

Expected-Delay-Summing Weak Bisimilarity for Markov Automata

Alessandro Aldini Marco Bernardo

Dipartimento di Scienze di Base e Fondamenti, Università di Urbino, Italy

A new weak bisimulation semantics is defined for Markov automata that, in addition to abstracting from internal actions, sums up the expected values of consecutive exponentially distributed delays possibly intertwined with internal actions. The resulting equivalence is shown to be a congruence with respect to parallel composition for Markov automata. Moreover, it turns out to be comparable with weak bisimilarity for timed labeled transition systems, thus constituting a step towards reconciling the semantics for stochastic time and deterministic time.


1 Introduction

Markov automata [5] integrate Segala’s simple probabilistic automata [13] and Hermanns’ interactive Markov chains [7], thus resulting in very expressive models. Markov automata feature two types of transitions, one for *action execution* and one for *time passing*. The choice among the actions enabled in a state is nondeterministic, the execution of the selected action is instantaneous, and the reached state is established according to a probability distribution. Time passing is described by means of exponentially distributed delays governed by the race policy, with the execution of internal actions taking precedence over such delays so to enforce maximal progress.

Markov automata come equipped with two compositional semantics, respectively based on *strong* and *weak* bisimilarity. While the former is the obvious combination of strong bisimilarity for probabilistic automata and strong bisimilarity for interactive Markov chains, the latter is more complicated due to certain desirable identifications that should be achieved when abstracting from internal actions. This has been accomplished by suitably defining weak bisimilarity over state probability subdistributions rather than over individual states [5]. The resulting equivalence has been shown to provide a sound and complete proof methodology for a touchstone equivalence called reduction barbed congruence [4].

Weak bisimilarity for Markov automata abstracts from *internal instantaneous actions*. However, in the setting of labeled transition systems enriched with deterministic delays, which are at the basis of models such as timed automata [2], the weak bisimulation semantics appeared in the literature (e.g., [15, 11, 1]) are also capable of abstracting from *sequences of delays* possibly intertwined with internal actions, in the sense that those delays can be reduced to a single one equal to the *sum* of the original delays. Some work in this direction has recently been done for Markovian process calculi with durational actions that, unlike Markov automata, feature neither nondeterminism nor probabilistic branching. More precisely, in [3] a weak bisimilarity has been proposed, which is capable of abstracting from internal actions that are exponentially timed by summing up their *expected delays*.

The purpose of this paper is to study an *expected-delay-summing* weak bisimulation semantics in the more expressive setting of Markov automata. As we will see, defining such a semantics is a challenging task due to the need of balancing disparate demands related to nondeterministic, probabilistic, and timing behaviors. Additionally, in [3] a tradeoff has emerged between compositionality, i.e., being a congruence with respect to parallel composition, and pseudo-aggregation exactness, i.e., preserving stationary-state

Figure 1: Merging τ -transitions, timed transitions, and nondeterministic choices


reward-based performance measures, in the sense that in the Markovian setting it is not possible to define an expected-delay-summing weak equivalence that enjoys both properties. These facts make it far from trivial to embody the expected-delay-summing capability of the weak semantics of [3] into the weak semantics originally developed for Markov automata in [5].

To clarify the additional identifications that we would like to obtain with respect to [5], let us consider a few illustrative examples. In these examples, we will depict states as circles, action transitions as arrows labeled with a, b, c for visible actions and τ for internal actions, timed transitions as arrows labeled with positive real rates λ, μ representing the inverses of expected delays, and probability distributions as dashed lines connecting states.

The first Markov automaton in Fig. 1(a) has an initial a -transition and a final b -transition, with a sequence of timed and τ -transitions in between. All the transitions in the sequence are reduced to a single timed transition in the second Markov automaton, whose rate has been computed as the inverse of the expected duration of the entire sequence: $(\frac{1}{\lambda} + \frac{1}{\mu})^{-1} = \frac{\lambda \cdot \mu}{\lambda + \mu}$. Also the third Markov automaton can be reduced to the second one, as both branches of the internal nondeterministic choice at state u_3 have the same expected duration $\frac{1}{\mu}$.

The two Markov automata in Fig. 1(b) can be identified as well. The timed and τ -transitions preceding and following the internal probabilistic choice after state s_3 can be reduced to two alternative timed transitions, whose basic rate $\frac{\lambda \cdot \mu}{\lambda + \mu}$ is respectively multiplied by p and $1 - p$.

In Fig. 2, the focus is on stochastic choices governed by the race policy, according to which the execution probability of a timed transition is proportional to its rate. The first Markov automaton has a stochastic choice at state s_3 between a μ_1 -transition and a μ_2 -transition, whilst the third Markov automaton has the same stochastic choice at state s'_2 . Both automata can be seen as equivalent to the second one, in which the stochastic choice and the timed and τ -transitions preceding and following it are reduced to two alternative timed transitions, whose basic rate $(\frac{1}{\lambda} + \frac{1}{\mu_1 + \mu_2})^{-1}$ is respectively multiplied by $\frac{\mu_1}{\mu_1 + \mu_2}$ – probability of taking the left branch – and $\frac{\mu_2}{\mu_1 + \mu_2}$ – probability of taking the right branch.

Figure 2: Merging τ -transitions, timed transitions, and stochastic choices

The paper is organized as follows. In Sect. 2, we present some background material and recall the definition of weak bisimilarity for Markov automata provided in [5]. In Sect. 3, we introduce our expected-delay-summing weak bisimilarity for Markov automata by extending the existing one with the construction of [3], we prove that it is a congruence with respect to parallel composition, and we compare its discriminating power with that of the weak bisimilarity of [5]. In Sect. 4, we propose a comparison between our relation and the weak bisimilarity for timed labeled transition systems defined in [15, 11], as both equivalences are capable of summing up expected delays while abstracting from internal actions. Finally, in Sect. 5 we provide some concluding remarks.

2 Background

2.1 Discrete Probability Subdistributions

Let Δ be a function from a nonempty, at most countable set S to $\mathbb{R}_{[0,1]}$. The support of Δ is defined as $\text{supp}(\Delta) = \{s \in S \mid \Delta(s) > 0\}$, while the size of Δ is defined as $\text{size}(\Delta) = \Delta(S)$ where, in general, $\Delta(S') = \sum_{s' \in S'} \Delta(s')$ for all $S' \subseteq S$.

Function Δ is a *discrete probability (sub)distribution* over S iff $\text{size}(\Delta) = 1$ (resp. $\text{size}(\Delta) \leq 1$). We denote by $\text{Subdistr}(S)$ and $\text{Distr}(S)$ the sets of subdistributions and distributions over S . Furthermore, we indicate with δ_s the Dirac distribution for $s \in S$, i.e., $\delta_s(s) = 1$ and $\delta_s(s') = 0$ for all $s' \in S \setminus \{s\}$.

Given $x \in \mathbb{R}_{\geq 0}$ and $\Delta \in \text{Subdistr}(S)$, we denote by $x \odot \Delta$ the function defined by $(x \odot \Delta)(s) = x \cdot \Delta(s)$ for all $s \in S$. Given $\Delta_1, \Delta_2 \in \text{Subdistr}(S)$, we denote by $\Delta_1 \oplus \Delta_2$ the function defined by $(\Delta_1 \oplus \Delta_2)(s) = \Delta_1(s) + \Delta_2(s)$. These functions are subdistributions when their size does not exceed 1. Moreover, given $\Delta_1 \in \text{Subdistr}(S_1)$ and $\Delta_2 \in \text{Subdistr}(S_2)$, we denote by $\Delta_1 \otimes \Delta_2$ the subdistribution over $S_1 \times S_2$ defined by $(\Delta_1 \otimes \Delta_2)(s_1, s_2) = \Delta_1(s_1) \cdot \Delta_2(s_2)$.

A subdistribution Δ over S can be viewed as a subset of $S \times \mathbb{R}_{[0,1]}$, in which only elements of $\text{supp}(\Delta)$ occur, each once with the corresponding probability. In other words, subdistribution Δ may be written as $\llbracket (s, p) \mid s \in \text{supp}(\Delta) \wedge p = \Delta(s) \rrbracket$. We also denote by $\Delta \ominus s$ the subdistribution that is obtained from Δ by removing the pair $(s, \Delta(s))$ when $s \in \text{supp}(\Delta)$.

2.2 Markov Automata

Markov automata [5] have two distinct types of transitions: *action transitions* and *Markov timed transitions*. The choice among the action transitions departing from a given state is nondeterministic. This choice can be influenced by the external environment, except for transitions labeled with the internal action τ . Once an action transition is chosen, the next state is internally selected according to some probability distribution, as in probabilistic automata [13].

Each Markov timed transition is labeled with a real number called *rate*, which uniquely identifies an exponentially distributed delay. As with interactive Markov chains [7], the choice among the Markov timed transitions departing from a given state is governed by the *race policy*, which means that the Markov timed transition that is executed is the one sampling the least duration. Therefore, the execution probability of a Markov timed transition is proportional to its rate, and the sojourn time associated with a state having outgoing Markov timed transitions is exponentially distributed with rate given by the sum of the rates of those transitions.

Different from [5], where Markov automata were introduced for the first time, in the definition below we explicitly build some assumptions into model.

Definition 2.1 A *Markov automaton (MA)* is a tuple $(S, A, \longrightarrow, \dashrightarrow)$ where:

- S is a nonempty, at most countable set of states.
- A is a set of actions containing at least the internal action τ .
- $\longrightarrow \subseteq S \times A \times \text{Distr}(S)$ is an action-transition relation.
- $\dashrightarrow \subseteq S \times \mathbb{R}_{\geq 0} \times S$ is a time-transition relation such that for all $s \in S$:
 - If $s \xrightarrow{0} s'$ for some $s' \in S$, then $s' = s$ (*zero speed*).
 - $\sum_{(s, \lambda, s') \in \dashrightarrow} \lambda < \infty$ (*speed boundedness*).
 - If $s \xrightarrow{\tau} \Delta$ with $\Delta \in \text{Distr}(S)$, then $s \xrightarrow{\lambda} s'$ for all $\lambda \in \mathbb{R}_{\geq 0}$ and $s' \in S$ (*maximal progress*). ■

Also the notion of parallel composition, although equivalent to the one in [5], is formulated in a slightly different way. We recall that the first of the three conditions below about the time-transition relation ensures that the parallel composition of two Markov timed selfloops, each having the same rate λ , results in a Markov timed selfloop with rate $\lambda + \lambda$, as established by the race policy, instead of λ . In the following, $s \not\rightarrow$ and $s \not\rightarrow$ stand for the absence of action transitions and Markov timed transitions, respectively, out of state s ; an action or rate decoration of the negative arrow means the absence of transitions labeled with that action or rate.

Definition 2.2 Let $\mathcal{M}_k = (S_k, A_k, \longrightarrow_k, \dashrightarrow_k)$ be an MA with initial state $s_{0,k}$ for $k = 1, 2$. The *parallel composition* of \mathcal{M}_1 and \mathcal{M}_2 with respect to an action set $A \subseteq (A_1 \cup A_2) \setminus \{\tau\}$ is the MA $\mathcal{M}_1 \parallel_A \mathcal{M}_2 = (S_1 \times S_2, A_1 \cup A_2, \longrightarrow, \dashrightarrow)$ with initial state $(s_{0,1}, s_{0,2})$ such that:

- $(s_1, s_2) \xrightarrow{a} \Delta$ iff one of the following conditions is fulfilled:
 - $a \in A$, $s_1 \xrightarrow{a}_1 \Delta_1$, $s_2 \xrightarrow{a}_2 \Delta_2$, and $\Delta = \Delta_1 \otimes \Delta_2$.
 - $a \notin A$, $s_1 \xrightarrow{a}_1 \Delta_1$, and $\Delta = \Delta_1 \otimes \delta_{s_2}$.
 - $a \notin A$, $s_2 \xrightarrow{a}_2 \Delta_2$, and $\Delta = \delta_{s_1} \otimes \Delta_2$.

- $(s_1, s_2) \xrightarrow{\lambda} (s'_1, s'_2)$ iff one of the following conditions is fulfilled:
 - $s_1 \xrightarrow{\mu}_1 s'_1, s_2 \xrightarrow{\gamma}_2 s'_2, (s'_1, s'_2) = (s_1, s_2)$, and $\lambda = \sum_{(s_1, \mu', s'_1) \in \xrightarrow{\mu}_1} \mu' + \sum_{(s_2, \gamma', s'_2) \in \xrightarrow{\gamma}_2} \gamma'$.
 - $s_1 \xrightarrow{\mu}_1 s'_1, s'_2 = s_2, s'_1 \neq s_1$ or $s_2 \not\xrightarrow{\gamma}_2, s_2 \xrightarrow{\tau}_2$, and $\lambda = \sum_{(s_1, \mu', s'_1) \in \xrightarrow{\mu}_1} \mu'$.
 - $s_2 \xrightarrow{\gamma}_2 s'_2, s'_1 = s_1, s'_2 \neq s_2$ or $s_1 \not\xrightarrow{\mu}_1, s_1 \xrightarrow{\tau}_1$, and $\lambda = \sum_{(s_2, \gamma', s'_2) \in \xrightarrow{\gamma}_2} \gamma'$. ■

Similar to [5], from now on we uniformly treat the action transitions and the Markov timed transitions of an MA $(S, A, \xrightarrow{\cdot}, \xrightarrow{\cdot})$ by considering time passage as a special action χ and viewing the MA as a triple $(S, A^\chi, \xrightarrow{\cdot})$ where:

- $A^\chi = A \cup \{\chi(\lambda) \mid \lambda \in \mathbb{R}_{\geq 0}\}$.
- $s \xrightarrow{\alpha} \Delta$ iff either $\alpha \in A$ with an identical action transition being present in the original MA, or the following conditions are met:
 - $s \not\xrightarrow{\tau}$.
 - $\alpha = \chi(\lambda)$ with $\lambda = \sum_{(s, \mu, s') \in \xrightarrow{\mu}} \mu$.
 - If $\lambda > 0$, then $\Delta(s') = \sum_{(s, \mu', s') \in \xrightarrow{\mu'}} \mu' / \lambda$ for all $s' \in S$, else $\Delta = \delta_s$.

Notice that a state $s \in S$ can have at most one outgoing transition of the form $s \xrightarrow{\chi(\lambda)} \Delta$, with λ and Δ being consistent with the race policy governing $\xrightarrow{\cdot}$. This can be considered a probabilistic form of time determinism. In case of parallel composition, it yields $(s_1, s_2) \xrightarrow{\chi(\lambda)} \Delta$ iff one of the following holds:

- $s_1 \xrightarrow{\chi(\mu)}_1 \Delta_1, s_2 \xrightarrow{\chi(\gamma)}_2 \Delta_2, \lambda = \mu + \gamma$, and $\Delta = (\frac{\mu}{\lambda} \odot (\Delta_1 \otimes \delta_{s_2})) \oplus (\frac{\gamma}{\lambda} \odot (\delta_{s_1} \otimes \Delta_2))$.
- $s_1 \xrightarrow{\chi(\lambda)}_1 \Delta_1, s_2 \not\xrightarrow{\chi(\gamma)}_2$ for all $\gamma \in \mathbb{R}_{\geq 0}, s_2 \not\xrightarrow{\tau}_2$, and $\Delta = \Delta_1 \otimes \delta_{s_2}$.
- $s_2 \xrightarrow{\chi(\lambda)}_2 \Delta_2, s_1 \not\xrightarrow{\chi(\mu)}_1$ for all $\mu \in \mathbb{R}_{\geq 0}, s_1 \not\xrightarrow{\tau}_1$, and $\Delta = \delta_{s_1} \otimes \Delta_2$.

2.3 Internal Transition Trees and Weak Transitions

The definition of weak bisimilarity for Markov automata given in [5] relies on *labeled trees*. For each such tree \mathcal{T} , we denote by $nodes(\mathcal{T})$ the set of its nodes, by $leaves(\mathcal{T})$ the set of its leaves, and by $\varepsilon_{\mathcal{T}}$ its root. When \mathcal{T} only contains $\varepsilon_{\mathcal{T}}$, the node $\varepsilon_{\mathcal{T}}$ is considered a leaf. Given $\sigma \in nodes(\mathcal{T})$, we denote by $children(\sigma)$ the set of nodes reachable in one step from σ .

For an MA $(S, A^\chi, \xrightarrow{\cdot})$, we use $(S \times \mathbb{R}_{[0,1]} \times (A^\chi \cup \{\perp\}))$ -labeled trees that somehow represent *resolutions of nondeterminism*. Each node σ is labeled with the corresponding state $Sta(\sigma) \in S$, the execution probability $Prob(\sigma) \in \mathbb{R}_{[0,1]}$ of the only path from $\varepsilon_{\mathcal{T}}$ to σ , and the action $Act(\sigma) \in A^\chi \cup \{\perp\}$ chosen to proceed.

Definition 2.3 Let $\mathcal{M} = (S, A^\chi, \xrightarrow{\cdot})$ be an MA. A *transition tree* \mathcal{T} for \mathcal{M} is an $(S \times \mathbb{R}_{[0,1]} \times (A^\chi \cup \{\perp\}))$ -labeled tree that satisfies the following conditions:

1. $Prob(\varepsilon_{\mathcal{T}}) = 1$.
2. For each $\sigma \in leaves(\mathcal{T})$, $Act(\sigma) = \perp$.

3. For each $\sigma \in \text{nodes}(\mathcal{T}) \setminus \text{leaves}(\mathcal{T})$, there is Δ such that $\text{Sta}(\sigma) \xrightarrow{\text{Act}(\sigma)} \Delta$ with:
 $\text{Prob}(\sigma) \odot \Delta = \llbracket (\text{Sta}(\sigma'), \text{Prob}(\sigma')) \mid \sigma' \in \text{children}(\sigma) \rrbracket$

The distribution induced by \mathcal{T} on its leaves is defined as:

$$\Delta_{\mathcal{T}} = \bigoplus_{\sigma \in \text{leaves}(\mathcal{T})} \llbracket (\text{Sta}(\sigma), \text{Prob}(\sigma)) \rrbracket$$

We say that \mathcal{T} is *internal* iff, for each $\sigma \in \text{nodes}(\mathcal{T})$, $\text{Act}(\sigma) \in \{\tau, \perp\}$. ■

Given $s \in S$ and $\Delta \in \text{Distr}(S)$, *weak transitions* based on internal transition trees and variants thereof are introduced as follows:

- $s \Longrightarrow \Delta$ iff Δ is induced by an internal transition tree \mathcal{T} with $\text{Sta}(\varepsilon_{\mathcal{T}}) = s$.
- $s \xrightarrow{\alpha} \Delta$ iff Δ is induced by a transition tree \mathcal{T} with $\text{Sta}(\varepsilon_{\mathcal{T}}) = s$, such that along every maximal path from $\varepsilon_{\mathcal{T}}$:
 - the action label of at least one inner node is α if $\alpha = \tau$,
 - the action label of precisely one inner node is α if $\alpha \neq \tau$,

while the action label of all the other inner nodes is τ .

- $s \xrightarrow{\hat{\alpha}} \Delta$ iff either $\alpha = \tau$ and $s \Longrightarrow \Delta$, or $\alpha \neq \tau$ and $s \xrightarrow{\alpha} \Delta$.

Requiring the target Δ of a weak transition to be a *full* distribution ensures that all the paths of the tree inducing the weak transition are of finite length, or that all of its infinite paths have probability 0. In the first case, the tree is not necessarily finite, because some node may have countably many children.

Convex combinations of identically labeled weak transitions are defined as follows: $s \xrightarrow{\alpha}_c \Delta$ iff there exist $n \in \mathbb{N}_{\geq 1}$, $(p_i \in \mathbb{R}_{[0,1]} \mid 1 \leq i \leq n)$, and $(s \xrightarrow{\alpha} \Delta_i \mid 1 \leq i \leq n)$ such that $\sum_{1 \leq i \leq n} p_i = 1$ and $\Delta = \bigoplus_{1 \leq i \leq n} p_i \odot \Delta_i$. Combined weak transition relations \Longrightarrow_c and $\xrightarrow{\hat{\alpha}}_c$ are defined similarly.

2.4 Strong and Weak Bisimilarities

Strong bisimilarity for Markov automata is a straightforward combination of strong bisimilarity for probabilistic automata [13] and strong bisimilarity for interactive Markov chains [7].

Definition 2.4 Let $(S, A^{\mathcal{X}}, \longrightarrow)$ be an MA. An equivalence relation \mathcal{B} over S is a *strong bisimulation* iff, whenever $(s_1, s_2) \in \mathcal{B}$, then for all $\alpha \in A^{\mathcal{X}}$ it holds that for each $s_1 \xrightarrow{\alpha} \Delta_1$ there exists $s_2 \xrightarrow{\alpha} \Delta_2$ such that $\Delta_1(C) = \Delta_2(C)$ for all $C \in S/\mathcal{B}$. We write $s_1 \sim s_2$ to denote that (s_1, s_2) is contained in some strong bisimulation. ■

The mix of the weak bisimilarities for the two classes of models is too fine for Markov automata. In [5], this drawback has been overcome by using combined weak transitions lifted to *subdistributions*. Each such α -transition is obtained by weighting the target distribution of the α -transition from each source state with the probability assigned to that state by the source subdistribution: $\Delta \xrightarrow{\alpha}_c \Psi$ iff $s \xrightarrow{\alpha}_c \Psi_s$ for all $s \in \text{supp}(\Delta)$ and $\Psi = \bigoplus_{s \in \text{supp}(\Delta)} \Delta(s) \odot \Psi_s$. Combined weak transition relations \Longrightarrow_c and $\xrightarrow{\hat{\alpha}}_c$ are lifted similarly.

Definition 2.5 Let $(S, A^{\mathcal{X}}, \longrightarrow)$ be an MA. A relation \mathcal{B} over $\text{Subdistr}(S)$ is a *weak bisimulation* iff, whenever $(\Delta_1, \Delta_2) \in \mathcal{B}$, then $\text{size}(\Delta_1) = \text{size}(\Delta_2)$ and for all $\alpha \in A^{\mathcal{X}}$ it holds that:

- (a) For each $s_1 \in \text{supp}(\Delta_1)$ there exist $\Delta'_2, \Delta''_2 \in \text{Subdistr}(S)$ such that:

1. $\Delta_2 \Longrightarrow_c \Delta'_2 \oplus \Delta''_2$ with $(\llbracket (s_1, \Delta_1(s_1)) \rrbracket, \Delta'_2) \in \mathcal{B}$ and $((\Delta_1 \ominus s_1), \Delta''_2) \in \mathcal{B}$.
2. For each $s_1 \xrightarrow{\alpha} \Psi_1$ there exists $\Delta'_2 \xrightarrow{\hat{\alpha}}_c \Psi_2$ such that $(\Delta_1(s_1) \odot \Psi_1, \Psi_2) \in \mathcal{B}$.

(b) Symmetric clause with the roles of Δ_1 and Δ_2 interchanged.

We write $\Delta_1 \approx \Delta_2$ to denote that (Δ_1, Δ_2) is contained in some weak bisimulation. Moreover, we let $s_1 \approx s_2$ iff $\delta_{s_1} \approx \delta_{s_2}$. \blacksquare

3 Expected-Delay-Summing Weak Bisimilarity

In this section, we introduce a new weak bisimilarity \approx_{eds} for Markov automata that, in addition to abstracting from τ -actions as \approx , sums up the expected values of consecutive exponentially distributed delays possibly intertwined with τ -actions. This is accomplished by relying on *reducible* projected transition trees. We prove that \approx_{eds} is an equivalence relation and a congruence with respect to parallel composition, then we investigate the relationships between \approx_{eds} and \approx .

3.1 Projected Transition Trees: Components and Durations

In general, a system is made out of several interacting sequential components. Therefore, we view a (global) state of the MA at hand as a vector of local states. We denote by $\text{Sta}(\sigma)[\ell]$ the state related to sequential component ℓ that occurs in the label of node σ of a transition tree associated with the MA. As shown in [3] for exponentially timed actions, this component view is necessary to achieve the congruence property with respect to parallel composition in the case of a weak bisimilarity that adds up the expected values of exponentially distributed delays.

Furthermore, we extend transition trees by considering labels taken from the set $S \times \mathbb{R}_{[0,1]} \times \mathbb{R}_{\geq 0} \times (A^\chi \cup \{\perp\})$, where each node σ is additionally labeled with the expected duration $\text{Expd}(\sigma) \in \mathbb{R}_{\geq 0}$ of the only path from the root to σ . To be precise, the states labeling the various nodes are *global*. In contrast, to support compositionality, the probability, the expected duration, and the action chosen to proceed labeling a node are *local* to the behavior of the considered component ℓ in isolation. Every path of such a tree corresponds to a computation of component ℓ that is not forbidden by synchronization constraints or maximal progress; the local transitions in that computation will be decorated with ℓ .

Definition 3.1 Let $\mathcal{M} = (S, A^\chi, \longrightarrow)$ be an MA resulting from the parallel composition of $n \in \mathbb{N}_{\geq 1}$ MAs; we say that \mathcal{M} is *sequential* when $n = 1$. Let $\ell \in \{1, \dots, n\}$. An ℓ -*projected transition tree* \mathcal{T} for \mathcal{M} is an $(S \times \mathbb{R}_{[0,1]} \times \mathbb{R}_{\geq 0} \times (A^\chi \cup \{\perp\}))$ -labeled tree that satisfies the following conditions:

1. $\text{Prob}(\varepsilon_{\mathcal{T}}) = 1$.
2. $\text{Expd}(\varepsilon_{\mathcal{T}}) = 0$.
3. For each $\sigma \in \text{leaves}(\mathcal{T})$, $\text{Act}(\sigma) = \perp$.
4. For each $\sigma \in \text{nodes}(\mathcal{T}) \setminus \text{leaves}(\mathcal{T})$, there is Δ_ℓ such that the local transition $\text{Sta}(\sigma)[\ell] \xrightarrow{\text{Act}(\sigma)}_\ell \Delta_\ell$ contributes to the derivation of some global transition $\text{Sta}(\sigma) \xrightarrow{\alpha} \Delta$ such that $\sigma' \in \text{children}(\sigma)$ iff $\text{Sta}(\sigma') \in \text{supp}(\Delta)$ with:

$$\text{Prob}(\sigma) \odot \Delta_\ell = \llbracket (\text{Sta}(\sigma')[\ell], \text{Prob}(\sigma')) \mid \sigma' \in \text{children}(\sigma) \rrbracket$$

and for each $\sigma' \in \text{children}(\sigma)$ it holds that:

$$\text{Expd}(\sigma') = \begin{cases} \text{Expd}(\sigma) + \frac{1}{\lambda} & \text{if } \text{Act}(\sigma) = \chi(\lambda) \text{ and } \lambda > 0 \\ \text{Expd}(\sigma) & \text{if } \text{Act}(\sigma) \in A \cup \{\chi(0)\} \end{cases}$$

\blacksquare

Variants of weak transitions based on internal ℓ -projected transition trees are defined as in Sect. 2.3 and are respectively denoted by \Longrightarrow_ℓ , $\xRightarrow{\alpha}_\ell$, and $\xRightarrow{\tilde{\alpha}}_\ell$.

Since the leaves of \mathcal{T} are labeled with the expected duration of the corresponding paths from the root, the distribution $\Delta_{\mathcal{T}}$ induced by \mathcal{T} on its leaves can be decomposed into duration-indexed subdistributions. To this aim, we define the subdistribution of $\Delta_{\mathcal{T}}$ associated with $t \in \mathbb{R}_{\geq 0}$ as:

$$\Delta_{\mathcal{T}}^t = \bigoplus_{\sigma \in \{\sigma' \in \text{leaves}(\mathcal{T}) \mid \text{Expd}(\sigma') = t\}} \llbracket (Sta(\sigma), Prob(\sigma)) \rrbracket$$

so that:

$$\Delta_{\mathcal{T}} = \bigoplus_{t \in \{\text{Expd}(\sigma) \mid \sigma \in \text{leaves}(\mathcal{T})\}} \Delta_{\mathcal{T}}^t.$$

For the sake of convenience, we will often aggregate probabilities associated with leaves that are labeled with the same expected duration by employing the notation:

$$\Delta_{\mathcal{T}} = \bigoplus_{i \in I} \Upsilon^i$$

for some indexed set $\{(t_i, \Upsilon^i)\}_{i \in I}$.

The above decomposition also extends to lifted combined weak transitions induced by internal variants of ℓ -projected transition trees. If $\Delta \xRightarrow{\alpha}_{\ell, c} \Psi$, with $s \xRightarrow{\alpha}_{\ell, c} \Psi_s$ for all $s \in \text{supp}(\Delta)$ and $\Psi = \bigoplus_{s \in \text{supp}(\Delta)} \Delta(s) \odot \Psi_s$, then, assuming $\Psi_s = \bigoplus_{i \in I_s} \Psi_s^i$, we have that $\Psi = \bigoplus_{s \in \text{supp}(\Delta)} \bigoplus_{i \in I_s} \Delta(s) \odot \Psi_s^i$. This can again be expressed in a more compact form as $\Psi = \bigoplus_{i \in I} \Upsilon^i$ for some indexed set $\{(t_i, \Upsilon^i)\}_{i \in I}$, provided that we join all subdistributions Ψ_s^i with the same t_i .

3.2 Reducible Projected Transition Trees: Intertwining τ and $\chi(\lambda)$

We now extend the notion of internal ℓ -projected transition tree – and add the corresponding weak transitions – by admitting nodes labeled with actions of the form $\chi(\lambda)$ that alternate with nodes whose action label is τ . The construction of this kind of tree proceeds as long as, in ℓ , the traversed local states have no alternative local transitions labeled with visible actions. In contrast, the local states of ℓ contributing to the global states associated with the leaves cannot have local transitions labeled with τ or $\chi(\lambda)$, unless they have a visible transition too. Following the terminology of [3], we call the resulting tree reducible.

Taking into account alternative local transitions labeled with visible actions is crucial to achieve an equivalence that does not abstract away from observable behaviors. For instance, should the context change along the sequence of transitions between s_1 and s_6 of Fig. 1, fusing those transitions into a unique Markov timed transition with the same expected duration would not be appropriate. If there were an a' -transition from s_2 and a b' -transition from s_3 , after an exponentially distributed delay with rate λ we would notice a context change in the first MA that cannot take place in the second one, thus preventing the transitions in the sequence from being merged.

Definition 3.2 An ℓ -projected transition tree \mathcal{T} for an MA $(S, A^\chi, \longrightarrow)$ is *reducible* iff:

1. For each $\sigma \in \text{nodes}(\mathcal{T}) \setminus \text{leaves}(\mathcal{T})$, $\text{Act}(\sigma) \in \{\tau\} \cup \{\chi(\lambda) \mid \lambda \in \mathbb{R}_{\geq 0}\}$.
2. For each $\sigma \in \text{nodes}(\mathcal{T}) \setminus \text{leaves}(\mathcal{T})$, $\text{Sta}(\sigma)[\ell] \xrightarrow{a}_\ell$ for all $a \in A \setminus \{\tau\}$.
3. For each $\sigma \in \text{leaves}(\mathcal{T})$, $\text{Sta}(\sigma)[\ell] \xrightarrow{\alpha}_\ell$ for all $\alpha \in \{\tau\} \cup \{\chi(\lambda) \mid \lambda \in \mathbb{R}_{\geq 0}\}$ or $\text{Sta}(\sigma)[\ell] \xrightarrow{a}_\ell$ for some $a \in A \setminus \{\tau\}$.
4. Along every maximal path from $\varepsilon_{\mathcal{T}}$, the action label of at least one inner node belongs to $\{\chi(\lambda) \mid \lambda \in \mathbb{R}_{\geq 0}\}$. ■

Given $s \in S$ and $\Delta \in \text{Distr}(S)$, we write $s \xRightarrow{\chi}_\ell \Delta$ iff Δ is induced by a reducible ℓ -projected transition tree \mathcal{T} with $\text{Sta}(\varepsilon_{\mathcal{T}}) = s$. Combined and lifted variants of $\xRightarrow{\chi}_\ell$ are defined as usual.




Figure 3: Reducing sequences of exponentially distributed delays

Example 3.3 We start by considering systems made out of a single sequential component, for which we omit decoration ℓ from transitions. In Fig. 3(a), we have $s_0 \xrightarrow{\chi} \llbracket (s_1, 1) \rrbracket$, with the expected duration of the unique path being equal to $\frac{1}{\lambda}$. In Fig. 3(b), we have $t_0 \xrightarrow{\chi} \llbracket (t_2, 1) \rrbracket$ with the expected duration of the path leading from t_0 to t_2 being equal to $\frac{1}{2\lambda} + \frac{1}{2\lambda} = \frac{1}{\lambda}$.

The two distributions induced by the two reducible projected transition trees will allow us to identify s_0 and t_0 by disregarding the intermediate state t_1 . This would not be possible in the absence of the third condition of Def. 3.2. In that case, we would also have $t_0 \xrightarrow{\chi} \llbracket (t_1, 1) \rrbracket$, whose associated expected duration $\frac{1}{2\lambda}$ could not be matched by any weak transition from s_0 .

Also the second condition of Def. 3.2, which does not admit alternative visible transitions along reducible computations, plays a role here. This is especially important for the intermediate state t_1 . If this state had an alternative a -transition, in the absence of the second condition s_0 and t_0 would again be identified, in spite of the fact that an external observer could see, at a certain point in time, the execution of a visible action only in the second system. ■

Example 3.4 Let us go back to Sect. 1. The weak transition $s_1 \xrightarrow{\chi} \llbracket (s_6, 1) \rrbracket$ in Fig. 1(a) is induced by a reducible projected transition tree that associates the expected duration $\frac{1}{\lambda} + \frac{1}{\mu} = \frac{\lambda + \mu}{\lambda \cdot \mu}$ with the path from the root to the only leaf. This is the same expected duration associated with the weak transition $t_1 \xrightarrow{\chi} \llbracket (t_2, 1) \rrbracket$, as well as the two occurrences of the weak transition $u_1 \xrightarrow{\chi} \llbracket (z_0, 1) \rrbracket$ induced by two alternative reducible projected transition trees deriving from the nondeterministic choice at u_3 . Therefore, we will be able to relate s_0 , t_0 , and u_0 .

In Fig. 1(b), we have the weak transition $s_1 \xrightarrow{\chi} \llbracket (u_2, p), (v_2, 1-p) \rrbracket$. The expected duration associated with the two leaves labeled with states u_2 and v_2 is $\frac{1}{\lambda} + \frac{1}{\mu} = \frac{\lambda + \mu}{\lambda \cdot \mu}$. This is also the expected sojourn time associated with t_1 , i.e., the inverse of $p \cdot \frac{\lambda \cdot \mu}{\lambda + \mu} + (1-p) \cdot \frac{\lambda \cdot \mu}{\lambda + \mu} = \frac{\lambda \cdot \mu}{\lambda + \mu}$. Thus, $t_1 \xrightarrow{\chi} \llbracket (w_0, p), (z_0, 1-p) \rrbracket$ will allow us to equate s_0 and t_0 .

In Fig. 2, the reducible projected transition tree whose root is associated with state s_1 has two leaves, labeled with states u_1 and v_1 , respectively. By virtue of the race condition at s_3 , the probability of reaching such leaves is $\frac{\mu_1}{\mu_1 + \mu_2}$ and $\frac{\mu_2}{\mu_1 + \mu_2}$, respectively, while the expected duration of both paths is $\frac{1}{\lambda} + \frac{1}{\mu_1 + \mu_2}$. Hence, we have $s_1 \xrightarrow{\chi} \llbracket (u_1, \frac{\mu_1}{\mu_1 + \mu_2}), (v_1, \frac{\mu_2}{\mu_1 + \mu_2}) \rrbracket$. By anticipating the stochastic choice involving μ_1 and μ_2 , we obtain an analogous result, as $s'_1 \xrightarrow{\chi} \llbracket (u'_1, \frac{\mu_1}{\mu_1 + \mu_2}), (v'_1, \frac{\mu_2}{\mu_1 + \mu_2}) \rrbracket$ with the expected duration being as before.

Notice the analogy with the weak transition $t_1 \xrightarrow{\chi} \llbracket (w_0, \frac{\mu_1}{\mu_1 + \mu_2}), (z_0, \frac{\mu_2}{\mu_1 + \mu_2}) \rrbracket$. By applying the race policy, we have that the expected sojourn time in t_1 is again $\frac{1}{\lambda} + \frac{1}{\mu_1 + \mu_2}$, from which we derive that s_0 , s'_0 , and t_0 can be identified. ■

Example 3.5 We now consider systems built from several components like the two MAs in Figs. 3(c) and (d), which are respectively obtained through the parallel composition of the two MAs in Figs. 3(a) and (b) with another MA that has only an a -transition. Each of the four states of the MA in Fig. 3(c) can be paired with a state of the MA in Fig. 3(d) as follows: (u_0, v_0) , (u_1, v_2) , (u'_0, v'_0) , (u'_1, v'_2) .

This is possible because the quantities associated with the nodes of reducible projected transition trees are computed locally to the considered components. If the additional MA had a Markov timed transition with rate μ instead of an a -transition, and the probabilities were computed globally, then the global probability of going from u_0 to u_1 (which is $\frac{\lambda}{\lambda+\mu}$) and the global probability of going from v_0 to v_2 (which is $2 \cdot \frac{2\lambda}{2\lambda+\mu}$) would be considered, which are different from each other. ■

3.3 A New Weak Bisimilarity: Definition and Properties

We are finally in the position of defining a new weak bisimilarity for MAs that sums up expected values of exponentially distributed delays while abstracting from τ -actions. This is accomplished by considering reducible projected transition trees *in addition to* internal transition trees.

It is useful to extend to global states and distributions the notation \longrightarrow_ℓ for local transitions employed in Defs. 3.1 and 3.2. In the following definition, we write $s \xrightarrow{\chi(\lambda)}_\ell \Delta$ to intend that Δ is induced by an ℓ -projected transition tree \mathcal{T} such that $\text{Sta}(\varepsilon_{\mathcal{T}}) = s$, $\text{Act}(\varepsilon_{\mathcal{T}}) = \chi(\lambda)$, and $\text{children}(\varepsilon_{\mathcal{T}}) = \text{leaves}(\mathcal{T})$. Weak, combined, and lifted variants of the extended notation are as expected.

Definition 3.6 Let $(S, A^\chi, \longrightarrow)$ be an MA. A relation \mathcal{B} over $\text{Subdistr}(S)$ is an *expected-delay-summing weak bisimulation* iff, whenever $(\Delta_1, \Delta_2) \in \mathcal{B}$, then $\text{size}(\Delta_1) = \text{size}(\Delta_2)$ and for all transition labels in $A^\chi \cup \{\chi\}$ it holds that:

- (a) For each $s_1 \in \text{supp}(\Delta_1)$ there exist $\Delta'_2, \Delta''_2 \in \text{Subdistr}(S)$ such that:
 1. $\Delta_2 \Longrightarrow_c \Delta'_2 \oplus \Delta''_2$ with $(\llbracket (s_1, \Delta_1(s_1)) \rrbracket, \Delta'_2) \in \mathcal{B}$ and $((\Delta_1 \ominus s_1), \Delta''_2) \in \mathcal{B}$.
 2. For each $s_1 \xrightarrow{a} \Psi_1$ there exists $\Delta'_2 \xrightarrow{\hat{a}}_c \Psi_2$ such that $(\Delta_1(s_1) \odot \Psi_1, \Psi_2) \in \mathcal{B}$.
 3. For each $s_1 \xrightarrow{\chi(\lambda)}_{\ell_1} \Psi_1$ such that $s_1 \not\xrightarrow{\chi}_{\ell_1}$, there exists $\Delta'_2 \xrightarrow{\chi(\lambda)}_{\ell_2, c} \Psi_2$ such that $(\Delta_1(s_1) \odot \Psi_1, \Psi_2) \in \mathcal{B}$.
 4. For each $s_1 \xrightarrow{\chi}_{\ell_1} \bigoplus_{i \in I} \Upsilon_1^{t_i}$ there exists $\Delta'_2 \xrightarrow{\chi}_{\ell_2, c} \bigoplus_{i \in I} \Upsilon_2^{t_i}$ such that $(\Delta_1(s_1) \odot \Upsilon_1^{t_i}, \Upsilon_2^{t_i}) \in \mathcal{B}$ for all $i \in I$.
- (b) Symmetric clause with the roles of Δ_1 and Δ_2 interchanged.

We write $\Delta_1 \approx_{\text{eds}} \Delta_2$ to denote that (Δ_1, Δ_2) is contained in some expected-delay-summing weak bisimulation. Moreover, we let $s_1 \approx_{\text{eds}} s_2$ iff $\delta_{s_1} \approx_{\text{eds}} \delta_{s_2}$. ■

Condition 1 of Def. 3.6 coincides with condition 1 of Def. 2.5, while condition 2 of Def. 2.5 is split into conditions 2 and 3 of Def. 3.6. In particular, by virtue of condition 3 above, Markov timed transitions locally enabled by component ℓ_1 are treated according to \approx whenever component ℓ_1 does not enable weak transitions induced by a *reducible* ℓ_1 -projected transition tree. This is the case when the Markov timed transition is locally alternative to a visible action transition, or in the presence of time-divergence, i.e., an infinite sequence of Markov timed transitions, because the third condition of Def. 3.2 prevents reducible projected transition trees from being generated as long as no state is encountered that enables a visible action or has no outgoing transitions.

On the other hand, condition 4 above states that every weak transition induced by a *reducible* ℓ_1 -projected transition tree must be matched by a weak transition induced by a *reducible* ℓ_2 -projected transition tree, where ℓ_1 is a component of the first process while ℓ_2 is a component of the second process. Notice that the target distributions are decomposed into duration-indexed subdistributions prior to the application of the expected-delay-summing weak bisimulation check, so to ensure that the matching also takes expected durations into account.

Example 3.7 Following the discussions in Exs. 3.3, 3.4, and 3.5, we can establish that:

- $s_0 \approx_{\text{eds}} t_0 \approx_{\text{eds}} u_0$ in Fig. 1(a). For instance, since $t_3 \approx_{\text{eds}} z_1$, it is easy to verify that $\mathcal{B} = \{(\delta_{t_0}, \delta_{u_0}), (\delta_{t_1}, \delta_{u_1}), (\delta_{t_2}, \delta_{z_0}), (\delta_{t_3}, \delta_{z_1})\}$ is an expected-delay-summing weak bisimulation. First, when applying Def. 3.6, for any pair $(\Delta_1, \Delta_2) \in \mathcal{B}$ no splitting of Δ_2 is needed (i.e., $\Delta_2 = \Delta'_2$). Then, $(\delta_{t_2}, \delta_{z_0}) \in \mathcal{B}$ as a consequence of condition 2 of Def. 3.6 and of the fact that $t_3 \approx_{\text{eds}} z_1$; $(\delta_{t_1}, \delta_{u_1}) \in \mathcal{B}$ follows by applying condition 4 of Def. 3.6 (in particular, consider the weak transitions shown in Ex. 3.4); $(\delta_{t_0}, \delta_{u_0})$ follows by applying again condition 2 of Def. 3.6.
- $s_0 \approx_{\text{eds}} t_0$ in Fig. 1(b).
- $s_0 \approx_{\text{eds}} t_0 \approx_{\text{eds}} s'_0$ in Fig. 2.
- $s_0 \approx_{\text{eds}} t_0$ and $u_0 \approx_{\text{eds}} v_0$ in Fig. 3. For instance, consider the two MAs in Figs. 3(c) and (d), which are obtained as discussed in Ex. 3.5. In particular, the former results from the parallel composition of components ℓ_1 and ℓ_3 , which are the MA of Fig. 3(a) and an MA that has only an a -transition, respectively, while the latter from the parallel composition of components ℓ_2 , which is the MA of Fig. 3(b), and ℓ_3 . Since $u'_1 \approx_{\text{eds}} v'_2$, it holds that $\mathcal{B} = \{(u_0, v_0), (u_1, v_2), (u'_0, v'_0), (u'_1, v'_2)\}$ is an expected-delay-summing weak bisimulation. On one hand, $u_0 \xrightarrow{\chi}_{\ell_1} \llbracket (u_1, 1) \rrbracket$, with the expected duration of the unique path being equal to $\frac{1}{\lambda}$, while $v_0 \xrightarrow{\chi}_{\ell_2} \llbracket (v_2, 1) \rrbracket$, with the expected duration of the unique path being equal to $\frac{1}{2\lambda} + \frac{1}{2\lambda} = \frac{1}{\lambda}$. Then, u_1 (resp., v_2) enables an a -transition leading to u'_1 (resp., v'_2). On the other hand, both u_0 and v_0 enable an a -transition, leading to u'_0 and v'_0 , respectively. Then, similarly as above, $u'_0 \xrightarrow{\chi}_{\ell_1} \llbracket (u'_1, 1) \rrbracket$ is matched by $v'_0 \xrightarrow{\chi}_{\ell_2} \llbracket (v'_2, 1) \rrbracket$. Notice that \mathcal{B} does not include pairs containing the intermediate states v_1 and v'_1 . On one hand, this is necessary to match the weak transitions of the two MAs. On the other hand, this is sufficient as, by virtue of the interleaving semantics of parallel composition, the visible transitions enabled in v_1 (resp., v'_1) are the same as those enabled in v_0 and v_2 (resp., v'_0 and v'_2). ■

The relation \approx_{eds} turns out to be reflexive, symmetric, transitive, and substitutive with respect to parallel composition.

Theorem 3.8 Let $\mathcal{M} = (S, A^\chi, \longrightarrow)$ be an MA. Then \approx_{eds} is an equivalence relation over $\text{Subdistr}(S)$. ■

In order for the congruence property to hold, as pointed out in [12] it is essential to generate $\chi(0)$ -selfloops for those states of the MAs at hand having neither τ -transitions nor Markov timed transitions. If such $\chi(0)$ -selfloops were not generated by the transformation at the end of Sect. 2.2, then the interplay among maximal progress, τ -divergence (i.e., an infinite sequence of τ -transitions), and Markov timed transitions would break compositionality. This can be seen by considering a τ -convergent MA with two states connected by a τ -transition, see Fig. 4(a), and a τ -divergent MA with a single state featuring a τ -selfloop, see Fig. 4(b), each composed in parallel with the MA illustrated in Fig. 4(c), which has two states connected by a Markov timed transition. If no $\chi(0)$ -selfloop were added to the final state




Figure 4: Interplay among maximal progress, τ -divergence, and Markov timed transitions

of the first MA, as shown in Fig. 4(d), then the two MAs of Figs. 4(a) and (b) would be identified, but the two composed MAs would be told apart. In fact, notice that the parallel composition of the MAs of Figs. 4(a) and (c) enables the Markov timed transition after the execution of the τ -transition, while the parallel composition of the MAs of Figs. 4(b) and (c) executes τ -transitions only. On the other hand, the two MAs of Figs. 4(a) and (d) are not identified by \approx_{eds} , as the Markov timed transition of the latter cannot be matched by the former.

Analogous considerations on compositionality and sensitivity to τ -divergence were also made in the IMC setting of [7] and the pure Markovian setting of [3]. In contrast, in the MA setting of [4] the problem was circumvented by using a parallel composition operator that requires all components to let time advance, in the same spirit as the deterministically timed model of [10, 15].

Theorem 3.9 Let $\mathcal{M}_k = (S_k, A_k^\lambda, \longrightarrow_k)$ be an MA for $k = 1, 2$. Let $A \subseteq (A_1 \cup A_2) \setminus \{\tau\}$ and consider the parallel composition $\mathcal{M}_1 \parallel_A \mathcal{M}_2$. Let $s_1, s'_1 \in S_1$ and $s_2 \in S_2$. If $s_1 \approx_{\text{eds}} s'_1$, then $(s_1, s_2) \approx_{\text{eds}} (s'_1, s_2)$. ■

The following theorem states that \approx_{eds} is a conservative extension of \approx for sequential MAs. These MAs constitute the common ground of the two equivalences, given that summing up expected delays can be done compositionally only if the component structure is elicited. Investigating the relation between the two equivalences in the case of generic MAs remains an open challenge unless renouncing to the compositionality result of Thm. 3.9. Indeed, on one hand, condition 4 of Def. 3.6 is critical to achieve the congruence property. On the other hand, it imposes local conditions over the reducible behaviors of matching components that are completely ignored by \approx , as this equivalence abstracts away from the component-based structure of the system.

Theorem 3.10 Let $(S, A^\lambda, \longrightarrow)$ be a sequential MA. Let $\Delta_1, \Delta_2 \in \text{Subdistr}(S)$. If $\Delta_1 \approx \Delta_2$, then $\Delta_1 \approx_{\text{eds}} \Delta_2$. ■

4 Application to Timed Labeled Transition Systems

So far, we have considered time passing described through *exponentially distributed delays*, which is typical of shared-resource systems. In this section, we consider a timed extension of labeled transition systems inspired by [10, 15], which is based on *fixed delays* as in real-time systems.

Definition 4.1 A *timed labeled transition system (TLTS)* is a tuple $(S, A, \longrightarrow, \dashrightarrow)$ where:

- S is a nonempty, possibly uncountable set of states.
- A is a set of actions containing at least the internal action τ .
- $\longrightarrow \subseteq S \times A \times S$ is an action-transition relation.
- $\dashrightarrow \subseteq S \times \mathbb{R}_{\geq 0} \times S$ is a time-transition relation such that for all $s \in S$:

- If $s \xrightarrow{0} s'$ for some $s' \in S$, then $s' = s$ (*zero delay*).
- If $s \xrightarrow{t} s'_1$ and $s \xrightarrow{t} s'_2$ for some $s'_1, s'_2 \in S$ and $t \in \mathbb{R}_{\geq 0}$, then $s'_1 = s'_2$ (*time determinism*).
- $s \xrightarrow{t_1+t_2} s''$ iff $s \xrightarrow{t_1} s'$ and $s' \xrightarrow{t_2} s''$ (*time additivity*).
- If $s \xrightarrow{\tau} s'$ for some $s' \in S$, then $s \not\rightarrow$ (*maximal progress*). ■

By analogy with MA, the passage of time $t \in \mathbb{R}_{\geq 0}$ can be viewed as a special action that, instead of $\varepsilon(t)$ as in [15], we denote by $\chi(t)$. Under this view, from now on we consider a TLTS as a triple $(S, A^\chi, \longrightarrow)$.

A notion of weak bisimilarity for TLTSs was studied in [15, 11]. It is essentially based on Milner's weak bisimilarity plus the capability of summing up fixed delays while abstracting from τ -actions. Weak transitions are defined as follows:

- $\Longrightarrow = (\xrightarrow{\tau})^*$.
- $\xrightarrow{a} = \Longrightarrow \xrightarrow{a} \Longrightarrow$.
- $\xrightarrow{\hat{a}} = \begin{cases} \Longrightarrow & \text{if } a = \tau \\ \xrightarrow{a} & \text{if } a \neq \tau \end{cases}$.
- $\xrightarrow{\chi(t)} = \Longrightarrow \xrightarrow{\chi(t_1)} \Longrightarrow \dots \Longrightarrow \xrightarrow{\chi(t_n)} \Longrightarrow$ where $t = \sum_{1 \leq i \leq n} t_i$ for $n \in \mathbb{N}_{\geq 1}$.

Definition 4.2 Let $(S, A^\chi, \longrightarrow)$ be a TLTS. A symmetric relation \mathcal{B} over S is a *timed weak bisimulation* iff, whenever $(s_1, s_2) \in \mathcal{B}$, then for all actions $a \in A$ and amounts of time $t \in \mathbb{R}_{\geq 0}$ it holds that:

- For each $s_1 \xrightarrow{a} s'_1$ there exists $s_2 \xrightarrow{\hat{a}} s'_2$ such that $(s'_1, s'_2) \in \mathcal{B}$.
- For each $s_1 \xrightarrow{\chi(t)} s'_1$ there exists $s_2 \xrightarrow{\chi(t)} s'_2$ such that $(s'_1, s'_2) \in \mathcal{B}$.

We write $s_1 \approx_t s_2$ to denote that (s_1, s_2) is contained in some timed weak bisimulation. ■

Relations \approx_{eds} and \approx_t share the idea of summing up delays while abstracting from τ -actions. The theorem below performs a more precise comparison based on an adaptation of the \approx_{eds} construction to TLTS models suitably modified according to the following considerations:

- We regard TLTS transitions as leading to Dirac distributions over states.
- We assume that actions can only be instantaneously (as opposed to continuously) enabled.
- We assume that all $\chi(t)$ -transitions have $t \in \mathbb{R}_{>0}$ and are expressed up to time decomposability (an aspect of time additivity), in the sense that a transition of the form $s \xrightarrow{\chi(t)} s'$ with $t \in \mathbb{R}_{>0}$ subsumes all the possible computations from s to s' whose total duration is t , such that each of the intermediate states at time distance $t' \in \mathbb{R}_{]0,t[}$ from s does not enable any action and has only computations to s' of total duration $t - t'$. Due to time determinism, this guarantees that every TLTS state will have at most one outgoing transition labeled with $\chi(t)$, where $t \in \mathbb{R}_{>0}$.
- When building transition trees, the expected duration of the only path from the root to a node is computed as the sum of the time delays – rather than their inverses as in the fourth condition of Def. 3.1 – occurring along that path. The reason is that, in a deterministically timed setting, a delay coincides with its expected value.

Theorem 4.3 Let $(S, A^\chi, \longrightarrow)$ be a modified TLTS originated from a system made out of a single sequential component, whose states have at most one outgoing τ -transition each. Let $s_1, s_2 \in S$. If $s_1 \approx_{\text{eds}} s_2$, then $s_1 \approx_t s_2$. ■

The result does not hold in the presence of choices among τ -transitions, because \approx_t is more sensitive to them than \approx_{eds} . For example, given $a, b \in A \setminus \{\tau\}$, the systems described in process algebraic style as $\chi(t_1 + t_2) \cdot (\tau \cdot a + \tau \cdot b)$ and $\chi(t_1) \cdot (\tau \cdot \chi(t_2) \cdot a + \tau \cdot \chi(t_2) \cdot b)$ are identified by \approx_{eds} , but distinguished by \approx_t . Intuitively, there is no reason to distinguish them on the basis of the instant of time in which the internal choice is solved, as in both cases an external observer should wait $t_1 + t_2$ time units before interacting with the system.

Moreover, the above implication cannot be reversed. For instance, given $a \in A \setminus \{\tau\}$, the process $a + \chi(t_1 + t_2)$ and the process $a + \chi(t_1) \cdot \chi(t_2)$ are identified by \approx_t , but distinguished by \approx_{eds} . The reason is that the latter cannot sum up delays in the presence of *locally* alternative visible actions.

5 Conclusions

Building on previous work [5, 3], we have incrementally extended the identification power of a weak semantics for MAs by defining a new weak bisimulation congruence \approx_{eds} that, in addition to abstracting from τ -actions, sums up the expected values of consecutive exponentially distributed delays possibly intertwined with τ -actions. From an application viewpoint, \approx_{eds} can thus serve as the semantical basis for state space reduction techniques more aggressive than those recently developed for MAs in [14].

The relation \approx_{eds} has also been compared with the weak bisimilarity defined in [15, 11] for TLTSs; in these models, it is standard to be capable of adding up expected delays interleaved with τ -actions. Therefore, the definition of \approx_{eds} constitutes a step towards reconciling the semantics for stochastic time and deterministic time, a subject recently addressed in [9].

As far as future work is concerned, we plan to investigate equational and logical characterizations of \approx_{eds} . Furthermore, since we have privileged the achievement of the congruence property with respect to the tradeoff emerged in [3], we intend to examine the preservation of quantitative properties. This has been addressed in [3] for stationary-state reward-based performance measures in the case of pure Markovian models. An analogous result in the case of MAs needs to take into account nondeterminism and hence the fact that only maximum and minimum values of performance measures can be computed after applying suitable schedulers [6]. Finally, we would like to adapt \approx_{eds} to a probabilistic extension of the TLTS model, in which the target of an action transition can be a general probability distribution over states as in the MA model, so to study the relationships with the weak bisimilarity defined in [8].

Acknowledgment: This work has been funded by MIUR-PRIN project CINA.

References

- [1] L. Aceto & D. Murphy (1996): *Timing and Causality in Process Algebra*. *Acta Informatica* 33, pp. 317–350.
- [2] R. Alur & D.L. Dill (1994): *A Theory of Timed Automata*. *Theoretical Computer Science* 126, pp. 183–235.
- [3] M. Bernardo (2015): *On the Tradeoff between Compositionality and Exactness in Weak Bisimilarity for Integrated-Time Markovian Process Calculi*. *Theoretical Computer Science*. To appear, available at www.sti.uniurb.it/bernardo/documents/tcs.pdf.
- [4] Y. Deng & M. Hennessy (2013): *On the Semantics of Markov Automata*. *Information and Computation* 222, pp. 139–168.

- [5] C. Eisentraut, H. Hermanns & L. Zhang (2010): *On Probabilistic Automata in Continuous Time*. In: *Proc. of the 25th IEEE Symp. on Logic in Computer Science (LICS 2010)*, IEEE-CS Press, pp. 342–351. Full version available at www.avacs.org/Publikationen/Open/avacs_technical_report_062.pdf.
- [6] D. Guck, H. Hatefi, H. Hermanns, J.-P. Katoen & M. Timmer (2013): *Modelling, Reduction and Analysis of Markov Automata*. In: *Proc. of the 10th Int. Conf. on the Quantitative Evaluation of Systems (QEST 2013)*, LNCS 8054, Springer, pp. 55–71.
- [7] H. Hermanns (2002): *Interactive Markov Chains*. Springer. Volume 2428 of LNCS.
- [8] R. Lanotte, A. Maggiolo-Schettini & A. Troina (2010): *Weak Bisimulation for Probabilistic Timed Automata*. *Theoretical Computer Science* 411, pp. 4291–4322.
- [9] J. Markovski, P.R. D’Argenio, J.C.M. Baeten & E.P. de Vink (2012): *Reconciling Real and Stochastic Time: The Need for Probabilistic Refinement*. *Formal Aspects of Computing* 24, pp. 497–518.
- [10] F. Moller & C. Tofts (1990): *A Temporal Calculus of Communicating Systems*. In: *Proc. of the 1st Int. Conf. on Concurrency Theory (CONCUR 1990)*, LNCS 458, Springer, pp. 401–415.
- [11] F. Moller & C. Tofts (1992): *Behavioural Abstraction in TCCS*. In: *Proc. of the 19th Int. Coll. on Automata, Languages and Programming (ICALP 1992)*, LNCS 623, Springer, pp. 559–570.
- [12] J. Schuster & M. Siegle (2014): *Markov Automata: Deciding Weak Bisimulation by Means of Non-Naïvely Vanishing States*. *Information and Computation* 237, pp. 151–173.
- [13] R. Segala (1995): *Modeling and Verification of Randomized Distributed Real-Time Systems*. PhD Thesis.
- [14] M. Timmer, J. van de Pol & M. Stoelinga (2013): *Confluence Reduction for Markov Automata*. In: *Proc. of the 11th Int. Conf. on Formal Modeling and Analysis of Timed Systems (FORMATS 2013)*, LNCS 8053, Springer, pp. 243–257.
- [15] Wang Yi (1991): *CCS + Time = An Interleaving Model for Real Time Systems*. In: *Proc. of the 18th Int. Coll. on Automata, Languages and Programming (ICALP 1991)*, LNCS 510, Springer, pp. 217–228.