models with and without nondeterminism have been defined during the last decade, and various logics to reason about such systems have been proposed (see e.g., [1, 4, 10, 26]). This holds for both discrete probabilistic systems and variants thereof, as well as systems that describe continuous-time stochastic phenomena. In particular, in the discrete setting several slight variants of (bi)simulations have been defined, and their logical characterisations studied, e.g., [3, 17, 21, 19, 40]. Although the relationship between (bi)simulations is fragmentarily known, a clear, concise classification is – in our opinion – lacking. Moreover, continuous-time and discrete-time semantics have largely been developed in isolation, and their connection has received scant attention, if at all.

This paper attempts to study the comparative semantics of branching-time relations for probabilistic systems that do not exhibit any nondeterminism. In particular, time-abstract (or discrete-time) fully probabilistic systems (FPS) and continuous-time Markov chains (CTMCs) are considered. Strong and weak (bi)simulation relations are covered together with their characterisation in terms of the temporal logics PCTL [26] and CSL [4, 10] for the discrete and continuous setting, respectively. Apart from presenting various existing branching-time relations and their connection in a uniform manner, several new results are provided. For FPSs, weak bisimulation [7] is shown to coincide with $\text{PCTL}_{\setminus X}$-equivalence, weak simulation is introduced whose kernel agrees with weak bisimulation, and the preorder weakly preserves a safe (live) fragment of $\text{PCTL}_{\setminus X}$. In the continuous-time setting, strong simulation is defined and is shown to coincide with a preorder on CSL. These results are pieced together with various results known from the literature, forming a uniform characterisation of the semantic spectrum of FPSs, CTMCs and of their interrelation.

Organisation of the paper. Section 2 provides the necessary background. Section 3 defines strong and weak (bi)simulations. Section 4 introduces PCTL and CSL and presents the logical characterisations. Section 5 presents the branching-time spectrum. Section 6 concludes the paper. Some proofs are included in this paper; for remaining proofs, see [9].

2 Preliminaries

This section introduces the basic concepts of the Markov models considered within this paper; for a more elaborate treatment see e.g., [25, 33, 34]. Let AP be a fixed, finite set of atomic propositions.

Definition 1. *A* fully probabilistic system *(FPS) is a tuple* $\mathcal{D} = (S, \mathbf{P}, L)$ *where:*

- *S is a countable set of states*
- $\mathbf{P} : S \times S \to [0,1]$ *is a probability matrix satisfying* $\sum_{s' \in S} \mathbf{P}(s,s') \in [0,1]$ *for all* $s \in S$
- $L : S \to 2^{AP}$ *is a labelling function which assigns to each state* $s \in S$ *the set* $L(s)$ *of atomic propositions that are valid in* s. ■

If $\sum_{s' \in S} \mathbf{P}(s,s') = 1$, state s is called stochastic, if this sum equals zero, state s is called absorbing; otherwise, s is called sub-stochastic.

Definition 2. *A (labelled)* DTMC *is an FPS where any state is either stochastic or absorbing, i.e.,* $\sum_{s' \in S} \mathbf{P}(s,s') \in \{0,1\}$ *for all* $s \in S$. ■

For $C \subseteq S$, $\mathbf{P}(s,C) = \sum_{s' \in C} \mathbf{P}(s,s')$ denotes the probability for s to move to a C-state. For technical reasons, $\mathbf{P}(s,\perp) = 1 - \mathbf{P}(s,S)$. Intuitively, $\mathbf{P}(s,\perp)$ denotes the probability to stay forever in s without performing any transition; although $\perp$ is not a "real" state (i.e., $\perp \notin S$), it may be regarded as a deadlock. In the context of simulation relations later on, $\perp$ is treated as an auxiliary state that is simulated by any other state. Let $S_\perp = S \cup \{\perp\}$. $\mathsf{Post}(s) = \{ s' \mid \mathbf{P}(s,s') > 0 \}$ denotes the set of direct successor states of s, and $\mathsf{Post}_\perp(s) = \{ s' \in S_\perp \mid \mathbf{P}(s,s') > 0 \}$, i.e., $\mathsf{Post}(s) \cup \{ \perp \mid \mathbf{P}(s,\perp) > 0 \}$.

We consider FPSs and therefore also DTMCs as *time-abstract* models. The name DTMC has historical reasons. A (discrete-)timed interpretation is appropriate in settings where all state changes occur at equidistant time points. For weak relations the time-abstract view will be decisive. In contrast, CTMCs are considered as *time-aware*, as they have an explicit reference to (real-)time, in the form of transition rates which determine the stochastic evolution of the system in time.

Definition 3. *A (labelled) CTMC is a tuple $\mathcal{C} = (S, \mathbf{R}, L)$ with S and L as before, and* rate matrix $\mathbf{R} : S \times S \to I\!R_{\geqslant 0}$ *such that the exit rate $E(s) = \sum_{s' \in S} \mathbf{R}(s,s')$ is finite.* ■

As in the discrete case, $\mathsf{Post}(s) = \{ s' \mid \mathbf{R}(s,s') > 0 \}$ denotes the set of direct successor states of s, and for $C \subseteq S$, $\mathbf{R}(s,C) = \sum_{s' \in C} \mathbf{R}(s,s')$ denotes the rate of moving from state s to C via a single transition.

The meaning of $\mathbf{R}(s,s') = \lambda > 0$ is that with probability $1 - e^{-\lambda \cdot t}$ the transition $s \to s'$ is enabled within the next t time units (provided that the current state is s). If $\mathbf{R}(s,s') > 0$ for more than one state s', a *race* between the outgoing transitions from s exists. The probability of s' winning this race before time t is $\frac{\mathbf{R}(s,s')}{E(s)} \cdot (1 - e^{-E(s) \cdot t})$. With $t \to \infty$ we get the time-abstract behaviour by the so-called embedded DTMC:

Definition 4. *The* embedded *DTMC of CTMC $\mathcal{C} = (S, \mathbf{R}, L)$ is given by $\mathsf{emb}(\mathcal{C}) = (S, \mathbf{P}, L)$, where $\mathbf{P}(s,s') = \mathbf{R}(s,s')/E(s)$ if $E(s) > 0$ and $\mathbf{P}(s,s') = 0$ otherwise.* ■

A CTMC is called *uniformised* if all states in $\mathcal{C}$ have the same exit rate. Each CTMC can be transformed into a uniformised CTMCs by adding self-loops [39]:

Definition 5. *Let $\mathcal{C} = (S, \mathbf{R}, L)$ be a CTMC and let (uniformisation rate) E be a real such that $E \geqslant \max_{s \in S} E(s)$. Then, $\mathit{unif}(\mathcal{C}) = (S, \overline{\mathbf{R}}, L)$ is a uniformised CTMC with $\overline{\mathbf{R}}(s,s') = \mathbf{R}(s,s')$ for $s \neq s'$, and $\overline{\mathbf{R}}(s,s) = \mathbf{R}(s,s) + E - E(s)$.* ■

In $\mathit{unif}(\mathcal{C})$ all rates of self-loops are "normalised" with respect to E, such that state transitions occur with an average "pace" of E, uniform for all states of the chain. We will later see that $\mathcal{C}$ and $\mathit{unif}(\mathcal{C})$ are related by weak bisimulation.

Paths and the probability measures on paths in FPSs and CTMCs are defined by a standard construction, e.g., [25, 33, 34], and are omitted here.

3 Bisimulation and simulation

We will use the subscript "d" to identify relations defined in the discrete setting (FPSs or DTMCs), and "c" for the continuous setting (CTMCs).

Definition 6. **[33, 35, 32, 24]** *Let $\mathcal{D} = (S, \mathbf{P}, L)$ be a FPS and R an equivalence relation on S. R is a* strong bisimulation *on $\mathcal{D}$ if for $s_1 \, R \, s_2$: $L(s_1) = L(s_2)$ and $\mathbf{P}(s_1, C) = \mathbf{P}(s_2, C)$ for all C in S/R. s_1 and s_2 in $\mathcal{D}$ are strongly bisimilar, denoted $s_1 \sim_d s_2$, if there exists a strong bisimulation R on $\mathcal{D}$ with $s_1 \, R \, s_2$.* ■

Definition 7. **[14,28]** *Let $\mathcal{C} = (S, \mathbf{R}, L)$ be a CTMC and R an equivalence relation on S. R is a* strong bisimulation *on $\mathcal{C}$ if for $s_1 \, R \, s_2$: $L(s_1) = L(s_2)$ and $\mathbf{R}(s_1, C) = \mathbf{R}(s_2, C)$ for all C in S/R. s_1 and s_2 in $\mathcal{C}$ are strongly bisimilar, denoted $s_1 \sim_c s_2$, if there exists a strong bisimulation R on $\mathcal{C}$ with $s_1 \, R \, s_2$.* ■

As $\mathbf{R}(s,C) = \mathbf{P}(s,C) \cdot E(s)$, the condition on the cumulative rates can be reformulated as (i) $\mathbf{P}(s_1,C) = \mathbf{P}(s_2,C)$ for all $C \in S/R$ and (ii) $E(s_1) = E(s_2)$. Hence, $\sim_c$ agrees with $\sim_d$ in the embedded DTMC provided that exit rates are treated as additional atomic propositions. By the standard construction, it can be shown that $\sim_d$ and $\sim_c$ are the coarsest strong bisimulations.

Proposition 1. *For CTMC $\mathcal{C} = (S, \mathbf{R}, L)$:*

1. *$s_1 \sim_c s_2$ implies $s_1 \sim_d s_2$ in $\mathrm{emb}(\mathcal{C})$, for any state $s_1, s_2 \in S$.*
2. *if $\mathcal{C}$ is uniformised then $\sim_c$ coincides with $\sim_d$ in $\mathrm{emb}(\mathcal{C})$.*

Definition 8. *A* distribution *on set S is a function $\mu : S \to [0,1]$ with $\sum_{s \in S} \mu(s) \leqslant 1$.* ■

We put $\mu(\perp) = 1 - \sum_{s \in S} \mu(s)$. $Distr(S)$ denotes the set of all distributions on S. Distribution μ on S is called stochastic if $\mu(\perp) = 0$. For simulation relations, the concept of weight functions is important.

Definition 9. **[29,31]** *Let S be a set, $R \subseteq S \times S$, and $\mu, \mu' \in Distr(S)$. A* weight function *for μ and μ' with respect to R is a function $\Delta : S_\perp \times S_\perp \to [0,1]$ such that:*

1. *$\Delta(s,s') > 0$ implies $s \, R \, s'$ or $s = \perp$*
2. *$\mu(s) = \sum_{s' \in S_\perp} \Delta(s,s')$ for any $s \in S_\perp$*
3. *$\mu'(s') = \sum_{s \in S_\perp} \Delta(s,s')$ for any $s' \in S_\perp$*

We write $\mu \sqsubseteq_R \mu'$ (or simply $\sqsubseteq$, if R is clear from the context) iff there exists a weight function for μ and μ' with respect to R. $\sqsubseteq_R$ is the lift of R to distributions. ■

Definition 10. **[31]** *Let $\mathcal{D} = (S, \mathbf{P}, L)$ be a FPS and $R \subseteq S \times S$. R is a* strong simulation *on $\mathcal{D}$ if for all $s_1 \, R \, s_2$: $L(s_1) = L(s_2)$ and $\mathbf{P}(s_1, \cdot) \sqsubseteq_R \mathbf{P}(s_2, \cdot)$. s_2 strongly simulates s_1 in $\mathcal{D}$, denoted $s_1 \precsim_d s_2$, iff there exists a strong simulation R on $\mathcal{D}$ such that $s_1 \, R \, s_2$.* ■

It is not difficult to see that $s_1 \sim_d s_2$ implies $s_1 \precsim_d s_2$. For a DTMC without absorbing states, $\precsim_d$ is symmetric and coincides with $\sim_d$, see [31].

Proposition 2. **[5,16]** *For any FPS, $\precsim_d \cap \precsim_d^{-1}$ coincides with $\sim_d$.*

Definition 11. *Let $\mathcal{C} = (S, \mathbf{R}, L)$ be a CTMC and $R \subseteq S \times S$. R is a* strong simulation *on $\mathcal{C}$ if for all $s_1 \, R \, s_2$: $L(s_1) = L(s_2)$, $\mathbf{P}(s_1, \cdot) \sqsubseteq_R \mathbf{P}(s_2, \cdot)$ and $E(s_1) \leqslant E(s_2)$. s_2 strongly simulates s_1 in $\mathcal{C}$, denoted $s_1 \precsim_c s_2$, iff there exists a strong simulation R on $\mathcal{C}$ such that $s_1 \, R \, s_2$.* ■

Proposition 3. *For any CTMC $\mathcal{C}$:*

1. *$s_1 \sim_c s_2$ implies $s_1 \precsim_c s_2$, for any state $s_1, s_2 \in S$.*
2. *$s_1 \precsim_c s_2$ implies $s_1 \precsim_d s_2$ in $\mathrm{emb}(\mathcal{C})$, for any state $s_1, s_2 \in S$.*
3. *$\precsim_c \cap \precsim_c^{-1}$ coincides with $\sim_c$.*
4. *if $\mathcal{C}$ is uniformised then $\precsim_c$ is symmetric and coincides with $\sim_c$.*

Weak bisimulation. In this paper, we only consider weak bisimulation which relies on branching bisimulation in the style of van Glabbeek and Weijland and only abstracts from stutter-steps inside the equivalence classes. While for ordinary transition systems branching bisimulation is strictly finer than Milner's observational equivalence, they agree for FPSs [7], and thus for CTMCs.

Let $\mathcal{D} = (S, \mathbf{P}, L)$ be a DTMC and $R \subseteq S \times S$ an equivalence relation. Any transition $s \to s'$ where s and s' are R-equivalent is an R-silent move. Let Silent_R denote the set of states $s \in S$ for which $\mathbf{P}(s, [s]_R) = 1$, i.e., all stochastic states that do not have a successor state outside their R-equivalence class. For any state $s \notin \mathsf{Silent}_R$, $s' \in S$ with $s' \notin [s]_R$:

$$\mathbf{P}(s, s' \mid \text{no } R\text{-silent move}) \;=\; \frac{\mathbf{P}(s,s')}{1 - \mathbf{P}(s,[s]_R)}$$

denotes the conditional probability to move from s to s' via a single transition under the condition that from s *no* transition inside $[s]_R$ is taken. Thus, either a transition is taken to another equivalence class under R or, for sub-stochastic states, the system deadlocks. For $C \subseteq S$ with $C \cap [s]_R = \varnothing$ let $\mathbf{P}(s, C \mid \text{no } R\text{-silent move}) = \sum_{s' \in C} \mathbf{P}(s, s' \mid \text{no } R\text{-silent move})$.

Definition 12. **[7]** *Let $\mathcal{D} = (S, \mathbf{P}, L)$ be a FPS and R an equivalence relation on S. R is a* weak bisimulation *on $\mathcal{D}$ if for all $s_1 \, R \, s_2$:*

1. $L(s_1) = L(s_2)$
2. *If s_1, $s_2 \notin \mathsf{Silent}_R$ then:* $\mathbf{P}(s_1, C \mid \text{no } R\text{-silent move}) = \mathbf{P}(s_2, C \mid \text{no } R\text{-silent move})$ *for all $C \in S/R$, $C \neq [s_1]_R$.*
3. *If $s_1 \in \mathsf{Silent}_R$ and $s_2 \notin \mathsf{Silent}_R$ then s_1 can reach a state $s' \in [s_1]_R \setminus \mathsf{Silent}_R$ with positive probability.*

s_1 and s_2 in $\mathcal{D}$ are weakly bisimilar, denoted $s_1 \approx_d s_2$, iff there exists a weak bisimulation R on $\mathcal{D}$ such that $s_1 \, R \, s_2$. ■

By the third condition, for any R-equivalence class C, either all states in C are R-silent (i.e., $\mathbf{P}(s, C) = 1$ for $s \in C$) or for $s \in C$ there is a path fragment that ends in an equivalence class that differs from C.

Example 1. For the following DTMC (where equally shaded states are equally labeled) the reachability condition is needed to establish a weak bisimulation for states s_1 and s_2:

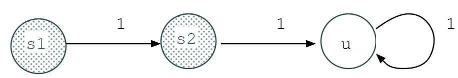

We have $s_1 \approx_d s_2$, and s_1 is $\approx_d$-silent while s_2 is not. Here, the reachability condition is obviously fulfilled. This condition can, however, not be dropped: otherwise s_1 and s_2 would be weakly bisimilar to an absorbing state with the same labeling. ■

Definition 13. **[12]** *Let $\mathcal{C} = (S, \mathbf{R}, L)$ be a CTMC and R an equivalence relation on S. R is a* weak bisimulation *on $\mathcal{C}$ if for all $s_1 \, R \, s_2$: $L(s_1) = L(s_2)$ and $\mathbf{R}(s_1, C) = \mathbf{R}(s_2, C)$ for all $C \in S/R$ with $C \neq [s_1]_R$. s_1 and s_2 in $\mathcal{C}$ are weakly bisimilar, denoted $s_1 \approx_c s_2$, iff there exists a weak bisimulation R on $\mathcal{C}$ such that $s_1 \, R \, s_2$.* ■

Proposition 4. *For any CTMC $\mathcal{C}$:*

1. $\sim_c$ *is strictly finer than* $\approx_c$.
2. *if* $\mathcal{C}$ *is uniformised then* $\approx_c$ *coincides with* $\sim_c$.
3. $\approx_c$ *coincides with* $\approx_c$ *in* $\mathsf{unif}(\mathcal{C})$.

The last result can be strengthened as follows. Any state s in $\mathcal{C}$ is weakly bisimilar to s considered as a state in $unif(\mathcal{C})$. (For this, consider the disjoint union of $\mathcal{C}$ and $unif(\mathcal{C})$ as a single CTMC.)

Proposition 5. *For CTMC* $\mathcal{C}$ *with* $s_1, s_2 \in S$*:* $s_1 \approx_c s_2$ *implies* $s_1 \approx_d s_2$ *in* $\mathsf{emb}(\mathcal{C})$.

Proof. Let R be a weak bisimulation on $\mathcal{C}$. We show that R is a weak bisimulation on $\mathsf{emb}(\mathcal{C})$ as follows. First, observe that all R-equivalent states have the same labelling. Assume $s_1 \, R \, s_2$ and $B = [s_1]_R = [s_2]_R$. Distinguish two cases. (i) s_1 is R-silent, i.e., $\mathbf{P}(s_1, B) = 1$. Hence, $\mathbf{R}(s_1, B) = E(s_1)$ and therefore $0 = \mathbf{R}(s_1, C) = \mathbf{R}(s_2, C)$ for all $C \in S/R$ with $C \neq B$. So, $\mathbf{P}(s_2, B) = 1$. (ii) Neither s_1 nor s_2 is R-silent, i.e., $\mathbf{P}(s_i, B) < 1$, for i=1,2. Note that:

$$E(s_i) = \sum_{\substack{C \in S/R \\ C \neq B}} \mathbf{R}(s_i, C) \; + \; \mathbf{R}(s_i, B)$$

As $s_1 \approx_c s_2$, $\mathbf{R}(s_1, C) = \mathbf{R}(s_2, C)$ for all $C \in S/R$ with $C \neq B$. Hence, $\sum_{C \in S/R, C \neq B} \mathbf{R}(s_1, C) = \sum_{C \in S/R, C \neq B} \mathbf{R}(s_2, C)$ and therefore $E(s_1) - \mathbf{R}(s_1, B) = E(s_2) - \mathbf{R}(s_2, B)$ (*). For any $C \in S/R$ with $C \neq B$ we derive:

$$\begin{aligned} \mathbf{P}(s_1, C \mid \text{no } R\text{-silent move}) \quad &\stackrel{def}{=} \quad \frac{\mathbf{P}(s_1, C)}{1 - \mathbf{P}(s_1, B)} \quad = \quad \frac{E(s_1) \cdot \mathbf{P}(s_1, C)}{E(s_1) - E(s_1) \cdot \mathbf{P}(s_1, B)} \\ \stackrel{def.\mathbf{R}}{=} \quad \frac{\mathbf{R}(s_1, C)}{E(s_1) - \mathbf{R}(s_1, B)} \quad &\stackrel{(*), s_1 \approx_c s_2}{=} \quad \frac{\mathbf{R}(s_2, C)}{E(s_2) - \mathbf{R}(s_2, B)} \quad = \quad \frac{\mathbf{P}(s_2, C)}{1 - \mathbf{P}(s_2, B)} \end{aligned}$$

which, by definition, equals $\mathbf{P}(s_2, C \mid \text{no } R\text{-silent move})$. So, $s_1 \approx_d s_2$. ■

Remark 1. Prop. 1.2 states that for a uniformised CTMC, $\sim_c$ coincides with $\sim_d$ on the embedded DTMC. The analogue for $\approx_c$ does not hold, as, e.g., in the uniformised CTMC of Example 1 we have $s_1 \approx_d s_2$ but $s_1 \not\approx_c s_2$ as $\mathbf{R}(s_1, [u]) \neq \mathbf{R}(s_2, [u])$. Intuitively, although s_1 and s_2 have the same time-abstract behaviour (up to stuttering) they have distinct timing behaviour. s_1 is "slower than" s_2 as it has to perform a stutter step prior to an observable step (from s_2 to u) while s_2 can immediately perform the latter step. Note that by Prop 4.2 and Prop. 1.2, $\approx_c$ coincides with $\sim_d$ for uniformised CTMCs. In fact, Prop. 5 can be strengthened in the following way: $\approx_c$ is the coarsest equivalence finer than $\approx_d$ such that $s_1 \approx_c s_2$ implies $\mathbf{R}(s_1, S \setminus [s_1]) = \mathbf{R}(s_2, S \setminus [s_2])$. ■

Weak simulation. Weak simulation on FPSs is inspired by our work on CTMCs [8]. Roughly speaking, $s_1 \precapprox s_2$ if the successor states of s_1 and s_2 can be grouped into subsets U_i and V_i (assume, for simplicity, $U_i \cap V_i = \varnothing$). All transitions from s_i to V_i are viewed as stutter-steps, i.e., internal transitions that do not change the labelling and respect $\precapprox$. To that end, any state in V_1 is required to be simulated by s_2 and, symmetrically, any state in V_2 simulates s_1. Transitions from s_i to U_i are regarded as visible steps. Accordingly, we require that the distributions for the conditional probabilities $u_1 \mapsto \mathbf{P}(s_1, u_1)/K_1$ and $u_2 \mapsto \mathbf{P}(s_2, u_2)/K_2$ to move from s_i to U_i are related via a weight function (as for $\precsim_d$). K_i denotes the total probability to move from s_i to a state in U_i in a single step. For technical reasons, we allow $\perp \in U_i$ and $\perp \in V_i$.

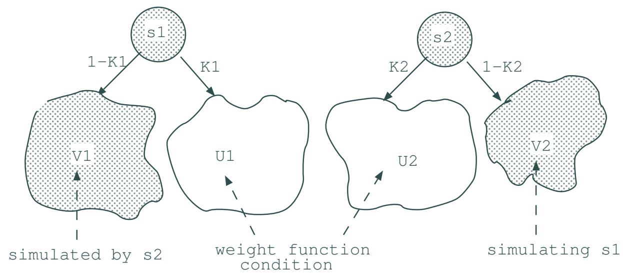

Definition 14. *Let $\mathcal{D} = (S, \mathbf{P}, L)$ be a FPS and $R \subseteq S \times S$. R is a* weak simulation *on $\mathcal{D}$ iff for $s_1 \, R \, s_2$: $L(s_1) = L(s_2)$ and there exist functions $\delta_i : S_\perp \to [0,1]$ and sets $U_i, V_i \subseteq S_\perp$ (i=$1,2$) with*

$$U_i \; = \; \{ \, u_i \in \mathsf{Post}_\perp(s_i) \mid \delta_i(u_i) > 0 \, \} \text{ and } V_i \; = \; \{ \, v_i \in \mathsf{Post}_\perp(s_i) \mid \delta_i(v_i) < 1 \, \}$$

such that:

1. *(a) $v_1 \, R \, s_2$ for all $v_1 \in V_1$, $v_1 \neq \perp$, and (b) $s_1 \, R \, v_2$ for all $v_2 \in V_2$, $v_2 \neq \perp$*
2. *there exists a function $\Delta : S_\perp \times S_\perp \to [0,1]$ such that:*
   (a) $\Delta(u_1, u_2) > 0$ implies $u_1 \in U_1$, $u_2 \in U_2$ and either $u_1 \, R \, u_2$ or $u_1 = \perp$,
   (b) if $K_1 > 0$ and $K_2 > 0$ then for all states $w \in S$:

$$K_1 \cdot \sum_{u_2 \in U_2} \Delta(w, u_2) = \delta_1(w) \cdot \mathbf{P}(s_1, w), \quad K_2 \cdot \sum_{u_1 \in U_1} \Delta(u_1, w) = \delta_2(w) \cdot \mathbf{P}(s_2, w)$$

   where $K_i = \sum_{u_i \in U_i} \delta_i(u_i) \cdot \mathbf{P}(s_i, u_i)$ for i=$1,2$
3. *for $u_1 \in U_1$, $u_1 \neq \perp$ there exists a path fragment $s_2, w_1, \ldots, w_n, u_2$ such that $n \geqslant 0$, $s_1 \, R \, w_j$, $0 < j \leqslant n$, and $u_1 \, R \, u_2$.*

s_2 weakly simulates s_1 in $\mathcal{D}$, denoted $s_1 \precsim_d s_2$, iff there exists a weak simulation R on $\mathcal{D}$ such that $s_1 \, R \, s_2$. ■

Note the correspondence to $\approx_d$ (cf. Def. 12), where $[s_1]_R$ plays the role of V_1, while the successors outside $[s_1]_R$ play the role of U_1, and the same for s_2, V_2 and U_2.

Example 2. In the following FPS we have $s_1 \precsim_d s_2$:

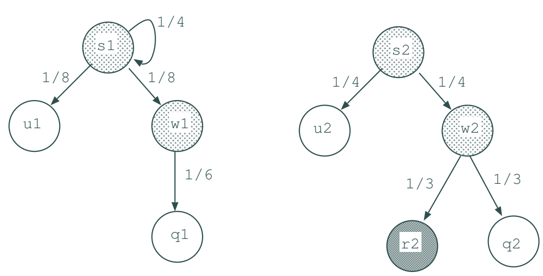

First, observe that $w_1 \precsim_d w_2$ since $R = \{ \, (q_1, q_2), (w_1, w_2) \, \}$ is a weak simulation, as we may deal with

- δ_1 the characteristic function of $U_1 = \{ q_1, \perp \}$ (and, thus, $V_1 = \varnothing$ and $K_1 = 1$)
- δ_2 the characteristic function of $U_2 = \{ r_2, q_2, \perp \}$ (and $V_2 = \varnothing$ and $K_2 = 1$)

and the weight function $\Delta(q_1,q_2) = \Delta(\perp,q_2) = \frac{1}{6}, \Delta(\perp,r_2) = \Delta(\perp,\perp) = \frac{1}{3}$. To establish a weak simulation for (s_1,s_2) consider the relation:

$$R = \{ (s_1,s_2),\ (u_1,u_2),\ (w_1,w_2),\ (q_1,q_2) \}$$

and put $V_1 = \{ \perp, s_1 \}$ and $V_2 = \varnothing$ while $U_i = \{ u_i, w_i, \perp \}$ where $\delta_1(\perp) = 1/2$, $\delta_i(u_i) = \delta_i(w_i) = \delta_2(\perp) = 1$. Then, $K_1 = \frac{1}{8} + \frac{1}{8} + \frac{1}{2} \cdot \frac{1}{2} = \frac{1}{2}$, $\quad K_2 = \frac{1}{4} + \frac{1}{4} + \frac{1}{2} = 1$. This yields the following distribution for the U-successors of s_1 and s_2: $u_1 : \frac{1}{4}$, $w_1 : \frac{1}{4}$, $\perp : \frac{1}{2}$, $u_2 : \frac{1}{4}$, $w_2 : \frac{1}{4}$, and $\perp : \frac{1}{2}$. Note that, e.g., $\frac{\delta_1(u_1) \cdot \mathbf{P}(s_1,u_1)}{K_1} = \frac{1}{4}$ and $\frac{\delta_1(\perp) \cdot \mathbf{P}(s_1,\perp)}{K_1} = \frac{1}{2}$. Hence, an appropriate weight function is: $\Delta(u_1,u_2) = \Delta(w_1,w_2) = \frac{1}{4}$, $\Delta(\perp,\perp) = \frac{1}{2}$, and $\Delta(\cdot) = 0$ for the remaining cases. Thus, according to Def. 14, R is a weak simulation. ■

Proposition 6. *For any FPS $\mathcal{D}$: $s_1 \approx_d s_2$ implies $s_1 \precapprox_d s_2$, and $s_1 \precsim_d s_2$ implies $s_1 \precapprox_d s_2$.*

Definition 15. **[8]** *Let $\mathcal{C} = (S, \mathbf{R}, L)$ be a CTMC and $R \subseteq S \times S$. R is a* weak simulation *on $\mathcal{C}$ iff for $s_1 R s_2$: $L(s_1) = L(s_2)$ and there exist $\delta_i : S \to [0,1]$ and $U_i, V_i \subseteq S$ (i=1,2) satisfying conditions 1. and 2. of Def. 14 (ignoring $\perp$) and the rate condition:*

$$\sum_{u_1 \in U_1} \delta_1(u_1) \cdot \mathbf{R}(s_1,u_1) \ \leqslant \sum_{u_2 \in U_2} \delta_2(u_2) \cdot \mathbf{R}(s_2,u_2)$$

s_2 weakly simulates s_1 in $\mathcal{C}$, denoted $s_1 \precapprox_c s_2$, iff there exists a weak simulation R on $\mathcal{C}$ such that $s_1 R s_2$. ■

The condition on the rates which replaces the reachability condition in FPSs states that s_2 is "faster than" s_1 in the sense that the total rate to move from s_2 to (the δ_2-part of) the U_2-states is at least the total rate to move from s_1 to (the δ_1-part of) the U_1-states. Note that $K_i \cdot E(s_i) = \sum_{u_i \in U_i} \delta_i(u_i) \cdot \mathbf{R}(s_i,u_i)$. Hence, the condition in Def. 15 can be rewritten as $K_1 \cdot E(s_1) \leqslant K_2 \cdot E(s_2)$. In particular, $K_2 = 0$ implies $K_1 = 0$. Therefore, a reachability condition as for weak simulation on FPSs is not needed here.

Proposition 7. *For CTMC $\mathcal{C}$ and states $s_1, s_2 \in S$:*

1. *$s_1 \precapprox_c s_2$ implies $s_1 \precapprox_d s_2$ in* emb$(\mathcal{C})$.
2. *$s_1 \approx_c s_2$ implies $s_1 \precapprox_c s_2$.*
3. *$\precapprox_c$ coincides with $\precsim_c$ in* unif$(\mathcal{C})$.

A few remarks are in order. Although $\precsim_c$ and $\precsim_d$ coincide for uniformised CTMCs (as $\precsim_c$ agrees with $\sim_c$, $\sim_c$ agrees with $\sim_d$, and $\sim_d$ agrees with $\precsim_d$), this does not hold for $\precapprox_d$ and $\precapprox_c$. For example, in:

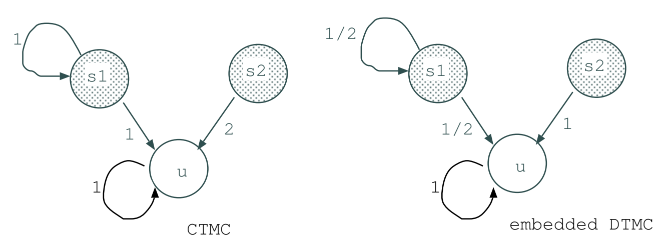

$s_2 \precapprox_d s_1$ in the embedded DTMC (on the right), but $s_2 \not\precapprox_c s_1$ in the CTMC (on the left), as the rate condition in Def. 15 is violated. Secondly, note that the analogue of Prop. 7.3 for $\precsim_c$ does not hold. This can be seen by considering the above embedded DTMC (on the right) as a uniformised CTMC. Finally, we note that although for uniformised CTMCs, $\sim_c$ and $\approx_c$ agree, a similar result for the simulation preorders does not hold. An example CTMC for which $s_1 \precapprox_c s_2$ but $s_1 \not\precsim_c s_2$ is depicted on the left. The fact that $s_1 \not\precsim_c s_2$ follows from the weight function condition in Def. 11. To see that $s_1 \precapprox_c s_2$, consider the reflexive closure R of $\{(s_1, s_2)\}$ and the partitioning $V_1 = \{s_2\}$, $V_2 = \{s_1\}$ and $U_1 = U_2 = \{u\}$ for which the conditions of a weak simulation are fulfilled.

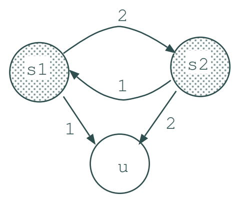

Theorem 1.

1. *For any FPS, weak simulation equivalence* $\precapprox_d \cap \precapprox_d^{-1}$ *coincides with* $\approx_d$.
2. *For any CTMC, weak simulation equivalence* $\precapprox_c \cap \precapprox_c^{-1}$ *coincides with* $\approx_c$.

4 Logical characterisations

PCTL. In Probabilistic CTL (PCTL) [26], state-formulas are interpreted over states of a FPS and path-formulas are interpreted over paths (i.e., sequences of states) in a FPS. The syntax of PCTL is as follows[1], where $\trianglelefteq \; \in \{\leqslant, \geqslant\}$:

$$\Phi ::= \mathrm{tt} \;\Big|\; a \;\Big|\; \Phi \wedge \Phi \;\Big|\; \neg\Phi \;\Big|\; \mathcal{P}_{\trianglelefteq p}(X\Phi) \;\Big|\; \mathcal{P}_{\trianglelefteq p}(\Phi\,\mathcal{U}\,\Phi) \;\Big|\; \mathcal{P}_{\trianglelefteq p}(\Phi\,\mathcal{W}\,\Phi)$$

where $p \in [0,1]$ and $a \in AP$. The satisfaction relation $\models$ is similar to CTL, where $s \models \mathcal{P}_{\trianglelefteq p}(\varphi)$ iff $\Pr(s,\varphi) \trianglelefteq p$. Here, $\Pr(s,\varphi)$ denotes the probability measure of the set of paths starting in state s fulfilling path-formula φ. As in CTL, X is the next-step operator, and the path-formula $\Phi\,\mathcal{U}\,\Psi$ asserts that Ψ will eventually be satisfied and that at all preceding states Φ holds (strong until). $\mathcal{W}$ is its weak counterpart, and does not require Ψ to eventually become true. The until-operator and the weak until-operator are closely related. For any PCTL-formula Φ and Ψ the following two formulae are equivalent:

$$\mathcal{P}_{\geqslant p}(\Phi\,\mathcal{W}\,\Psi) \;\equiv\; \mathcal{P}_{\leqslant 1-p}((\neg\Psi)\,\mathcal{U}\,\neg(\Phi \vee \Psi)).$$

A similar equivalence holds when the weak until- and the until-operator are swapped.

CSL. Continuous Stochastic Logic (CSL) [10] is a variant of the (identically named) logic by Aziz *et al.* [4] and extends PCTL by operators that reflect the real-time nature of CTMCs: a time-bounded until-operator and a steady-state operator. We focus here on a fragment of CSL where the time bounds of (weak) until are of the form "$\leqslant t$"; other time bounds can be handled by mappings on this case, cf. [6]. The syntax of CSL is, for real t, or $t = \infty$:

$$\Phi ::= \mathrm{tt} \;\Big|\; a \;\Big|\; \Phi \wedge \Phi \;\Big|\; \neg\Phi \;\Big|\; \mathcal{P}_{\trianglelefteq p}(X^{\leqslant t}\Phi) \;\Big|\; \mathcal{P}_{\trianglelefteq p}(\Phi\,\mathcal{U}^{\leqslant t}\,\Phi) \;\Big|\; \mathcal{P}_{\trianglelefteq p}(\Phi\,\mathcal{W}^{\leqslant t}\,\Phi) \;\Big|\; \mathcal{S}_{\trianglelefteq p}(\Phi)$$

[1] The bounded until-operator [26] is omitted here as for weak relations, FPSs are viewed as being time-abstract. For the strong relations on FPSs, this operator could, however, be considered without any problem.

To have a well-defined steady-state operator it is assumed that the steady-state probabilities in the CTMC do exist for any starting state. Intuitively, $\mathcal{S}_{\trianglelefteq p}(\Phi)$ asserts that on the long run, the probability for a Φ-state meets the bound $\trianglelefteq p$. The path-formula $\Phi\,\mathcal{U}^{\leqslant t}\,\Psi$ asserts that Ψ is satisfied at some time instant before t and that at all preceding time instants Φ holds (strong until). The connection between the until-operator and the weak until-operator is as in PCTL.

Logical characterisation of bisimulation. In both the discrete and the continuous setting, strong bisimulation ($\sim_d$ and $\sim_c$) coincides with logical equivalence (in PCTL and CSL, respectively) [3, 6, 19]. For weak bisimulation, the next-step operator is ignored, as it is not invariant with respect to stuttering. Let PCTL$_{\setminus X}$ denote the fragment of PCTL without the next-step operator; similarly, CSL$_{\setminus X}$ is defined. PCTL$_{\setminus X}$-equivalence (denoted $\equiv_{\text{PCTL}_{\setminus X}}$) and CSL$_{\setminus X}$-equivalence ($\equiv_{\text{CSL}_{\setminus X}}$) are defined in the obvious way.

Theorem 2. *For any FPS: $\approx_d$ coincides with PCTL$_{\setminus X}$-equivalence.*

Proof. By structural induction on the syntax of PCTL$_{\setminus X}$-formulae. We only consider the until operator. Let $\varphi = \Phi_1\,\mathcal{U}\,\Phi_2$. By the induction hypothesis we may assume that $Sat(\Phi_i)$ for $i{=}1,2$ is a disjoint union of equivalence classes under $\approx_d$. Let $B = [s]_{\approx_d}$. Then, $B \cap Sat(\Phi_i) = \varnothing$ or $B \subseteq Sat(\Phi_i)$. Only the cases $B \subseteq Sat(\Phi_1)$ and $B \cap Sat(\Phi_2) = \varnothing$ are of interest; otherwise, $\Pr(s_1,\varphi) = \Pr(s_2,\varphi) \in \{0,1\}$ for all s_1, $s_2 \in B$. Let S' be the set of states that reach a Φ_2-state via a (non-empty) Φ_1-path, i.e., $S' = \{\, s \in Sat(\Phi_1) \setminus Sat(\Phi_2) \mid \Pr(s,\varphi) > 0 \,\}$. It follows that S' is the disjoint union of equivalence classes under $\approx_d$.

We first observe the following. For $s \notin S'$, $\Pr(s,\varphi) \in \{0,1\}$. For $s \in S'$, the vector $\Big(\Pr(s,\varphi)\Big)_{s \in S'}$ is the *unique* solution of the equation system:

$$x_s = \mathbf{P}(s, Sat(\Phi_2)) \; + \sum_{s' \in Sat(\Phi_1) \setminus Sat(\Phi_2)} \mathbf{P}(s,s') \cdot x_{s'} \tag{1}$$

For any $\approx_d$-equivalence class $B \subseteq S'$, select $s_B \in B$ such that $\mathbf{P}(s_B,B) < 1$. Such state is guaranteed to exist, since if $\mathbf{P}(s,B)$ would equal 1 for any $s \in B$ then none of the B-states can reach a Φ_2-state, contradicting being in S'. Now consider the unique solution $(x_B)_{B \in S/\approx_d, B \subseteq S'}$ of the equation system:

$$x_B \;=\; \mathbf{P}(s_B, Sat(\Phi_2)) \;+ \sum_{\substack{C \in S/\approx_d \\ C \subseteq S'}} \mathbf{P}(s_B,C) \cdot x_C.$$

A calculation shows that the vector $(x_s)_{s \in S'}$ where $x_s = x_B$ if $s \in B$ is a solution to (1). Hence, $x_B = \Pr(s,\varphi)$ for all states $s \in B$.

The fact that PCTL$_{\setminus X}$-equivalence implies $\approx_d$ is proven as follows. W.l.o.g. we assume S to be finite and that any equivalence class C under $\equiv_{\text{PCTL}_{\setminus X}}$ is represented by a PCTL$_{\setminus X}$-formula Φ_C. (for infinite-state CTMCs approximations of master-formulae can be used). For PCTL$_{\setminus X}$ equivalence classes B and C with $B \neq C$, consider the path formulae $\varphi = \Phi_B\,\mathcal{U}\,\Phi_C$ and $\psi = \diamond\neg\Phi_B$. Then, $\Pr(s_1,\varphi) = \Pr(s_2,\varphi)$ and $\Pr(s_1,\psi) = \Pr(s_2,\psi)$ for any $s_1, s_2 \in B$. In particular, if $\mathbf{P}(s,B) < 1$ for some $s \in B$ then $\Pr(s,\psi) > 0$. Hence, for any $s' \in B$ there exists a path leading from s' to a state not in B. Assume that s_1, $s_2 \in B$ and that $\mathbf{P}(s_i,B) < 1$ for $i{=}1,2$. Then:

$$\Pr(s_i,\varphi) \;=\; \frac{\mathbf{P}(s_i,C)}{1 - \mathbf{P}(s_i,B)}.$$

This is justified as follows. If $\Pr(s_i,\varphi) = 0$, then obviously $\mathbf{P}(s_i,C) = 0$. Otherwise, by instantiating the equation system in (1) with $S' = B$, $\Phi_2 = \Phi_C$, $\Phi_1 = \Phi_B$ it can easily be verified that the vector with the values $x_s = \frac{\mathbf{P}(s,C)}{1-\mathbf{P}(s,B)}$ (for $s \in B$) is a solution. ■

Proposition 8. *For CTMC $\mathcal{C}$, s in $\mathcal{C}$, and $CSL_{\setminus X}$-formula Φ: $s \models \Phi$ iff $s \models \Phi$ in $unif(\mathcal{C})$.*

Proof. By induction on the syntax of Φ. For the propositional fragment the result is obvious. For the $\mathcal{S}$- and $\mathcal{P}$-operator, we exploit the fact that steady-state and transient distributions in $\mathcal{C}$ and $unif(\mathcal{C})$ are identical, and that the semantics of $\mathcal{U}^{\leqslant t}$ and $\mathcal{W}^{\leqslant t}$ agrees with transient distributions [6]. ■

Proposition 9. *For any uniformised CTMC: $\equiv_{CSL}$ coincides with $\equiv_{CSL_{\setminus X}}$.*

Proof. The direction "⇒" is obvious. We prove the other direction. Assume CTMC $\mathcal{C}$ is uniformised and s_1, s_2 be states in $\mathcal{C}$. From Prop. 1.1 and the logical characterisations of $\sim_c$ and $\sim_d$ it follows:

$$s_1 \equiv_{\mathrm{CSL}} s_2 \text{ iff } s_1 \sim_c s_2 \text{ iff } s_1 \sim_d s_2 \text{ iff } s_1 \equiv_{\mathrm{PCTL}} s_2.$$

Hence, it suffices to show that $\equiv_{\mathrm{CSL}_{\setminus X}}$ implies $\equiv_{\mathrm{PCTL}_{\setminus X}}$ (for uniformised CTMC). This is done by structural induction on the syntax of PCTL-formulae. Clearly, only the next step operator is of interest. Consider PCTL-path formula $\varphi = X\Phi$. By induction hypothesis $Sat(\Phi)$ is a (countable) union of equivalence classes of $\equiv_{\mathrm{CSL}_{\setminus X}}$. In the following, we establish for $s_1 \equiv_{\mathrm{CSL}_{\setminus X}} s_2$:

$$\mathbf{P}(s_1, Sat(\Phi)) = \mathbf{P}(s_2, Sat(\Phi)) \text{ that is } \Pr(s_1, X\Phi) = \Pr(s_2, X\Phi).$$

Let $B = [s_1]_{\equiv_{\mathrm{CSL}_{\setminus X}}} = [s_2]_{\equiv_{\mathrm{CSL}_{\setminus X}}}$. First observe that $\mathbf{P}(s_1, B) = \mathbf{P}(s_2, B)$; otherwise, if, e.g., $\mathbf{P}(s_1, B) < \mathbf{P}(s_2, B)$ one would have $\Pr(s_1, \Diamond^{\leqslant t}\neg\Phi_B) < \Pr(s_2, \Diamond^{\leqslant t}\neg\Phi_B)$ for some sufficiently small t, contradicting $s_1 \equiv_{\mathrm{CSL}_{\setminus X}} s_2$. As in the proof of Theorem 2 we assume a finite state space and that any $\equiv_{\mathrm{CSL}_{\setminus X}}$-equivalence class C can be characterised by $\mathrm{CSL}_{\setminus X}$ formula Φ_C. Distinguish:

- $\mathbf{P}(s_1, B) = \mathbf{P}(s_2, B) < 1$. Using the same arguments as in the proof of Theorem 2 we obtain:

$$\Pr(s_i, \Phi_B\, \mathcal{U}\, \Phi) \;=\; \frac{\mathbf{P}(s_i, Sat(\Phi))}{1 - \mathbf{P}(s_1, B)}, \;\; i = 1, 2.$$

  As $s_1 \equiv_{\mathrm{CSL}_{\setminus X}} s_2$ and $\Phi_B\, \mathcal{U}\, \Phi$ is a $\mathrm{CSL}_{\setminus X}$-path formula we get: $\Pr(s_1, \Phi_B\, \mathcal{U}\, \Phi) = \Pr(s_2, \Phi_B\, \mathcal{U}\, \Phi)$. Since $\mathbf{P}(s_1, B) = \mathbf{P}(s_2, B)$, it follows $\mathbf{P}(s_1, Sat(\Phi)) = \mathbf{P}(s_2, Sat(\Phi))$.
- $\mathbf{P}(s_1, B) = \mathbf{P}(s_2, B) = 1$. As $Sat(\Phi)$ is the union of equivalence classes under $\equiv_{\mathrm{CSL}_{\setminus X}}$, the intersection with B is either empty or equals B. For $i = 1, 2$: $\mathbf{P}(s_i, Sat(\Phi)) = 1$ if $B \subseteq Sat(\Phi)$ and 0 if $B \cap Sat(\Phi) = \varnothing$. Hence, $\mathbf{P}(s_1, Sat(\Phi) = \mathbf{P}(s_2, Sat(\Phi))$.

Thus, $s_1 \equiv_{\mathrm{PCTL}} s_2$. ■

Theorem 3. *For any CTMC: $\approx_c$ coincides with $CSL_{\setminus X}$-equivalence.*

Proof.

$$
\begin{array}{lll}
 & s_1 \approx_c^{\mathcal{C}} s_2 & \\
\text{iff} & s_1 \approx_c^{unif(\mathcal{C})} s_2 & \text{(by Prop. 4.3)} \\
\text{iff} & s_1 \sim_c^{unif(\mathcal{C})} s_2 & \text{(by Prop. 4.2)} \\
\text{iff} & s_1 \equiv_{\mathrm{CSL}}^{unif(\mathcal{C})} s_2 & \text{(since } \sim_c \text{ and CSL-equivalence coincide)} \\
\text{iff} & s_1 \equiv_{\mathrm{CSL}_{\setminus X}}^{unif(\mathcal{C})} s_2 & \text{(by Prop. 9)} \\
\text{iff} & s_1 \approx_{\mathrm{CSL}_{\setminus X}}^{\mathcal{C}} s_2 & \text{(by Prop. 8)}
\end{array}
$$

■

Logical characterisation of simulation. $\precsim_d$ for DTMCs without absorbing states equals $\sim_d$ [31], and hence, equals $\equiv_{\text{PCTL}}$. For FPS where $\precsim_d$ is non-symmetric and strictly coarser than $\sim_d$, a logical characterisation is obtained by considering a fragment of PCTL in the sense that $s_1 \precsim_d s_2$ iff all PCTL-safety properties that hold for s_2 also hold for s_1. A similar result can be established for $\precsim_c$ and a safe fragment of CSL.

Safe and live fragments of PCTL and CSL. In analogy to the universal and existential fragments of CTL, safe and live fragments of PCTL and CSL are defined as follows. We consider formulae in positive normal form, i.e., negations may only be attached to atomic propositions. In addition, only a restriced class of probability bounds is allowed in the probabilistic operator. The syntax of PCTL-safety formulae (denoted by Φ_S) is as follows:

$$\mathrm{tt} \;\Big|\; \mathrm{ff} \;\Big|\; a \;\Big|\; \neg a \;\Big|\; \Phi_S \wedge \Phi_S \;\Big|\; \Phi_S \vee \Phi_S \;\Big|\; \mathcal{P}_{\leqslant p}(X \Phi_L) \;\Big|\; \mathcal{P}_{\geqslant p}(\Phi_S \,\mathcal{W}\, \Phi_S) \;\Big|\; \mathcal{P}_{\leqslant p}(\Phi_L \,\mathcal{U}\, \Phi_L)$$

PCTL-liveness formulae (denoted by Φ_L) are defined as follows:

$$\mathrm{tt} \;\Big|\; \mathrm{ff} \;\Big|\; a \;\Big|\; \neg a \;\Big|\; \Phi_L \wedge \Phi_L \;\Big|\; \Phi_L \vee \Phi_L \;\Big|\; \mathcal{P}_{\leqslant p}(X \Phi_L) \;\Big|\; \mathcal{P}_{\geqslant p}(\Phi_L \,\mathcal{W}\, \Phi_L) \;\Big|\; \mathcal{P}_{\leqslant p}(\Phi_S \,\mathcal{U}\, \Phi_S)$$

As a result of the aforementioned relationship between $\mathcal{U}$ and $\mathcal{W}$, there is a duality between safety and liveness properties for PCTL, i.e., for any formula Φ_S there is a liveness property equivalent to $\neg\Phi_S$, and the same applies to liveness property Φ_L. Safe and live fragments of CSL are defined in an analogous way, where the steady-state operator is not considered, see [8].

Logical characterisation of simulation. Let $s_1 \precsim^{safe}_{\text{PCTL}} s_2$ iff for all PCTL-safety formulae Φ_S: $s_2 \models \Phi_S$ implies $s_1 \models \Phi_S$. Likewise, $s_1 \precapprox^{safe}_{\text{PCTL}_{\setminus X}} s_2$ iff this implication holds for all $\text{PCTL}_{\setminus X}$-safety formulae. The preorders $\precsim^{live}_{\text{PCTL}}$ and $\precapprox^{live}_{\text{PCTL}_{\setminus X}}$ are defined similarly, and the same applies for the preorders corresponding to the safe and live fragments of CSL and $\text{CSL}_{\setminus X}$. The first of the following results follows from a result by [17] for a variant of Hennessy-Milner logic. The fourth result has been reported in [8]. The same proof strategy can be used to prove the second and third result [9]. We conjecture that the converse of the third and fourth result also holds.

Theorem 4.

1. *For any FPS:* $\precsim_d$ *coincides with* $\precsim^{safe}_{PCTL}$ *and with* $\precsim^{live}_{PCTL}$.
2. *For any CTMC:* $\precsim_c$ *coincides with* $\precsim^{safe}_{CSL}$ *and with* $\precsim^{live}_{CSL}$.
3. *For any FPS:* $\precapprox_d \subseteq \precapprox^{safe}_{PCTL_{\setminus X}}$ *and* $\precapprox_d \subseteq \precapprox^{live}_{PCTL_{\setminus X}}$.
4. *For any CTMC:* $\precapprox_c \subseteq \precapprox^{safe}_{CSL_{\setminus X}}$ *and* $\precapprox_c \subseteq \precapprox^{live}_{CSL_{\setminus X}}$.

5 The branching-time spectrum

Summarising the results obtained in the literature together with our results in this paper yields the 3-dimensional spectrum of branching-time relations depicted in Fig. 1. All strong bisimulation relations are clearly contained within their weak variants, i.e., $\sim_d \subseteq \approx_d$ and $\sim_c \subseteq \approx_c$. The plane in the "front" (black arrows) represents the continuous-time setting, whereas the plane in the "back" (light blue or gray arrows) represents the

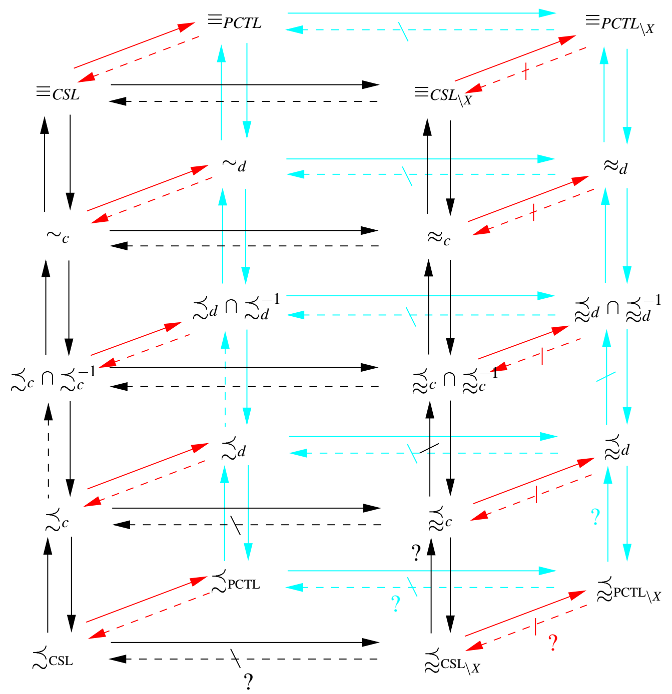

Fig. 1. Spectrum of branching-time relations for CTMCs and DTMCs

discrete-time setting. Arrows connecting the two planes (red or dark gray) relate CTMCs and their embedded DTMCs. $R \longrightarrow R'$ means that R is finer than R', while $R \not\longrightarrow R'$ means that R is not finer than R'. The dashed arrows in the continuous setting refer to uniformised CTMCs, i.e., if there is a dashed arrow from R to R', R is finer than R' for uniformised CTMCs. In the discrete-time setting the dashed arrows refer to DTMCs without absorbing states. Note that these models are obtained as embeddings of uniformised CTMCs (except for the pathological CTMC where all exit rates are 0, in which case all relations in the picture agree). If a solid arrow is labeled with a question mark, we claim the result, but have no proof (yet). For negated dashed arrows with a question mark, we claim that the implication does not hold even for uniformised CTMCs (DTMCs without absorbing states). The only difference between the discrete and continuous setting is that weak and strong bisimulation equivalence agree for uniformised CTMCs, but not for DTMCs without absorbing states.

The weak bisimulation proposed in [2] is strictly coarser than $\approx_d$, and thus does not preserve $\equiv_{\mathrm{PCTL}_{\setminus X}}$. The ordinary, non-probabilistic branching-time spectrum is more diverse, because there are many different weak bisimulation-style equivalences [23]. In the setting considered here, the spectrum spanned by Milner-style observational equivalence and branching bisimulation equivalence collapses to a single “weak bisimulation equivalence” [7]. Another difference is that for ordinary transition systems, simulation equivalence is strictly coarser than bisimulation equivalence. Further, in this non-probabilistic setting weak relations have to be augmented with aspects of divergence to obtain a logical characterisation by $CTL_{\setminus X}$ [37]. In the probabilistic setting, divergence occurs with probability 0 or 1, and does not need any distinguished treatment.

6 Concluding remarks

This paper has explored the spectrum of strong and weak (bi)simulation relations for countable fully probabilistic systems as well as continuous-time Markov chains. Based on a cascade of definitions in a uniform style, we have studied strong and weak (bi)simulations, and have provided logical characterisations in terms of fragments of PCTL and CSL. The definitions have three ingredients: (1) a condition on the labelling of states with atomic propositions, (2) a time-abstract condition on the probabilistic behaviour, and (3) a model-dependent condition: a rate condition for CTMCs (on the exit rates in the strong case, and on the total rates of "visible" moves in the weak case), and a reachability condition on the "visible" moves in the weak FPS case. The strong FPS case does not require a third condition.

As the rate conditions imply the corresponding reachability condition, the "continuous" relations are finer than their "discrete" counterparts, and the continuous-time setting excludes the possibility to abstract from stuttering occurring with probability 1.[2] While weak bisimulation in CTMCs (and FPSs) is a rather fine notion, it is the best abstraction preserving all properties that can be specified in CSL (PCTL) without next-step.

References

1. L. de Alfaro. Temporal logics for the specification of performance and reliability. *Symp. on Th. Aspects of Comp. Sc.*, LNCS 1200, pp. 165–176, 1997.
2. S. Andova and J. Baeten. Abstracion in probabilistic process algebra. *Tools and Algorithms for the Construction and Analysis of Systems*, LNCS 2031, pp. 204–219, 2001.
3. A. Aziz, V. Singhal, F. Balarin, R. Brayton and A. Sangiovanni-Vincentelli. It usually works: the temporal logic of stochastic systems. *CAV*, LNCS 939, pp. 155–165, 1995.
4. A. Aziz, K. Sanwal, V. Singhal and R. Brayton. Verifying continuous time Markov chains. *Computer-Aided Verification*, LNCS 1102, pp. 269–276, 1996.
5. C. Baier. On algorithmic verification methods for probabilistic systems. Habilitation thesis, University of Mannheim, 1998.
6. C. Baier, B.R. Haverkort, H. Hermanns and J.-P. Katoen. Model checking continuous-time Markov chains by transient analysis. *CAV*, LNCS 1855, pp. 358–372, 2000.
7. C. Baier, H. Hermanns. Weak bisimulation for fully probabilistic processes. *Computer-Aided Verification*, LNCS 1254, pp. 119-130. 1997.
8. C. Baier, J.-P. Katoen, H. Hermanns and B. Haverkort. Simulation for continuous-time Markov chains. *Concurrency Theory*, LNCS 2421, pp. 338-354, 2002.
9. C. Baier, H. Hermanns, J.-P. Katoen and V. Wolf. Comparative branching-time semantics for Markov chains. Tech. Rep., Univ. of Twente, 2003.
10. C. Baier, J.-P. Katoen and H. Hermanns. Approximate symbolic model checking of continuous-time Markov chains. *Concurrency Theory*, LNCS 1664, pp. 146–162, 1999.
11. M. Bernardo and R. Gorrieri. Extended Markovian process algebra. *Concurrency Theory*, LNCS 1119, pp. 315-330, 1996.

[2] In process-algebraic terminology, the reachability condition guarantees the law $\tau.P = P$ for FPS. This law cannot hold for CTMCs due to the advance of time while stuttering (performing τ).

12. M. Bravetti. Revisiting interactive Markov chains. *3rd Workshop on Models for Time-Critical Systems*, BRICS Notes NP-02-3, pp. 68–88, 2002.
13. M. Brown, E. Clarke, O. Grumberg. Characterizing finite Kripke structures in propositional temporal logic. *Th. Comp. Sc.*, **59**: 115–131, 1988.
14. P. Buchholz. Exact and ordinary lumpability in finite Markov chains. *J. of Appl. Prob.*, **31**: 59–75, 1994.
15. E. Clarke, O. Grumberg and D.E. Long. Model checking and abstraction. *ACM Tr. on Progr. Lang. and Sys.*, **16**(5): 1512–1542, 1994.
16. J. Desharnais. *Labelled Markov Processes.* PhD Thesis, McGill University, 1999.
17. J. Desharnais. Logical characterisation of simulation for Markov chains. *Workshop on Probabilistic Methods in Verification*, Tech. Rep. CSR-99-8, Univ. of Birmingham, pp. 33–48, ,1999.
18. J. Desharnais, A. Edalat and P. Panangaden. A logical characterisation of bisimulation for labelled Markov processes. In *IEEE Symp. on Logic in Comp. Sc.*, pp. 478–487, 1998.
19. J. Desharnais and P. Panangaden. Continuous stochastic logic characterizes bisimulation of continuous-time Markov processes. *J. of Logic and Alg. Progr.*, **56**: 99–115, 2003.
20. J. Desharnais, V. Gupta, R. Jagadeesan and P. Panangaden. Approximating labelled Markov processes. In *IEEE Symp. on Logic in Comp. Sc.*, pp. 95–106, 2000.
21. J. Desharnais, V. Gupta, R. Jagadeesan and P. Panangaden. Weak bisimulation is sound and complete for PCTL*. *Concurrency Theory*, LNCS 2421, pp. 355-370, 2002.
22. R.J. van Glabbeek. The linear time - branching time spectrum I. The semantics of concrete, sequential processes. Ch. 1 in *Handbook of Process Algebra*, pp. 3–100, 2001.
23. R.J. van Glabbeek. The linear time - branching time spectrum II. The semantics of sequential processes with silent moves. *Concurrency Theory*, LNCS 715, pp. 66–81, 1993.
24. R.J. van Glabbeek, S.A. Smolka and B. Steffen. Reactive, generative, and stratified models of probabilistic processes. *Inf. & Comp.*, **121**: 59–80, 1995.
25. W. Feller. *An Introduction to Probability Theory and its Applications.* John Wiley, 1968.
26. H. Hansson and B. Jonsson. A logic for reasoning about time and reliability. *Form. Asp. of Comp.* **6**: 512–535, 1994.
27. H. Hermanns. *Interactive Markov Chains.* LNCS 2428, 2002.
28. J. Hillston. *A Compositional Approach to Performance Modelling*. Cambr. Univ. Press, 1996.
29. C. Jones and G. Plotkin. A probabilistic powerdomain of evaluations. In *IEEE Symp. on Logic in Computer Science*, pp. 186–195, 1989.
30. B. Jonsson. Simulations between specifications of distributed systems. *Concurrency Theory*, LNCS 527, pp. 346–360, 1991.
31. B. Jonsson and K.G. Larsen. Specification and refinement of probabilistic processes. In *IEEE Symp. on Logic in Comp. Sc.*, pp. 266-277, 1991.
32. C.-C. Jou and S.A. Smolka. Equivalences, congruences, and complete axiomatizations for probabilistic processes. *Concurrency Theory*, LNCS 458, pp. 367–383, 1990.
33. J.G. Kemeny and J.L. Snell. *Finite Markov Chains*. Van Nostrand, 1960.
34. V.G. Kulkarni. *Modeling and Analysis of Stochastic Systems*. Chapman & Hall, 1995.
35. K.G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Inf. and Comp.*, **94**(1): 1–28, 1991.
36. R. Milner. *Communication and Concurrency*. Prentice-Hall, 1989.
37. R. De Nicola and F. Vaandrager. Three logics for branching bisimulation (extended abstract). In *IEEE Symp. on Logic in Comp. Sc.*, pp. 118–129, 1992.
38. A. Philippou, I. Lee, and O. Sokolsky. Weak bisimulation for probabilistic systems. *Concurrency Theory*, LNCS 1877, pp. 334–349, 2000.
39. M.L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, 1994.
40. R. Segala and N.A. Lynch. Probabilistic simulations for probabilistic processes. *Nordic J. of Computing*, **2**(2): 250–273, 1995.
41. M.I.A. Stoelinga. *Verification of Probabilistic, Real-Time and Parametric Systems.* PhD Thesis, University of Nijmegen, 2002.