

Model Checking Timed and Stochastic Properties with CSL^{TA}

Susanna Donatelli, *Member, IEEE*, Serge Haddad, and Jeremy Sproston

Abstract—Markov chains are a well-known stochastic process that provide a balance between being able to adequately model the system's behavior and being able to afford the cost of the model solution. Systems can be modeled directly as Markov chains, or with a higher-level formalism for which Markov chains represent the underlying semantics. Markov chains are widely used to study the performance of computer and telecommunication systems. The definition of stochastic temporal logics like Continuous Stochastic Logic (CSL) and its variant asCSL, and of their model-checking algorithms, allows a unified approach to the verification of systems, allowing the mix of performance evaluation and probabilistic verification. In this paper, we present the stochastic logic CSL^{TA}, which is more expressive than CSL and asCSL, and in which properties can be specified using automata (more precisely, timed automata with a single clock). The extension with respect to expressiveness allows the specification of properties referring to the probability of a finite sequence of timed events. A typical example is the responsiveness property “with probability at least 0.75, a message sent at time 0 by a system A will be received before time 5 by system B and the acknowledgment will be back at A before time 7”, a property that cannot be expressed in either CSL or asCSL. Furthermore, the choice of using automata rather than the classical temporal operators Next and Until should help in enlarging the accessibility of model checking to a larger public. We also present a model-checking algorithm for CSL^{TA}.

Index Terms—Verification, model checking, Markov processes.

1 INTRODUCTION

SOFTWARE performance engineering [1] (SPE) is a discipline that advocates an integrated approach to system design and system analysis. SPE emphasizes the importance of obtaining performance measures early in the development process, when appropriate decisions can be taken at a lower cost, usually through modeling, given that the target system is not available yet. Over the years, SPE has evolved to encompass other nonfunctional requirements such as dependability and quality of service (QoS) aspects.

Part of the work of a software performance engineer consists of defining appropriate performance and dependability properties. In this paper, we introduce a new stochastic logic that allows us to define and check nonfunctional properties over Continuous-Time Markov Chains (CTMC).

CTMCs are a well-known stochastic process that balances the need to have a model that is representative enough of the real system being modeled while still allowing affordable solution costs: There are standard and widespread solution methods for the computation of performance measures of a CTMC. The relevance of CTMCs for SPE is clearly evidenced by recent works on automatic translations from UML diagrams, annotated according to some standard or quasi-

standard profile such as UML-SPT [2] and MARTE [3], to various performance evaluation formalisms like queuing networks [4], stochastic Petri nets [5], and stochastic process algebras [6], whose underlying stochastic process is, in most cases, a CTMC. A complete overview of the tools and algorithms for the translation from UML diagrams to performance models can be found in [7].

Historically, CTMCs have been analyzed using steady state and transient analysis to compute the probability of finding the system in a given state assuming it has reached equilibrium or the probability of finding a system in a given state at time t . From these basic methods, the use of state and/or transition rewards allows the computation of performance/dependability properties.

More recently, the definition of Continuous Stochastic Logic (CSL) [8], [9] and variants, such as asCSL [10], has introduced a new approach for the definition of the performance and dependability properties of a system. Temporal-logic-based approaches are particularly useful when the measure of interest depends on the execution path. Given a formal description of the system and its requirements, we can then execute a model-checking algorithm which establishes automatically whether the system model meets the requirements expressed in CSL or asCSL.

To illustrate the advantages of stochastic logics, consider a system whose stochastic behavior is described by a CTMC, whose states (of which there can be millions) are partitioned into “system is working properly” (*work*-states), “system is working in degraded mode” (*degr*-states), or “system is not working properly” (*fail*-states). The CTMC can move from *work* to *degr* states and to *fail* states (either directly or through *degr* states). A simple example of CTMC exhibiting this behavior is shown in Fig. 1. A classical dependability property requires the computation of the probability of failing within the time interval I , or later than a given threshold t : these probabilities can be easily computed using classical solution methods for CTMCs.

- S. Donatelli and J. Sproston are with the Dipartimento di Informatica, Università degli studi di Torino, Corso Svizzera 185, 10149 Torino, Italy. E-mail: {susi, sproston}@di.unito.it
- S. Haddad is with the Laboratoire Spécification et Vérification, CNRS UMR 8643, Ecole Normale Supérieure de Cachan, 61, avenue du Président Wilson, 94235 Cachan Cedex, France. E-mail: haddad@lsv.ens-cachan.fr.

Manuscript received 27 Jan. 2008; revised 2 Sept. 2008; accepted 16 Sept. 2008; published online 10 Dec. 2008.

Recommended for acceptance by J. Hillston, M. Kwiatkowska, and M. Telek. For information on obtaining reprints of this article, please send e-mail to: tse@computer.org, and reference IEEECS Log Number TSESI-2008-01-0039. Digital Object Identifier 10.1109/TSE.2008.108.




Fig. 1. A simple CTMC.

Instead, if we are interested in only those failures in which the system fails within the time interval I , without first entering the degraded mode, we have to compute the probability of reaching a *fail* state within I , while passing only through *work* states. The stochastic temporal logic CSL has temporal operators that allow a simple and semantically-clear description of such a property using the Until operator: $\mathcal{P}_{\leq \lambda}(\text{work } \mathcal{U}^I \text{ fail})$. The property is satisfied in a state s if the set of timed paths of the CTMC that start in s and visit only *work* states before entering a *fail* state at a time in the interval I have a probability at most λ .

The logic asCSL permits the specification of paths in terms of state labels (such as *work*, *degr*, and *fail*) and action labels. For example, $\mathcal{P}_{\leq \lambda}((\text{work}, \text{Act})^*; (\text{work}, \text{failure1}); (\text{fail}, \sqrt{\quad})^I)$, is similar to the CSL formula above, with the additional restriction that the change from *work*-states to *fail*-states is due to action *failure1*.

In this paper, we propose the new stochastic temporal logic Continuous Stochastic Logic with Timed Automata (CSL^{TA}), which builds on CSL and asCSL by enriching the set of properties that can be defined and verified, and presents its associated model-checking algorithm. Let us first explain the main motivations for introducing a new logic. Modifications are along two lines: Properties are specified using deterministic one-clock timed automata and the defined logic is at least as expressive as CSL and asCSL (and strictly more expressive than CSL and asCSL without nested formulas). For the time being, we shall be informal, at the risk of being slightly imprecise, to convey the main ideas. The rest of the paper will provide the required formal development.

The idea of using automata for specifying system behavior is familiar to computer scientists in general and to software engineers in particular. The use of automata to specify temporal logic properties is not new: Vardi and Wolper [11] define a linear-time temporal logic with Büchi automata operators, while Clarke et al. [12] introduce the temporal logic ECTL, which uses Muller automata to specify linear and branching temporal properties, and develop an associated model-checking algorithm.

Timed automata [13] are a widespread formalism for the specification of timed systems, and are supported by tools such as UPPAAL [14]. Previous work has considered the use of timed automata to specify both the system and the properties that it should satisfy [13], or, more typically, the use of timed temporal logic to specify the properties of a timed automata system [15], [16]; observe that, in [16], the model-checking algorithm involves the transformation of temporal logic properties into timed automata. In this paper, we use timed automata to specify timed and probabilistic properties of the system and not to specify the system itself.

We illustrate the limitations of CSL and asCSL with respect to CSL^{TA} using again the CTMC of Fig. 1. Assume that we are interested in the probability of the system exhibiting the following behavior: The system goes from




Fig. 2. Automata for two QoS properties.

work states to *degr* states and then from *degr* states to *fail* states in a time greater than or equal to t . This property can be expressed in CSL^{TA} using the timed automaton of Fig. 2a, where x is a clock (a variable whose value increases at the same rate as time). This property can also be expressed in asCSL, using a formula similar to the one given above: $\mathcal{P}_{\leq \lambda}((\text{work}, \text{Act})^*; (\text{degr}, \text{Act})^*; (\text{fail}, \sqrt{\quad})^I)$, for the interval $I = [t, \infty)$. This property cannot be expressed in CSL, which can only express the probability of being in *work* or *degr* states until, at a time at least t , the system moves to *fail* states: This property is obviously not equivalent to the original one since it also includes paths that cycle between *work* and *degr* states.

Now, assume that the QoS requirements imposed on the system are more stringent and detailed, requiring that, with probability at least λ , the system goes from *work* states to *degr* states in no less than t' time units and then from *degr* states to *fail* states in no less than t'' time units. Paths are, therefore, characterized in terms of states and in terms of two time constraints. This property can be expressed in CSL^{TA} using the timed automaton of Fig. 2b, where the edge label $\{x\}$ indicates that the clock x is reset to zero on traversal of the edge. Observe that the two properties are different because 1) checks only the global time to get to *fail* states, while 2) also “looks” inside the composition of this duration: indeed, the automaton of Fig. 2b is not equivalent to the automaton of Fig. 2a with $t = t' + t''$.

This QoS requirement cannot be expressed either in CSL or in asCSL, as will be explained in later sections. What can indeed be expressed in CSL and asCSL is that, with probability at most λ , the system moves not before t' from *work* states to the subset of *degr* states from which, with probability at most λ' , the system moves to *fail* not before t'' . Note that we have to utilize a “fake” probability, introducing an overspecification with respect to the original QoS requirement because timed intervals cannot be nested directly in CSL and asCSL, while probabilistic operators can be nested. Note that taking $\lambda' = 1$ does not solve the problem.

In addition to introducing CSL^{TA}, we present its associated model-checking algorithm. Contrary to the previous approaches that perform *ad hoc* transformations of the CTMC before a transient or steady-state analysis, this algorithm generates a Markov regenerative process and then computes a reachability probability on this process. Furthermore, we prove that CSL^{TA} is at least as expressive as CSL and asCSL: It is possible to transform any CSL or asCSL formula into an equivalent CSL^{TA} formula. We note that the CSL^{TA} model-checking algorithm, when executed on CSL^{TA} properties transformed from CSL properties, is no more expensive in terms of computational complexity

than the CSL model-checking algorithm of [9]. Finally, we show that CSL^{TA} is strictly more expressive than CSL: Note that the proof technique used is different from those used in the nonstochastic context, for example, in [17]. We also show that CSL^{TA} is more expressive than asCSL when restricting to the case of formulas without nesting.

With regard to related work, performance metrics that depend on paths have also been studied in [18], [19]. In particular, the work in [19] uses automata for the specification of the set of paths of interest of a CTMC: Rewards, which are usually associated with states or transitions of the CTMC, are instead associated with locations and transitions of the automaton, thus providing a wide range of performance measures based on states and/or events of the CTMC. We also note that the logic CSL^{TA} is similar to the logic TECTL_{\exists}^* [20] from the nonprobabilistic model-checking literature. One-clock timed automata have been studied in, for example, [21], [22]. Finally, we recall that the original definition of CSL permitted the description of a sequence of timed Until formulas within a single probabilistic operator $\mathcal{P}_{\sim\lambda}$ [8], in contrast to the more established definition in which only one time Until formula can be included within a probabilistic operator; however, the decidability results of [8] are based on results from algebraic and transcendental number theory, whereas established performance evaluation techniques are used as the foundation of the algorithms for CSL in [9] and for CSL^{TA} in this paper.

The rest of the paper is organized as follows: Section 2 defines the syntax and semantics of CSL^{TA}, illustrated with the help of small examples. Section 3 presents the model-checking algorithm for CSL^{TA} and gives an example on a simple CTMC, while Section 4 compares the expressiveness of CSL^{TA}, CSL, and asCSL. Section 5 summarizes the paper and discusses future work. A conference version of this paper appears as [23]. We extend that version in the following ways: The semantics of CSL^{TA} is improved in order to make the modeling of properties more intuitive, Section 3 is expanded, and a more formal approach, including proofs of the main results, is taken up in Section 4.

2 SYNTAX AND SEMANTICS OF CSL^{TA}

In this section, we first introduce a number of preliminary concepts. After defining a class of labeled CTMCs, we recall the notion of execution path of a CTMC as a finite or infinite sequence of transitions from state to state. We then introduce a restricted class of timed automata, and consider the manner in which such automata can be used to express properties of CTMC execution paths. Finally, we introduce the syntax of CSL^{TA}, which is similar to CSL but which uses timed automata to express properties of paths, and present its semantics.

2.1 Labeled Markov Chains

We first introduce continuous-time Markov chains labeled both by atomic propositions on states and by actions on transitions.

Atomic propositions can refer to basic properties that are observed when the system is in a state (such as *idle* or *error*), whereas actions refer to basic properties that are observed when the system makes a transition from state to state (such as *activate* or *send_message*). Such labeled Markov chains can be used as the underlying semantic

model of high-level formalisms such as stochastic Petri nets and stochastic process algebras. Let $\mathbb{R}_{\geq 0}$ ($\mathbb{R}_{> 0}$) be the set of nonnegative (positive) reals and let \mathbb{N} be the set of natural numbers.

Definition 2.1 (Action- and State-Labeled Markov Chain).

An action and state-labeled continuous-time Markov chain (ASMC) is a tuple $\mathcal{M} = \langle S, Act, AP, lab, \mathbf{R} \rangle$, where S is a finite set of states, Act is a finite set of action labels, AP is a finite set of atomic propositions, $lab : S \rightarrow 2^{AP}$ is a state labeling function, and $\mathbf{R} : S \times Act \times S \rightarrow \mathbb{R}_{\geq 0}$ is a rate matrix. We require that, for any state s , there exists a pair $(a, s') \in Act \times S$ with $\mathbf{R}(s, a, s') > 0$.

Intuitively, the rate matrix \mathbf{R} describes the transitions that can be made between states of the ASMC, on which actions, and with which rate. A transition from state s to state s' performing action a , exists if $\mathbf{R}(s, a, s') > 0$. A transition from s to s' performing a and of duration $\tau \in \mathbb{R}_{> 0}$ is denoted by $s \xrightarrow{a, \tau} s'$.

Definition 2.2 (Paths of \mathcal{M}).

A finite path of an ASMC \mathcal{M} is a finite sequence of transitions $\sigma = s_0 \xrightarrow{a_0, \tau_0} s_1 \xrightarrow{a_1, \tau_1} \dots s_{n-1} \xrightarrow{a_{n-1}, \tau_{n-1}} s_n$ where $\mathbf{R}(s_i, a_i, s_{i+1}) > 0$ for $i = 0, \dots, n-1$. An infinite path of \mathcal{M} is an infinite sequence of transitions $\sigma = s_0 \xrightarrow{a_0, \tau_0} s_1 \xrightarrow{a_1, \tau_1} \dots$, where $\mathbf{R}(s_i, a_i, s_{i+1}) > 0$ for all $i \geq 0$ and such that $\sum_{i \geq 0} \tau_i = \infty$.

Notation. Given $s \in S$, let $Path^{\mathcal{M}}(s)$ be the set of infinite paths $s_0 \xrightarrow{a_0, \tau_0} s_1 \xrightarrow{a_1, \tau_1} \dots$ such that $s_0 = s$. Let $\text{Pr}_s^{\mathcal{M}}$ be the probability measure on $Path^{\mathcal{M}}(s)$ defined in the standard manner (for example, see [9], [10]). Let $\sigma = s_0 \xrightarrow{a_0, \tau_0} s_1 \xrightarrow{a_1, \tau_1} \dots \xrightarrow{a_{n-1}, \tau_{n-1}} s_n$ be a finite path. Then $|\sigma| = n$ denotes the length of σ and $\tau(\sigma) = \sum_{i=0}^{n-1} \tau_i$ is the total duration of σ . By convention, $\tau_{-1} = 0$. For an infinite path σ , we let $|\sigma| = \infty$ and $\tau(\sigma) = \infty$.

As usual, we can describe an ASMC by a graph. An example of an ASMC is given in Fig. 3. The vertices of this graph are its states whereas the edges represent its transitions. The atomic propositions (here p and q) that are satisfied in a state are indicated near the corresponding node. Finally, the rate of a transition labels the corresponding edge.

2.2 Timed Automata

We now present a restricted variant of timed automata [13], which are used in CSL^{TA} to describe properties of ASMC paths. More precisely, in our context, timed automata are used as acceptors of finite ASMC paths. The class of timed automata that we consider are deterministic (i.e., given a path σ of an ASMC, there is at most one path of the timed automaton which reads σ), and have a single clock. In the same manner as in classical analysis techniques for timed automata [13], we present our timed automata using natural-numbered constants (rational-numbered constants can also be considered through rescaling). We proceed to define deterministic (one-clock) timed automata. We use the symbol \sharp to denote a pseudoaction that is not included in the action set Act of any ASMC ($\sharp \notin Act$). Clock variables are real-valued variables whose value increases linearly with time. We consider a single clock variable x . A valuation $\bar{x} \in \mathbb{R}_{\geq 0}$ is interpreted as assigning a nonnegative real value to x . A constraint is of the form $\alpha \prec x \prec \beta$ or $\alpha \prec x$ where $\alpha, \beta \in \mathbb{N}$, $\alpha \leq \beta$, and \prec stands for either $<$ or \leq . An inner




Fig. 3. An ASMC.

constraint is a constraint $\alpha < x < \beta$ such that $\alpha < \beta$. The set of inner constraints is denoted Inner. A *boundary constraint* is a constraint $\alpha \leq x \leq \beta$ such that $\alpha = \beta$; we generally write boundary constraints as $x = \alpha$. The set of boundary constraints is denoted Boundary. Let γ be a constraint and \bar{x} be a clock valuation. Then, we write $\bar{x} \models \gamma$ if γ is satisfied when \bar{x} is substituted for x in γ .



Locations of a timed automaton are labeled with *state propositions*. A state proposition is a proposition which either holds, or does not hold, in an ASMC state. For a set Σ of state propositions, let \models_{Σ} be its associated satisfaction relation: hence, we write $\mathcal{M}, s \models_{\Sigma} \Phi$ to denote that the state s of the ASMC \mathcal{M} satisfies Φ . We omit \mathcal{M} and write $s \models_{\Sigma} \Phi$ when clear from the context. We also consider Boolean expressions of state propositions: For example, $s \models_{\Sigma} \Phi_1 \wedge \Phi_2$ denotes that s satisfies Φ_1 and Φ_2 . Let $\mathcal{B}(\Sigma)$ be the set of Boolean expressions over state propositions of Σ . Later in the paper we will make the set of state propositions Σ used in CSL^{TA} precise. For the purposes of the current explanation, the reader can consider the case in which $\Sigma = AP$ with $s \models_{AP} p$ if and only if $p \in \text{lab}(s)$, for a state s and $p \in AP$.

Definition 2.3 (Deterministic Timed Automaton). A deterministic timed automaton (DTA) $\mathcal{A} = \langle \Sigma, \text{Act}, L, \Lambda, \text{Init}, \text{Final}, \rightarrow \rangle$ comprises:

- Σ , a finite alphabet of state propositions;
- Act , a finite alphabet of actions;
- L , a finite set of locations;
- $\Lambda : L \rightarrow \mathcal{B}(\Sigma)$, a location labeling function;
- Init , a subset of L called the initial locations;
- Final , a subset of L called the final locations;
- $\rightarrow \subseteq L \times ((\text{Inner} \times 2^{\text{Act}}) \cup (\text{Boundary} \times \{\#\})) \times \{\emptyset, x\} \times L$, a set of edges, where $l \xrightarrow{\gamma, A, r} l'$ denotes that $(l, \gamma, A, r, l') \in \rightarrow$.

Furthermore, \mathcal{A} fulfils the following conditions.

- *Initial determinism:* $\forall l, l' \in \text{Init}, \Lambda(l) \wedge \Lambda(l') \Leftrightarrow \text{false}$.
- *Determinism on actions:* $\forall A, A' \subseteq \text{Act}$ s.t. $A \cap A' \neq \emptyset, \forall l, l', l'' \in L$, if $l'' \xrightarrow{\gamma, A, r} l \wedge l'' \xrightarrow{\gamma, A', r'} l'$, then either $\Lambda(l) \wedge \Lambda(l') \Leftrightarrow \text{false}$ or $\gamma \wedge \gamma' \Leftrightarrow \text{false}$.
- *Determinism on #:* $\forall l, l', l'' \in L$, if $l'' \xrightarrow{\gamma, \#, r} l \wedge l'' \xrightarrow{\gamma', \#, r'} l'$, then either $\Lambda(l) \wedge \Lambda(l') \Leftrightarrow \text{false}$ or $\gamma \wedge \gamma' \Leftrightarrow \text{false}$.

Fig. 4. The DTA $\mathcal{A}_{X^{[\alpha, \beta]} \Phi_1}$.Fig. 5. The DTA $\mathcal{A}_{\Phi_1, U^{[\alpha, \beta]} \Phi_2}$.

- *No #-labeled loops:* For all sequences $l_0 \xrightarrow{\gamma_0, A_0, r_0} l_1 \xrightarrow{\gamma_1, A_1, r_1} \dots \xrightarrow{\gamma_{n-1}, A_{n-1}, r_{n-1}} l_n$ such that $l_0 = l_n$, there exists $i \leq n$ such that $A_i \neq \#$.

Let Φ_1 and Φ_2 be state propositions in the alphabet Σ . Fig. 4 depicts a DTA, using the usual conventions for the graphical representation of timed automata (i.e., nodes represent locations, and edges represent edges labeled with their guards, actions sets, and the set of clocks to be reset to 0, respectively). Initial locations are denoted by an incoming arrow with no source, and final locations by a double border. Edges labeled by $\#$ are called *boundary edges*, while the other edges are called *inner edges*. For the DTA of Fig. 4 the determinism is obvious because there is no choice allowed. Fig. 5 shows a more complex DTA, with inner edges (the self-loops on l_0, l_1 , and the arc from l_0 to l_1) and boundary edges (the arcs from l_0 to l_1 and from l_0 to l_2). The DTA respects the determinism constraints of the definition because the two boundary edges out of location l_0 lead to two locations whose labeling cannot be both satisfied by any state of an ASMC (indeed, $\Lambda(l_1) \wedge \Lambda(l_2) = (\Phi_1 \wedge \neg \Phi_2) \wedge \Phi_2 = \text{false}$).

The semantics of DTA, expressed in terms of paths, is standard [13], apart from the case of boundary edges, which are *urgent* and have *priority* over other edges. Urgency specifies that time cannot elapse if a boundary edge is enabled and it is a feature of the variants of timed automata used in the tools UPPAAL [14] and KRONOS [24]. In our context, the notions of urgency and priority are not relevant when considering a DTA in isolation. They will be introduced later when we define the notion of a path of a DTA that reads an ASMC path, and the notion of path acceptance (Definition 2.7).

Examples of DTA: Next and Until. The DTA $\mathcal{A}_{X^{[\alpha, \beta]} \Phi_1}$ in Fig. 4 specifies behaviors in which the first transition of \mathcal{M} must be taken to a state satisfying Φ_1 after at least α time units, but not after β time units, and corresponds to the Next path formula $X^{[\alpha, \beta]} \Phi_1$ of CSL [9]. The action of the transition is not important; this fact is represented by the action set Act on the edge of the DTA.

We can use the DTA $\mathcal{A}_{\Phi_1 \mathcal{U}^{[\alpha, \beta]} \Phi_2}$ of Fig. 5 to represent the property of eventually reaching a state satisfying Φ_2 at some instant between α and β time units, remaining within states satisfying Φ_1 before that point (the timed Until path property $\Phi_1 \mathcal{U}^{[\alpha, \beta]} \Phi_2$ of CSL [9]). In contrast to the previous example, this DTA uses boundary edges which witness that the time interval $[\alpha, \beta]$ has been entered. In this way, we distinguish between the time interval $[0, \alpha]$, where the truth value of Φ_2 is irrelevant, and the time interval $[\alpha, \beta]$, where the truth value of Φ_2 becomes relevant.

Paths of a DTA. We now define a notion of path in a DTA, which represents a timed evolution of the automaton.

Definition 2.4 (Configurations of \mathcal{A}). A configuration of a DTA \mathcal{A} is a pair (l, \bar{x}) , where $l \in L$ and \bar{x} is a valuation.

Given an edge $e = (l, \gamma, A, r, l') \in \rightarrow$, let $\text{source}(e) = l$, $\text{guard}(e) = \gamma$, $\text{action}(e) = A$, $\text{reset}(e) = r$, and $\text{target}(e) = l'$. We let the valuation $\bar{x}[x := 0]$ be equal to 0 and let the valuation $\bar{x}[\emptyset := 0]$ be equal to \bar{x} .

Definition 2.5 (Step of \mathcal{A}). A step of a DTA \mathcal{A} from a configuration (l, \bar{x}) is $(l, \bar{x}) \xrightarrow{\delta, e} (l', \bar{x}')$ of \mathcal{A} with $\delta \geq 0$, $\text{source}(e) = l$, $\bar{x} + \delta \models \text{guard}(e)$, $\text{target}(e) = l'$, and $\bar{x}' = (\bar{x} + \delta)[\text{reset}(e) := 0]$.

A single step in the evolution of \mathcal{A} is a transition in which we let time elapse and then an inner or a boundary edge is taken. Note that $\delta = 0$ is also allowed.

Definition 2.6 (Paths of \mathcal{A}). A finite path of a DTA \mathcal{A} is a finite sequence of steps $(l_0, \bar{x}_0) \xrightarrow{\delta_0, e_0} (l_1, \bar{x}_1) \xrightarrow{\delta_1, e_1} \dots (l_{n-1}, \bar{x}_{n-1}) \xrightarrow{\delta_{n-1}, e_{n-1}} (l_n, \bar{x}_n)$.

An infinite path of a DTA \mathcal{A} is an infinite sequence of steps $(l_0, \bar{x}_0) \xrightarrow{\delta_0, e_0} (l_1, \bar{x}_1) \xrightarrow{\delta_1, e_1} \dots$

2.3 Acceptance of ASMC Paths

We now give an intuitive explanation of how a path $\sigma = s_0 \xrightarrow{a_0, \tau_0} s_1 \xrightarrow{a_1, \tau_1} \dots$ of an ASMC \mathcal{M} can be accepted by a DTA \mathcal{A} . The key idea is that \mathcal{A} evolves according to the states and actions that it “reads” along σ . Recalling that the value of clocks in timed automata increase at the same rate as real-time, as time elapses in \mathcal{M} the value of the clock x of \mathcal{A} changes accordingly. Steps corresponding to inner edges of \mathcal{A} are triggered by transitions of \mathcal{M} , whereas steps corresponding to boundary edges of \mathcal{A} are triggered by the elapse of time (without a corresponding transition of \mathcal{M}).

The DTA \mathcal{A} begins in a configuration $(l_0, 0)$ with location $l_0 \in \text{Init}$ such that the initial state s_0 of σ satisfies the expression $\Lambda(l_0)$ over state propositions (formally, $s_0 \models_{\Sigma} \Lambda(l_0)$). Note that, by initial determinism, there is at most one $l \in \text{Init}$ such that s_0 satisfies $\Lambda(l)$. If s_0 does not satisfy $\Lambda(l)$ for all $l \in \text{Init}$, then \mathcal{A} rejects σ .

Given the existence of $l_0 \in \text{Init}$ such that s_0 satisfies $\Lambda(l_0)$, the DTA \mathcal{A} then moves from $(l_0, 0)$ to another configuration depending on the first transition $s_0 \xrightarrow{a_0, \tau_0} s_1$ of σ . First, we consider the case in which there are no outgoing boundary edges from l_0 . If there exists a step $(l_0, 0) \xrightarrow{\tau_0, e_0} (l_1, \bar{x}_1)$ such that $a_0 \in \text{action}(e_0)$ and s_1 satisfies $\Lambda(l_1)$, then this step is taken. Note that, by determinism on actions, there exists at most one step satisfying these conditions. If no such step exists, then \mathcal{A} rejects σ .

Now we consider the case in which there exists at least one boundary edge from l_0 . Consider the step $(l_0, 0) \xrightarrow{\delta'_0, e'_0} (l'_1, \bar{x}'_1)$, where $\text{action}(e'_0) = \sharp$, which (by urgency of boundary edges) corresponds to the earliest boundary edge available by letting time elapse from $(l_0, 0)$. If $\delta'_0 > \tau_0$, then this step is available only after \mathcal{M} has performed the transition $s_0 \xrightarrow{a_0, \tau_0} s_1$; hence, the DTA “reads” the ASMC transition before the boundary edge is available, and this case is similar to the case in which there are no boundary edges from l_0 in the previous paragraph. If, however, $\delta'_0 \leq \tau_0$, the DTA takes the step before “reading” the ASMC transition. Note that, in this case, the remaining time before the transition of \mathcal{M} must be “read” by \mathcal{A} is $\tau_0 - \delta'_0$, rather than τ . This has implications for deciding whether a boundary edge can be taken from (l'_1, \bar{x}'_1) or whether the transition of \mathcal{M} must be “read” before any boundary edge is enabled for choice.

Unless \mathcal{A} has already rejected σ , the path of \mathcal{A} generated by σ then continues from (l_1, \bar{x}_1) or (l'_1, \bar{x}'_1) . Finally, if the path of \mathcal{A} generated by σ reaches a configuration with a location in Final , then σ is accepted. If, however, the path of \mathcal{A} generated by σ does not reach such a configuration, then σ is rejected. Hence, there are two ways in which \mathcal{A} can reject σ : if there does not exist a step corresponding to the “reading” of a transition of σ or if a final location is never reached.

We now formally describe the conditions for the acceptance of an ASMC path by a DTA.

Definition 2.7. Let \mathcal{M} be an ASMC, and let \mathcal{A} be a DTA. The infinite path $\sigma_{\mathcal{M}} = s_0 \xrightarrow{a_0, \tau_0} s_1 \xrightarrow{a_1, \tau_1} \dots$ of \mathcal{M} is accepted by \mathcal{A} if there exists

- a finite path $\sigma_{\mathcal{A}} = (l_0, \bar{x}_0) \xrightarrow{\delta_0, e_0} (l_1, \bar{x}_1) \xrightarrow{\delta_1, e_1} \dots \xrightarrow{\delta_{m-1}, e_{m-1}} (l_m, \bar{x}_m)$ of \mathcal{A} ,
- an index $n \in \mathbb{N}$,
- a time $\tau \leq \tau_n$, and
- a function $\kappa : \{0, \dots, m\} \rightarrow \{0, \dots, n\}$ which maps indices of $\sigma_{\mathcal{A}}$ to indices of $\sigma_{\mathcal{M}}$,

such that the following conditions are satisfied:

- C1. $l_0 \in \text{Init}$, $\bar{x}_0 = 0$, $\kappa(0) = 0$, and $\forall 0 \leq i \leq m$, $l_i \in \text{Final} \Leftrightarrow i = m$;
- C2. $\forall 0 \leq i \leq m$, $s_{\kappa(i)} \models_{\Sigma} \Lambda(l_i)$;
- C3. $\forall 0 \leq i < m$, if e_i is an inner edge, then $\kappa(i+1) = \kappa(i) + 1 \wedge a_{\kappa(i)} \in \text{action}(e_i)$ else $\kappa(i+1) = \kappa(i)$;
- C4. $\forall 0 \leq i < m$, if e_i is an inner edge then Boundary edges are urgent: for all $0 \leq \delta' < \delta_i$, there does not exist an edge $e' = (l_i, \gamma, \sharp, r, l') \in \rightarrow$ such that $\bar{x} + \delta' \models \gamma$ and $s_{\kappa(i)} \models_{\Sigma} \Lambda(l')$, and Boundary edges have priority: for all edges $e' \in \rightarrow$ such that $e \neq e'$, if $\text{source}(e') = l$, $\bar{x} + \delta_i \models \text{guard}(e')$, $\text{target}(e') = l'$, and $s_{\kappa(i)} \models_{\Sigma} \Lambda(l')$, then $\text{action}(e') \neq \sharp$;
- C5. $\forall 0 \leq i < n$, $\sum_{j|\kappa(j)=i} \delta_j = \tau_{\kappa(i)}$;
- C6. $\sum_{j|\kappa(j)=n} \delta_j = \tau$.

Condition C1 specifies that $\sigma_{\mathcal{A}}$ must start from an initial location and end in a final location. C2 requires that the state propositions satisfy the expressions labeling the corresponding locations in the sequence. C3 specifies that κ can map




Fig. 6. Examples of path acceptance.

consecutive indices of σ_A to the same index of σ_M , provided that the DTA edges corresponding to these indices are boundary edges. It also requires that a transition of the ASMC in σ can be matched by a traversal of an inner edge provided that the action of the transition is included in the action set of the edge. C4 limits the path of the DTA to paths whose steps respect the additional conditions on \sharp : urgency and priority of boundary edges. C5 “align times,” by requiring that the sum of durations in σ_A corresponding to a particular index i of σ_M is τ_i . C6 applies the reasoning of C5 to the case in which the path σ_A features boundary edges directly before reaching a final state (recall that $\tau \leq \tau_n$).

It should be clear that, given an ASMC \mathcal{M} and a DTA \mathcal{A} , due to our requirements for DTA and to the additional requirements of urgency and priority of boundary edges, there is at most one path of \mathcal{A} that accepts a given path of \mathcal{M} . Accordingly, if s is a state of \mathcal{M} , we let $AccPath^{\mathcal{M}}(s, \mathcal{A})$ be the set of infinite paths of \mathcal{M} starting from s that are accepted by \mathcal{A} .

Examples of path acceptance. In Fig. 6, we present two examples of the way in which a path of the ASMC \mathcal{M} of Fig. 3 can be accepted by the DTA $\mathcal{A}_{\rho\lambda^{[\alpha,\beta]}_q}$. We write e_{ij} to refer to the edge of $\mathcal{A}_{\rho\lambda^{[\alpha,\beta]}_q}$ from location l_i to location l_j , and we use dotted lines to represent the κ function. Note that the presence of more than one dotted line from a state s_i means that the DTA traverses a boundary edge. Example 1 of Fig. 6 has $\alpha = 2$ and $\beta = 6$, and depicts a case in which q does not hold at time α , but becomes true at time 5, which belongs to $[\alpha, \beta]$; therefore the DTA reaches l_2 through l_1 . Example 2 of Fig. 6 has $\alpha = 6$ and $\beta > 6$ and depicts a case in which q already holds before α ; therefore, the DTA reaches l_2 directly from l_0 .

We now describe briefly some examples of paths of \mathcal{M} which are rejected by $\mathcal{A}_{\rho\lambda^{[\alpha,\beta]}_q}$. If the first transition of \mathcal{M} is $s_0 \xrightarrow{a,7} s_1$, then the associated path of $\mathcal{A}_{\rho\lambda^{[\alpha,\beta]}_q}$ will consist of the single step $(l_0, 0) \xrightarrow{e_{01,2}} (l_1, 2)$: after the value of the clock x exceeds 6, it will not be possible to take further steps. If, on the other hand, the path of \mathcal{M} is $s_0 \xrightarrow{a,1} s_1 \xrightarrow{c,0.5} s_2$, then the associated path of $\mathcal{A}_{\rho\lambda^{[\alpha,\beta]}_q}$ will consist of the single step $(l_0, 0) \xrightarrow{e_{00,1}} (l_0, 1)$, after which it will not be possible to take any further steps: the state s_2 is not labeled by p , boundary edges are available only at time 2, and yet the transition $s_1 \xrightarrow{c,0.5} s_2$ occurs before time 2.

All the timed automata configurations considered in Fig. 6 are configurations in which the DTA spends a nonzero amount of time: this is not always the case if

boundary edges are involved. Indeed, a path obtained from the first example by splitting the step $(l_0, 0) \xrightarrow{e_{01,2}} (l_1, 2)$ into the two steps $(l_0, 0) \xrightarrow{e_{00,2}} (l_0, 2)$ and $(l_0, 2) \xrightarrow{e_{01,0}} (l_1, 2)$ is also an accepting path for the ASMC execution of the example. Another source of zero-time configurations in an accepting path is the presence in the path of more than one edge with the same clock constraints and no reset of clock in between.

2.4 CSL^{TA}

Given the definition of DTA, we can now present formally the syntax of CSL^{TA}. Note that the syntax of CSL^{TA} is essentially identical to that of CSL or asCSL [8], [9], [10], apart from the fact that properties of paths are specified using DTA (instead of being specified by timed temporal logic operators as, for example, in CSL).

Definition 2.8 (Syntax of CSL^{TA}). Let $\lambda \in [0, 1]$ be a real number and let $\sim \in \{<, \leq, \geq, >\}$ be a comparison operator. The syntax of CSL^{TA} is defined by:

$$\Phi ::= p \mid \neg\Phi \mid \Phi \wedge \Phi \mid \mathcal{S}_{\sim\lambda}(\Phi) \mid \mathcal{P}_{\sim\lambda}(\mathcal{A}(\Phi_1, \dots, \Phi_n))$$

where $p \in AP$ and $\mathcal{A}(\Phi_1, \dots, \Phi_n)$ is a DTA with a finite alphabet Σ of state propositions such that $\Sigma = \{\Phi_1, \dots, \Phi_n\}$ and Φ_1, \dots, Φ_n are CSL^{TA} formulas.

Note that CSL^{TA} is a CTL*-like language with nesting of path and state formulas [26]; in particular, the state propositions of a DTA are state formulas of CSL^{TA}. For example, we can write a CSL^{TA} formula such as $\mathcal{P}_{\geq 0.99}(\mathcal{A}_{\rho\lambda^{[\alpha,\beta]}_{p \geq 0.1}(\mathcal{A}_{\chi^{[\alpha,\beta]}_q})})$ (which corresponds to the CSL formula $\mathcal{P}_{\geq 0.99}(\rho\lambda^{[\alpha,\beta]} \mathcal{P}_{\geq 0.1}(\chi^{[\alpha,\beta]}_q))$).

Intuitively, state s satisfies the formula $\mathcal{S}_{\sim\lambda}(\Phi)$ if and only if $val \sim \lambda$, where val is the steady-state probability, computed assuming the ASMC starts in s , of being in an ASMC state that satisfies Φ . State s satisfies instead the formula $\mathcal{P}_{\sim\lambda}(\mathcal{A})$ if and only if $val \sim \lambda$, where val is the probability of all ASMC paths starting in s and accepted by \mathcal{A} .

We proceed to define the semantics of CSL^{TA} in terms of the satisfaction relation \models . For a given CSL^{TA} formula Φ and state s of \mathcal{M} , we write $\mathcal{M}, s \models \Phi$ to denote that Φ is satisfied in state s . We write $\pi(s, \cdot)$ for the steady-state distribution of \mathcal{M} , computed starting from state s .

Definition 2.9 (Semantics of CSL^{TA}). For $\mathcal{M} = \langle S, Act, AP, lab, \mathbf{R} \rangle$, and state $s \in S$, the satisfaction relation \models for CSL^{TA} is defined as follows:

$$\begin{aligned} \mathcal{M}, s \models p &\Leftrightarrow p \in lab(s), \\ \mathcal{M}, s \models \neg\Phi &\Leftrightarrow \mathcal{M}, s \not\models \Phi, \\ \mathcal{M}, s \models \Phi_1 \wedge \Phi_2 &\Leftrightarrow \mathcal{M}, s \models \Phi_1 \text{ and } \mathcal{M}, s \models \Phi_2, \\ \mathcal{M}, s \models \mathcal{S}_{\sim\lambda}(\Phi) &\Leftrightarrow \sum_{s' \in S \text{ s.t. } \mathcal{M}, s' \models \Phi} \pi(s, s') \sim \lambda, \\ \mathcal{M}, s \models \mathcal{P}_{\sim\lambda}(\mathcal{A}(\Phi_1, \dots, \Phi_n)) &\Leftrightarrow \Pr_s^{\mathcal{M}}(AccPath^{\mathcal{M}}(s, \mathcal{A}(\Phi_1, \dots, \Phi_n))) \sim \lambda. \end{aligned}$$

3 MODEL CHECKING FOR CSL^{TA}

As usual with CTL*^{*}-like languages [25], in order to evaluate the satisfaction of a formula Φ over an ASMC \mathcal{M} , we proceed by a bottom-up evaluation of the subformulas occurring in Φ over all the states of \mathcal{M} , labeling accordingly the states with the subformulas that they satisfy. Let Φ' be such a subformula.

- If Φ' is either an atomic proposition p , $\neg\Psi$, or $\Psi \wedge \Psi'$, then the evaluation is performed by a straightforward application of Definition 2.9.
- If $\Phi' = \mathcal{S}_{\sim\lambda}(\Psi)$, then the steady state of \mathcal{M} with respect to every state s is computed first. Afterward, the steady-state probability of the subset of states that fulfil Ψ is computed and the value compared with λ in order to check whether s satisfies Φ' .
- Finally, if $\Phi' = \mathcal{P}_{\sim\lambda}(\mathcal{A})$, for each state s , we compute the probability of $\text{AccPath}^{\mathcal{M}}(s, \mathcal{A})$ (the set of paths of \mathcal{M} accepted by \mathcal{A}), and we compare it to λ in order to check whether s satisfies Φ' . Observe that the state properties of \mathcal{A} label the states of the ASMC due to the bottom-up evaluation. The computation of $\text{Pr}_s^{\mathcal{M}}(\text{AccPath}^{\mathcal{M}}(s, \mathcal{A}))$ is the topic of the remainder of this section.

We use s_0 to denote the state for which we compute $\text{Pr}_{s_0}^{\mathcal{M}}(\text{AccPath}^{\mathcal{M}}(s_0, \mathcal{A}))$ and let $l_0 \in \text{Init}$ be the location of \mathcal{A} for which $s_0 \models_{\Sigma} \Lambda(l_0)$ (if no such location exists, then $\text{Pr}_{s_0}^{\mathcal{M}}(\text{AccPath}^{\mathcal{M}}(s_0, \mathcal{A})) = 0$).

3.1 The “Synchronized” Stochastic Process $\mathcal{M} \times \mathcal{A}$

The computation of $\text{Pr}_{s_0}^{\mathcal{M}}(\text{AccPath}^{\mathcal{M}}(s_0, \mathcal{A}))$ requires the definition of a stochastic process $\mathcal{M} \times \mathcal{A}$ that describes the joint evolution of \mathcal{M} and \mathcal{A} . The stochastic process has been enriched with two absorbing states: \top and \perp . At some instant of its execution, this process may be in one of three situations.

1. At some previous instant, \mathcal{A} has not been able to mimic the execution of \mathcal{M} and, thus, this process is in the absorbing state \perp whatever the subsequent transitions of \mathcal{M} .
2. At some previous instant, \mathcal{A} has reached a final location by following the execution of \mathcal{M} and, thus, this process is in the absorbing state \top whatever the subsequent transitions of \mathcal{M} .
3. Otherwise, the process is in some state of \mathcal{M} associated with a finite timed execution of \mathcal{A} not ending in a final location.

States of $\mathcal{M} \times \mathcal{A}$. If at some instant the execution of \mathcal{M} is neither rejected nor accepted, we observe that, for the future behavior of the process $\mathcal{M} \times \mathcal{A}$, only the current location of the path in the DTA and the value of clock x are relevant. This yields the following state description: $N(t) = (s(t), l(t), \bar{x}(t))$, where $s(t)$ is the state of \mathcal{M} at time $t \in \mathbb{R}_{\geq 0}$, $l(t)$ is the location of \mathcal{A} at time t , and $\bar{x}(t)$ is the value of the clock at time t . However, in $\mathcal{M} \times \mathcal{A}$, we consider only *tangible* states, i.e., states which do not trigger a boundary edge in zero time. Therefore, we introduce the following definition (which allows us to skip nontangible states), which is sound

because, by definition of DTA, there are no loops of \sharp transitions in \mathcal{A} .

Definition 3.1. Let $(s, l, \bar{x}) \in S \times L \times \mathbb{R}_{\geq 0}$. Then, *closure*(s, l, \bar{x}) is defined as follows:

- if $l \in \text{Final}$, then $\text{closure}(s, l, \bar{x}) = \top$;
- if $l \notin \text{Final}$ and there is a boundary edge $l \xrightarrow{\gamma, \sharp, r} l'$ with $\bar{x} \models \gamma$ and $s \models_{\Sigma} \Lambda(l')$, then $\text{closure}(s, l, \bar{x}) = \text{closure}(s, l', \bar{x}[r := 0])$;
- otherwise, $\text{closure}(s, l, \bar{x}) = (s, l, \bar{x})$.

The set of states of the process $\mathcal{M} \times \mathcal{A}$ is a subset of $\{\perp, \top\} \cup \{(s, l, \bar{x}) \mid \text{closure}(s, l, \bar{x}) = (s, l, \bar{x}), s \models_{\Sigma} \Lambda(l)\}$.

Behavior of $\mathcal{M} \times \mathcal{A}$. Let $C = \{c_0, \dots, c_m\}$ be the set of constants used in the clock constraints of \mathcal{A} enlarged with 0, ordered as follows: $0 = c_0 < c_1 < \dots < c_m$. We define $\text{next}(c_i) = c_{i+1}$ for all $i < m$ and $\text{next}(c_m) = \infty$.

Let (s, l, \bar{x}) be a state such that $\bar{x} \in [c_i, \text{next}(c_i))$ for some $i \leq m$. Process $\mathcal{M} \times \mathcal{A}$ can evolve from (s, l, \bar{x}) due to two reasons: 1) The ASMC \mathcal{M} changes its state or 2) time elapses and the value of x reaches $\text{next}(c_i)$.

In case 1, a transition $s \xrightarrow{a, \tau} s'$ is taking place in the ASMC \mathcal{M} before the next timing constant $\text{next}(c_i)$ is reached, i.e., $\bar{x} + \tau < \text{next}(c_i)$, for some τ . If this transition cannot be read by \mathcal{A} from its state (l, \bar{x}) , then the stochastic process makes a transition to \perp . If it can be read through an edge (l, γ, A, r, l') of \mathcal{A} (and, by definition, there is exactly either one or no such edge), then the process $\mathcal{M} \times \mathcal{A}$ moves to $\text{closure}(s', l', (v + \tau)[r := 0])$. Note that *closure* is needed since boundary transitions may be then triggered in zero time and/or \mathcal{A} may reach a final state so that the process enters the \top state.

Case 2 represents instead the situation in which the next clock barrier is reached by x before an ASMC transition takes place (an amount of time equal to $\text{next}(c_i) - \bar{x}$ has elapsed). Then, the process $\mathcal{M} \times \mathcal{A}$ evolves from (s, l, \bar{x}) to $\text{closure}(s, l, \text{next}(c_i))$.

It is straightforward to show that each path of $\mathcal{M} \times \mathcal{A}$ leading to \top corresponds to a (single) path in \mathcal{M} accepted by \mathcal{A} and vice versa. Furthermore, $\text{Pr}_{s_0}^{\mathcal{M}}(\text{AccPath}^{\mathcal{M}}(s_0, \mathcal{A}))$ can be computed as the probability of reaching \top in process $\mathcal{M} \times \mathcal{A}$ from $(s_0, l_0, 0)$. In the remainder of this section, we explain how the latter probability can be computed.

$\mathcal{M} \times \mathcal{A}$ is a Markov Renewal Process. We can rewrite a state of $\mathcal{M} \times \mathcal{A}$ (different from \perp, \top) in terms of the last clock constant reached as follows: $N(t) = (s(t), l(t), c(t), \bar{x}(t) - c(t))$, where $c(t)$ is the largest $c \in C$ such that $c \leq \bar{x}(t)$.

We now show that $\mathcal{M} \times \mathcal{A}$ is a Markov renewal process (MRP). For the definition of MRP and Markov renewal sequences, see, for example, [26]. Consider a sequence $\{T_k, k = 0, 1, 2, \dots\}$ of strictly increasing timing instants in the evolution of $\mathcal{M} \times \mathcal{A}$, with $N(T_k) = (s_k, l_k, c(T_k), \bar{x}_k - c(T_k))$. The timing instants are defined as follows:

1. $T_0 = 0$.
2. If $\bar{x}_k < c_m$, then T_{k+1} is the next time at which the next constant in C is reached, the clock x is reset to 0, or the process reaches $\{\top, \perp\}$.
3. If $\bar{x}_k \geq c_m$, then T_{k+1} is the first time after T_k that the clock x is reset to 0 or the process reaches $\{\top, \perp\}$.

When, for some T_k , $N(T_k) \in \{\perp, \top\}$, the sequence T_k is finite. Similarly, when $\bar{x}_k \geq c_m$, it could happen that the probability to reach a regeneration point is strictly less than 1. These particular cases do not raise any problem with respect to our computation (see the discussion at the end of this section).

Let $Y_k = N(T_k^+)$ be the state directly after all of the events at time T_k .

Theorem 3.2. $(Y, T) = \{(Y_k, T_k), k = 0, 1, 2, \dots\}$ is a Markov renewal sequence and $N(t)$ is an MRP.

The proof of Theorem 3.2 is straightforward given the definition of MRP (see [26]), because, due to the definition of T_k , we have that $Y_k = (s(T_k^+), l(T_k^+), c(T_k))$, where $c(T_k) \in C$ is the value of the clock at T_k and the joint distribution of Y_{k+1} and $T_{k+1} - T_k$ only depends on Y_k . Therefore, (Y, T) is a Markov renewal sequence. Moreover, $N(t) = (s(t), l(t), c(t), \bar{x}(t) - c(t))$, which is equal to $(s_k, l_k, c(T_k), \delta)$ for some k , and $0 \leq \delta \leq T_{k+1} - T_k$ is an MRP because $N(t)$, which is equal to $N(T_k + \delta)$, only depends on Y_k .

It is well known that $Y = \{Y_k, k = 0, 1, 2, \dots\}$ is a Discrete-Time Markov Chain (DTMC), namely, the embedded DTMC of the MRP. In general, the solution of an MRP requires the definition of the global and local kernel matrices (see [26]). The computation of the probability of reaching the absorbing state \top from the initial state can be performed on the DTMC $P_{i,j}$, which expresses the probability that, if i is the state at regeneration instant 0, then j is the state at the next regeneration instant T_1 (that is, $P_{i,j} = Pr\{Y_1 = j \mid Y_0 = i\}$).

3.2 Tangible Reachability Graph of $\mathcal{M} \times \mathcal{A}$

We next define a data structure that supports the definition of the DTMC Y and the computation of its transition probabilities. This data structure is called Tangible Reachability Graph (TRG) and is inspired by the identically named graph of Deterministic Stochastic Petri Nets [27], in which the elapsing of time between two consecutive timing constants c and $\text{next}(c)$ is interpreted as a deterministic “transition” of duration $\text{next}(c) - c$. Note that, in our case, a deterministic “transition” can only be preempted by a transition of $\mathcal{M} \times \mathcal{A}$ that includes a clock reset.

The nodes of the TRG take the form of elements of $(S \times L \times C) \cup \{\perp, \top\}$. For a constant $c \in C$ and a constraint γ , we write $(c, \text{next}(c)) \models \gamma$ if, for all $\bar{x} \in (c, \text{next}(c))$, we have $\bar{x} \models \gamma$. The arcs between nodes of the TRG are defined by the following four rules:

[M]: a simple Markovian move, in which the ASMC \mathcal{M} moves “according to” the DTA \mathcal{A} and there is no clock reset. Formally, there exists the arc $(s, l, c) \xrightarrow{M(a,e)} (s', l', c)$ if 1) $\mathbf{R}(s, a, s') > 0$; 2) $e = (l, \gamma, A, \emptyset, l')$ is an inner edge of \mathcal{A} such that $(c, \text{next}(c)) \models \gamma$, $a \in A$, and $s' \models_{\Sigma} \Lambda(l')$; and 3) $l' \notin \text{Final}$. Furthermore, there exists the arc $(s, l, c) \xrightarrow{M(a,e)} \top$ if the conditions 1 and 2 above are satisfied and $l' \in \text{Final}$.

[M_res]: as for a simple Markovian move, but with a clock reset that can start an evolution of \mathcal{A} over boundary transitions. Formally, there exists the arc $(s, l, c) \xrightarrow{M_res(a,e)}$

$\text{closure}(s', l', 0)$ if 1) $\mathbf{R}(s, a, s') > 0$ and 2) $e = (l, \gamma, A, x, l')$ is an inner edge of \mathcal{A} such that $(c, \text{next}(c)) \models \gamma$, $a \in A$, and $s' \models_{\Sigma} \Lambda(l')$.

[M_KO]: a Markovian move that is not accepted by \mathcal{A} . Formally, there exists the arc $(s, l, c) \xrightarrow{M_KO(a)} \perp$ if there exists $s' \in S$ such that $\mathbf{R}(s, a, s') > 0$ and there does not exist an inner edge $e = (l, \gamma, A, r, l')$ of \mathcal{A} such that $(c, \text{next}(c)) \models \gamma$, $a \in A$, and $s' \models_{\Sigma} \Lambda(l')$.

[D]: let time elapse. Formally, there exists the arc $(s, l, c) \xrightarrow{D} \text{closure}(s, l, \text{next}(c))$ if $c < c_m$.

Note that there is a single arc from a node (s, l, c) due to a transition (s, a, s') in the ASMC because of the assumption of determinism of \mathcal{A} and that there is at most one **D** arc from a node. Observe also that we evaluate the guard of a transition with respect to the open interval $(c, \text{next}(c))$ based on the straightforward result that, given a finite set of clock values (here, C), the probability that an ASMC performs a transition when the value of the clock belongs to this set is null.

We now define TRS as the set of nodes reachable from the set of states $(s, l, 0)$, for all $s \in S$ and $l \in \text{Init}$, with $s \models_{\Sigma} \Lambda(l)$, by traversing the arcs expressed by the four rules above (note that we consider all states $s \in S$ because satisfaction needs to be checked on all states of \mathcal{M}). Then, the TRG of $\mathcal{M} \times \mathcal{A}$ is defined as the graph over TRS where the arcs are described as above.

Observe that, if (s, l, c) is a node of the TRG, then any $(s, l, c + \delta)$ with $0 \leq \delta < \text{next}(c) - c$ is a state of the MRP $N(t)$ and that a **D**-arc (respectively, **M_res**-arc) to a node (s, l, c) means that upon event **D** (respectively, **M_res**) the state of $\mathcal{M} \times \mathcal{A}$ is exactly (s, l, c) , while if the same state is entered through an **M**-arc, the state of the process can be $(s, l, c + \delta)$ for any $0 < \delta < \text{next}(c)$.

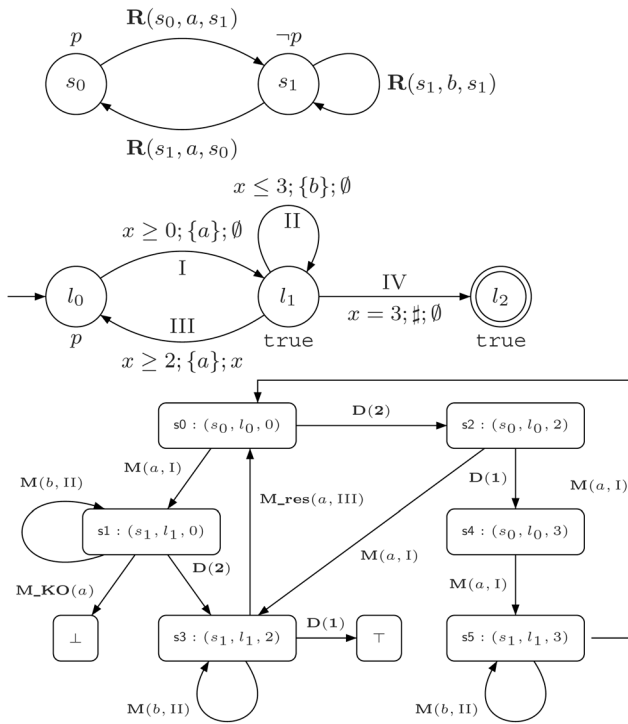
The upper part of Fig. 7 shows an ASMC \mathcal{M} and a DTA \mathcal{A} . DTA edges have been tagged with roman numerals to cross reference them in the TRG of $\mathcal{M} \times \mathcal{A}$ shown in the lower part. Let us consider two paths in the TRG and, for each path, the corresponding realizations in the stochastic process $\mathcal{M} \times \mathcal{A}$. The path $s0, s1, s3, \top$ corresponds to process evolutions in which an a -event occurs with clock x in $(0, 2)$ and then time elapses until clock reaches 3. Note that the intermediate state $s3$ corresponds to reaching time 2. The path $s0, s1, s1, \perp$ captures those process evolutions in which an a -event is followed by a b -event and then again an a -event, all occurring with a clock in $(0, 2)$. As the last event cannot be mimicked by the DTA, the process reaches \perp .

3.3 Building the Embedded DTMC

To compute the probability of reaching \top , we need to identify in the TRG the states of the DTMC Y and the associated transition probabilities. The states are defined according to the specification of the MRP.

Definition 3.3. Let $s \in TRS$. Then, s is a state of the DTMC embedded into the MRP (Y, T) if either

1. $s = (s, l, c)$ with $l \in \text{Init}$ and $c = 0$ (initial states),
2. s can be entered by an arc labeled **D** or **M_res**,
3. $s = \top$, or
4. $s = \perp$.

Fig. 7. An ASMC \mathcal{M} , a DTA \mathcal{A} , and their TRG.

Note that not all states of the TRG are states of the DTMC: Indeed, in the example of Fig. 7, state s_5 is not a state of the DTMC. Intuitively speaking, in order to reach s_5 , the ASMC must perform an a -event *after* the clock has reached 3.

We have represented in Fig. 8 the graph associated with the embedded DTMC of the example depicted in Fig. 7. An arc between two states means that there is a nonnull probability to reach the destination from the source without going through a regeneration point. For instance, there is an edge from s_4 to s_0 because one possible path (in the TRG) goes through a M move followed by a M_{res} move triggering the regeneration point.

To compute the probabilities of the DTMC (i.e., to label the edges of the associated graph), we need to define, for each state $(s, l, c) \in TRS \setminus \{\perp, \top\}$ of the DTMC, how the process $\mathcal{M} \times \mathcal{A}$ can evolve before reaching the next regeneration point. This (transient) behavior is driven by the *subordinated* CTMC $\mathcal{C}_{(s,l,c)}$, which describes the evolution of the process from (s, l, c) until a successive state of $\mathcal{M} \times \mathcal{A}$ is reached, either due to a state change in \mathcal{M} , due to the clock having reached next(c), or due to the clock being reset.

The states of the subordinated CTMC $\mathcal{C}_{(s,l,c)}$ can be computed using, again, the TRG. From (s, l, c) we take in the TRG the transitive closure over arcs of type M , possibly followed by a M_{res} -arc or a M_{KO} -arc. More formally, the states of the subordinated CTMC $\mathcal{C}_{(s,l,c)}$ are defined as follows:

- $(s, l, c) \in \mathcal{C}_{(s,l,c)}$;
- $(s', l', c) \in \mathcal{C}_{(s,l,c)}$ if there exists a path in the TRG from (s, l, c) to (s', l', c) in which all arcs are of type M ;
- $\top \in \mathcal{C}_{(s,l,c)}$ if there exists a path in the TRG from (s, l, c) to \top in which either all arcs are of type M , or

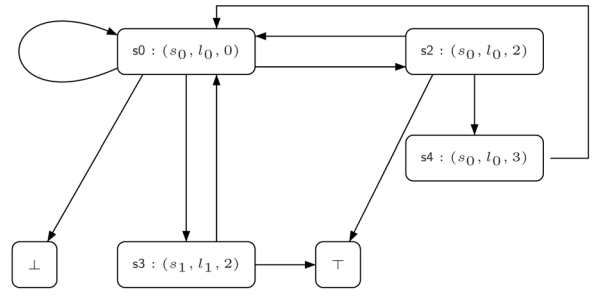


Fig. 8. The embedded DTMC.

the path ends in an M_{res} -arc and all preceding arcs (if any) are of type M ;

- $(s', l', 0)^{Reset} \in \mathcal{C}_{(s,l,c)}$ if there exists a (possibly empty) path in the TRG from (s, l, c) to (s', l', c) of arcs all of type M , and an arc from (s', l', c) to $(s', l', 0)$ of type M_{res} ;
- $\perp \in \mathcal{C}_{(s,l,c)}$ if there exists a (possibly empty) path in the TRG from (s, l, c) to (s', l', c) of arcs all of type M and an M_{KO} arc from (s', l', c) to \perp .

We distinguish in the subordinated CTMCs the state $(s, l, 0)^{Reset}$, entered upon a clock reset, from the state $(s, l, 0)$ entered through a Markovian transition. Indeed, the first state corresponds to the next regeneration point, whereas the second one is only an intermediate state (see below). Observe that, when $c > 0$ it is never the case that (s, l, c) can be entered through a non-Markovian transition.¹ The subset of “reset” states is denoted $Reset$. Summarizing, the states of the subordinated CTMC $\mathcal{C}_{(s,l,c)}$ are of the form \perp , \top , $(s', l', 0)^{Reset}$, or (s', l', c') for $s' \in S$, $l' \in L$, and $c' \in \{0, \dots, c\}$. We shall indicate with s a generic state of the DTMC, of the various forms indicated above. Note that there is a subordinated CTMC built only for those states s of form $s = (s, l, c)$ that we will denote as $\mathcal{C}_{(s,l,c)}$ or, equivalently, as \mathcal{C}_s .

Fig. 9 depicts the subordinated chain $\mathcal{C}_{(s_0, l_0, 0)}$ for state $(s_0, l_0, 0)$. Observe that this subordinated chain is derived from the TRG but is not a subgraph of it due to the duplication of state $(s_0, l_0, 0)$. This CTMC represents all behaviors during the evolution of clock x in $(0, 2)$ until a regeneration point is reached. The (non-time-triggered) regeneration points correspond to the absorbing states \perp and $(s_0, l_0, 0)^{Reset}$. Let us interpret the state of this CTMC at time 2. If it is \perp (resp., $(s_0, l_0, 0)^{Reset}$), then it means that the next regeneration point is \perp (respectively, $(s_0, l_0, 0)$) since we have reached it *before* 2. If this state is $(s_0, l_0, 0)$ (respectively, $(s_1, l_1, 0)$), it means that the next regeneration point corresponds to $x = 2$ and we follow the corresponding D edge in the TRG to determine the next regeneration point, here $(s_0, l_0, 2)$ (respectively, $(s_1, l_1, 2)$). Therefore, the probabilities of DTMC transitions from $(s_0, l_0, 0)$ are obtained from the transient probability distribution of the subordinated chain at time 2.

1. Note that, in the TRG, an M_{res} transition corresponds in Deterministic Stochastic Petri Nets to the case of an exponential transition that preempts a deterministic transition and then immediately reenables it, which, as explained in [26], requires a duplication of the states of the subordinated CTMC.

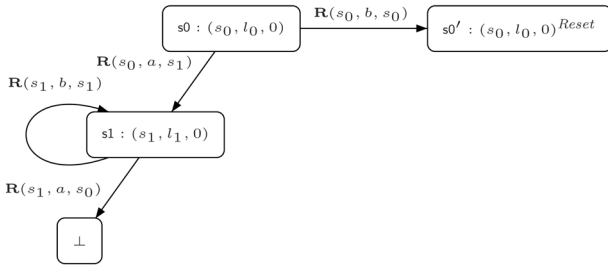


Fig. 9. The subordinated CTMC with respect to $(s_0, l_0, 0)$.

We can now generalize the example to consider the rates of the DTMC in the general case. Let $P_{s,s'}$ be the transition probabilities of the DTMC. The elements of $P_{s,s'}$ are computed using the subordinated CTMC \mathcal{C}_s and all of the elements of row s of P are computed on the same subordinated CTMC \mathcal{C}_s . Rows corresponding to the DTMC states \top and \perp are obviously identically zero since they are absorbing. Let $s = (s, l, c)$. We denote by $\pi_s(\tau)$ the transient (respectively, steady state) distribution of \mathcal{C}_s at time τ when τ is finite (respectively, when $\tau = \infty$).

We are now in a position to give the formulas for the nonnull entries of P . By convention, in the following formulas, $\perp^{Reset} = \perp$, $\top^{Reset} = \top$, and, when $s' \notin \mathcal{C}_s$, then $\pi_s(\tau)(s') = 0$.

- If $c < c_m$, then $P_{s,s'}$ equals:

$$\sum_{s'' \in \mathcal{C}_s \setminus Reset \wedge s'' \xrightarrow{D} s'} \pi_s(\text{next}(c) - c)(s'') + \pi_s(\text{next}(c) - c)(s'^{Reset}).$$

The first term (sum over s'') corresponds to the case where the next regeneration is triggered by $x = \text{next}(c)$ and we look for D edges to determine the next regeneration point. Remember that $Reset$ is the set of $(\cdot, \cdot, 0)^{Reset}$ states. The second term corresponds to a regeneration point obtained by a clock reset $((s', l', 0)^{Reset})$ or by reaching $\{\perp, \top\}$.

- If $c = c_m$, then

$$P_{s,s'} = \pi_s(\infty)(s'^{Reset}),$$

for $s' \in \{\perp, \top\} \cup \{(s', l', 0)\}$

When $c = c_m$, the only way to obtain a regeneration point is to reset the clock or to reach $\{\perp, \top\}$.

The first case requires transient analysis of the subordinated CTMCs, which is usually performed by uniformization [28]. The second case only requires steady-state analysis, which is generally less computationally expensive.

Note that there are two peculiarities of the embedded DTMC. First, we can reenter the same state due to a clock reset. This has no effect on the computation. Second, the transition matrix can be substochastic because, for some DTMC states, there is a nonnull probability to never reach another state of the MRP. Again, this is not problematic, because the reachability probability computation with a substochastic matrix is identical to that with a stochastic transition matrix.

Finally, as often in a probabilistic setting, checking whether the set of paths of \mathcal{M} accepted by \mathcal{A} has probability 0 or 1 can be performed without any numerical computation. The only relevant information in the DTMC, given its transition probability \mathbf{P} , is (for every pair of states (s, s')) whether $\mathbf{P}(s, s') > 0$ and this information is obtained by a simple examination of the TRG.

4 EXPRESSIVENESS OF CSL^{TA}

In this section, we study the relationship between CSL^{TA}, CSL [9], and aCSL [10]. Formulas interpreted on ASMCs are described as being equivalent if, for any ASMC, the same states of the ASMC satisfy the formulas. Formally, we say that the formula Φ_1 (of the logic L_1 , with the satisfaction relation \models_{L_1}) and Φ_2 (of the logic L_2 , with the satisfaction relation \models_{L_2}) are *equivalent* if, for any ASMC \mathcal{M} , and for any state s of \mathcal{M} , we have $\mathcal{M}, s \models_{L_1} \Phi_1$ if and only if $\mathcal{M}, s \models_{L_2} \Phi_2$. The logic L_1 is *at least as expressive as* the logic L_2 if, for each formula Φ_2 of the logic L_2 , there exists a formula Φ_1 of the logic L_1 such that Φ_1 and Φ_2 are equivalent. The logic L_1 is *strictly more expressive than* the logic L_2 if L_1 is at least as expressive as L_2 and there exists a formula Φ_1 of L_1 such that there does not exist an equivalent formula Φ_2 of logic L_2 . Given an ASMC \mathcal{M} with state set S and the satisfaction relation \models_L of the logic L , let $Sat_L^{\mathcal{M}}(\Phi) = \{s \in S \mid \mathcal{M}, s \models_L \Phi\}$ (when \mathcal{M} is clear from the context, we write $Sat_L(\Phi)$).

4.1 CSL^{TA} Is at Least as Expressive as CSL

In this section, we recall the definition of CSL [9].

Definition 4.1. *The syntax of CSL is defined as follows:*

$$\Phi ::= p \mid \Phi \wedge \Phi \mid \neg \Phi \mid \mathcal{S}_{\sim \lambda}(\Phi) \mid \mathcal{P}_{\sim \lambda}(\mathcal{X}^I \Phi) \mid \mathcal{P}_{\sim \lambda}(\Phi \mathcal{U}^I \Phi)$$

where $a \in AP$ is an atomic proposition, $I \subseteq \mathbb{R}_{\geq 0}$ is a nonempty interval, $\sim \in \{<, \leq, \geq, >\}$ is a comparison operator, and $\lambda \in [0, 1]$ is a probability.

For any infinite path $\sigma = s_0 \xrightarrow{a_0, \tau_0} s_1 \xrightarrow{a_1, \tau_1} \dots$ and $t \in \mathbb{R}_{\geq 0}$, if i is the smallest index such that $t \leq \sum_{j=0}^i \tau_j$, then we let $\sigma @ t = s_i$; that is, $\sigma @ t$ is used to denote the state along σ occupied at time t .

Definition 4.2. *For $\mathcal{M} = \langle S, Act, AP, lab, \mathbf{R} \rangle$ and state $s \in S$, the satisfaction relation \models_{CSL} is defined as follows:*

$$\begin{aligned} \mathcal{M}, s \models_{\text{CSL}} p &\Leftrightarrow p \in lab(s), \\ \mathcal{M}, s \models_{\text{CSL}} \Phi_1 \wedge \Phi_2 &\Leftrightarrow \mathcal{M}, s \models_{\text{CSL}} \Phi_1 \quad \text{and} \\ &\quad \mathcal{M}, s \models_{\text{CSL}} \Phi_2, \\ \mathcal{M}, s \models_{\text{CSL}} \neg \Phi &\Leftrightarrow \mathcal{M}, s \not\models_{\text{CSL}} \Phi, \\ \mathcal{M}, s \models_{\text{CSL}} \mathcal{S}_{\sim \lambda}(\Phi) &\Leftrightarrow \sum_{s' \in Sat_{\text{CSL}}^{\mathcal{M}}(\Phi)} \pi(s, s') \sim \lambda, \\ \mathcal{M}, s \models_{\text{CSL}} \mathcal{P}_{\sim \lambda}(\varphi) &\Leftrightarrow \\ &\quad \text{Pr}_s^{\mathcal{M}}\{\sigma \in Path^{\mathcal{M}}(s) \mid \mathcal{M}, \sigma \models_{\text{CSL}} \varphi\} \sim \lambda, \\ \mathcal{M}, \sigma \models_{\text{CSL}} \mathcal{X}^I \Phi &\Leftrightarrow \mathcal{M}, \sigma(1) \models_{\text{CSL}} \Phi \quad \text{and} \\ &\quad \sigma = s \xrightarrow{a, \tau} \sigma' \quad \text{with } \tau \in I, \\ \mathcal{M}, \sigma \models_{\text{CSL}} \Phi_1 \mathcal{U}^I \Phi_2 &\Leftrightarrow \exists t \in I. \mathcal{M}, \sigma @ t \models_{\text{CSL}} \Phi_2 \quad \text{and} \\ &\quad \forall t' \in [0, t). \mathcal{M}, \sigma @ t' \models_{\text{CSL}} \Phi_1. \end{aligned}$$

In the following, when clear from the context, we write $s \models_{\text{CSL}} \Phi$ for \mathcal{M} , $s \models_{\text{CSL}} \Phi$ and $\sigma \models_{\text{CSL}} \Phi$ for \mathcal{M} , $\sigma \models_{\text{CSL}} \Phi$.

The following proposition shows that CSL^{TA} is at least as expressive as CSL.

Proposition 4.3. *For any formula Φ of CSL, there is a formula Φ' of CSL^{TA} equivalent to Φ . The size of Φ' is linear with respect to the size of Φ .*

Proof. The semantics of constructors for state formulas are identical for CSL and CSL^{TA}; therefore, it suffices to prove that any path formula of CSL is equivalent to some path formula of CSL^{TA}. The idea of the proof is to translate the path operator $\mathcal{X}^{[\alpha, \beta]} \Phi$ of CSL with the DTA $\mathcal{A}_{\mathcal{X}^{[\alpha, \beta]}}$ of Fig. 4 and the path operator $\Phi_1 \mathcal{U}^{[\alpha, \beta]} \Phi_2$ with the DTA $\mathcal{A}_{\Phi_1 \mathcal{U}^{[\alpha, \beta]} \Phi_2}$ of Fig. 5. In the following, we concentrate on the case in which the time interval of a CSL formula is of the form $[\alpha, \beta]$, where $\alpha > 0$ (the translation of path operators with time intervals other than $[\alpha, \beta]$ is similar and will be discussed briefly at the end of the proof). Therefore, our task consists in showing that, for a given ASMC $\mathcal{M} = \langle S, Act, AP, lab, \mathbf{R} \rangle$, a given state $s \in S$, and an infinite path $\sigma \in Path^{\mathcal{M}}(s)$, we have:

1. $\sigma \models_{\text{CSL}} \Phi_1 \mathcal{U}^{[\alpha, \beta]} \Phi_2$, iff $\sigma \in AccPath^{\mathcal{M}}(s, \mathcal{A}_{\Phi_1 \mathcal{U}^{[\alpha, \beta]} \Phi_2})$, and
2. $\sigma \models_{\text{CSL}} \mathcal{X}^{[\alpha, \beta]} \Phi$, iff $\sigma \in AccPath^{\mathcal{M}}(s, \mathcal{A}_{\mathcal{X}^{[\alpha, \beta]} \Phi})$.

Consider Case 1. We show first that $\sigma \models_{\text{CSL}} \Phi_1 \mathcal{U}^{[\alpha, \beta]} \Phi_2$ implies $\sigma \in AccPath^{\mathcal{M}}(s, \mathcal{A}_{\Phi_1 \mathcal{U}^{[\alpha, \beta]} \Phi_2})$. Let $\sigma_{\mathcal{M}} = s_0 \xrightarrow{a_0, \tau_0} s_1 \xrightarrow{a_1, \tau_1} \dots$ be a path such that $\sigma_{\mathcal{M}} \models_{\text{CSL}} \Phi_1 \mathcal{U}^{[\alpha, \beta]} \Phi_2$. Then, by Definition 4.2, there exists $t \in [\alpha, \beta]$ such that $\sigma_{\mathcal{M}} @ t \models_{\text{CSL}} \Phi_2$ and $\sigma_{\mathcal{M}} @ t' \models_{\text{CSL}} \Phi_1$ for all $t' < t$. There are two cases to consider:

- 1.a. Φ_2 is satisfied when the value of the clock x reaches α ;
- 1.b. Φ_2 is not yet satisfied when the value of the clock x reaches α .

Consider Case 1.a. Observe that there exists some $i \in \mathbb{N}$ such that $s_i \models_{\text{CSL}} \Phi_1 \wedge \Phi_2$, $s_j \models_{\text{CSL}} \Phi_1$ for all $j < i$, and $\sum_{k=0}^{i-1} \tau_k < \alpha \leq \sum_{k=0}^i \tau_k$. We assert that the ASMC path $\sigma_{\mathcal{M}}$ is accepted by the path $\sigma_{\mathcal{A}} = (l_0, \bar{x}_0) \xrightarrow{\tau_0, e_{00}} \dots (l_0, \bar{x}_{i-1}) \xrightarrow{\tau_{i-1}, e_{00}} (l_0, \bar{x}_i) \xrightarrow{\tau_i, e_{02}} (l_2, \bar{x}_{i+1})$ of $\mathcal{A}_{\Phi_1 \mathcal{U}^{[\alpha, \beta]} \Phi_2}$ (that is, to accept $\sigma_{\mathcal{M}}$, the DTA performs i loops of the edge e_{00} , then traverses the edge e_{02}). To verify that $\sigma_{\mathcal{M}}$ is accepted by $\sigma_{\mathcal{A}}$, consider Definition 2.7. First, observe that $l_0 \in Init$, $\bar{x}_0 = 0$ and $l_2 \in Final$. Second, recalling that $\Lambda(l_0) = \Phi_1$, we observe that $s_j \models_{\text{CSL}} \Lambda(l_0)$ for all $j \leq i$. Furthermore, recalling that $\Lambda(l_2) = \Phi_2$, we observe that $s_i \models_{\text{CSL}} \Lambda(l_2)$. Third, we have that $a_j \in Act$, and hence, $a_j \in \text{action}(e_{00})$, for all $j < i$. The final requirements of Definition 2.7, which concern the durations of $\sigma_{\mathcal{M}}$ and $\sigma_{\mathcal{A}}$, follow directly from the observation that the durations of the transitions of $\sigma_{\mathcal{A}}$ are equal to the durations of the first i transitions of $\sigma_{\mathcal{M}}$.

Now consider Case 1.b. In this case, there exists $i \in \mathbb{N}$ such that $s_i \models_{\text{CSL}} \Phi_2$, $s_j \models_{\text{CSL}} \Phi_1$ for all $j < i$, and $\alpha < \sum_{k=0}^{i-1} \tau_k < \beta$. Furthermore, we let $i^\alpha < i$ be the largest index for which $\sum_{k=0}^{i^\alpha-1} \tau_k \leq \alpha$. Intuitively, the

state s_{i^α} will be the state along $\sigma_{\mathcal{M}}$ when α time units have elapsed. Let $\tau' = \alpha - \sum_{k=0}^{i^\alpha-1} \tau_k$ and $\tau'' = \tau_{i^\alpha} - \tau'$. We claim that the ASMC path $\sigma_{\mathcal{M}}$ is accepted by the path $\sigma_{\mathcal{A}} = (l_0, \bar{x}_0) \xrightarrow{\tau_0, e_{00}} \dots \xrightarrow{\tau_{i-1}, e_{00}} (l_0, \bar{x}_{i^\alpha}) \xrightarrow{\tau_{i^\alpha}, e_{01}} (l_1, \bar{x}_{i^\alpha+1}) \xrightarrow{\tau_{i^\alpha+1}, e_{11}} (l_1, \bar{x}_{i^\alpha+2}) \xrightarrow{\tau_{i^\alpha+2}, e_{11}} \dots \xrightarrow{\tau_{i-1}, e_{11}} (l_1, \bar{x}_i) \xrightarrow{\tau_i, e_{12}} (l_2, \bar{x}_{i+1})$ of $\mathcal{A}_{\Phi_1 \mathcal{U}^{[\alpha, \beta]} \Phi_2}$ (that is, to accept $\sigma_{\mathcal{M}}$, the DTA performs i^α loops of the edge e_{00} , then traverses the edge e_{01} , then performs $i - (i^\alpha + 1)$ loops of the edge e_{11} , then traverses the edge e_{12}). Consider Definition 2.7: We note that the index described in point 2 of Definition 2.7 is i and the function $\kappa : \{0, \dots, i+1\} \rightarrow \{0, \dots, i\}$ is defined by $\kappa(j) = j$ for all $j \leq i^\alpha$ and $\kappa(j) = j-1$ for all $i^\alpha < j \leq i+1$. We now verify that the choice of $\sigma_{\mathcal{A}}$, index i , and function κ satisfies the conditions of Definition 2.7. First, observe that $l_0 \in Init$, $\bar{x}_0 = 0$ and $l_2 \in Final$. Second, recalling that $\Lambda(l_0) = \Phi_1$ ($\Lambda(l_1) = \Phi_1 \wedge \neg \Phi_2$, $\Lambda(l_2) = \Phi_2$, respectively), we observe that $s_j \models_{\text{CSL}} \Lambda(l_0)$ for all $j \leq i^\alpha$ ($s_j \models_{\text{CSL}} \Lambda(l_1)$ for all $i^\alpha < j < i$, $s_i \models_{\text{CSL}} \Lambda(l_2)$, respectively). Third, we have that $a_j \in Act$ and, hence, $a_j \in \text{action}(e_{00})$, $a_j \in \text{action}(e_{11})$, and $a_i \in \text{action}(e_{12})$, for all $j \leq i$. The final requirements of Definition 2.7, concerning the durations of $\sigma_{\mathcal{M}}$ and $\sigma_{\mathcal{A}}$, follow by the following facts: for all $j \leq i$ such that $j \neq i^\alpha$, we have that the duration of the j th transition of $\sigma_{\mathcal{M}}$ is equal to the duration of the $\kappa(j)$ th transition of $\sigma_{\mathcal{A}}$; furthermore, $\tau_{i^\alpha} = \tau' + \tau''$.

The reverse direction of Case 1 follows in a similar manner and we omit the details.

Consider Case 2. We show that $\sigma \models_{\text{CSL}} \mathcal{X}^{[\alpha, \beta]} \Phi$ implies $\sigma \in AccPath^{\mathcal{M}}(s, \mathcal{A}_{\mathcal{X}^{[\alpha, \beta]} \Phi})$. Let $\sigma_{\mathcal{M}}$ be a path such that $\mathcal{M}, \sigma_{\mathcal{M}} \models_{\text{CSL}} \mathcal{X}^{[\alpha, \beta]} \Phi$. Then, by Definition 4.2, we have $\sigma_{\mathcal{M}} = s \xrightarrow{a, \tau} \sigma'$ for some $a \in Act$, $\tau \in I$, and path σ' , where $\sigma_{\mathcal{M}}(1) \models_{\text{CSL}} \Phi$. By the definition of the DTA $\mathcal{A}_{\mathcal{X}^{[\alpha, \beta]} \Phi}$, there exists a path $(l_0, 0) \xrightarrow{\tau, e} (l_1, \tau)$, where $e = (l_0, \alpha \leq x \leq \beta, Act, \emptyset, l_1)$ is the edge from l_0 to l_1 . From the fact that $l_0 \in Init$, $l_1 \in Final$, and $\sigma_{\mathcal{M}}(1) \models_{\text{CSL}} \Phi$, we have that $\sigma_{\mathcal{M}}$ is accepted by $\mathcal{A}_{\mathcal{X}^{[\alpha, \beta]} \Phi}$ according to the criteria of Definition 2.7. The reverse direction of Case 2 follows in a similar manner.

We now consider briefly the case for other types of time intervals. For the Next operator, the DTA of Fig. 4 requires only modifications to the guard of its single edge (for example, the time interval (α, ∞) is represented by the guard $x > \alpha$). Similarly, open or half-open time intervals of the Until operator can be represented by changing the associated inequalities of constraints from nonstrict to strict: for example, the time interval $(\alpha, \beta]$ can be represented by changing the constraint $x \leq \alpha$ to $x < \alpha$ in the DTA of Fig. 5. For a time interval of the form $[\alpha, \infty)$ or (α, ∞) for $\alpha > 0$, the guards of the form $x \leq \beta$ in the DTA of Fig. 5 are changed to true. Instead, for a time interval of the form $[0, \beta]$ or $[0, \beta)$, the location l_0 and its outgoing edges are removed and both l_1 and l_2 become initial locations.

Finally, the assertion on formula sizes is straightforward. \square

We observe that the verification of a CSL formula of the form $\mathcal{P}_{\sim\lambda}(\Phi_1 U^I \Phi_2)$ and a CSL^{TA} formula of the form $\mathcal{P}_{\sim\lambda}(\mathcal{A}_{\Phi_1 U^I \Phi_2})$ involve similar computation steps: for example, in the case of $I = [\alpha, \beta]$ with $\alpha > 0$, a transient analysis of two CTMCs is required, both in the CSL model-checking algorithm of [9] and in the CSL^{TA} model-checking algorithm of Section 3. The computational complexity of model checking CSL^{TA} properties transformed from equivalent CSL properties is the same as that for model checking the original CSL properties with the algorithm of [9].

4.2 CSL^{TA} Is at Least as Expressive as asCSL

In this section, we recall the stochastic temporal logic asCSL [10] and show that every asCSL formula can be expressed as a CSL^{TA} formula.

4.2.1 Definition of asCSL

First, we present the syntax and semantics of asCSL. In contrast to the original presentation of asCSL in [10], we consider *nondeterministic program automata* as path operators, as opposed to regular expressions (called programs in [10]). As asCSL programs can be translated into nondeterministic program automata, the presentation of asCSL is as general as the original presentation with regular expressions. Also note that we use the special action \surd , which in a similar way to \sharp , allows a transition in the automaton without a corresponding transition in the ASMC. Note the distinction between \sharp -labeled transitions of DTA and \surd -labeled transitions of nondeterministic program automata: The latter are not triggered by behavior of the ASMC, whereas, in contrast, \sharp -labeled transitions are triggered by the passage of time.

Definition 4.4. A nondeterministic program automaton (NPA) is a tuple $\mathcal{N} = \langle Z, \Xi, \delta, Z_{Init}, Z_F \rangle$, where Z is a finite set of states with the set $Z_{Init} \subseteq Z$ of initial states and the set $Z_F \subseteq Z$ of final states, Ξ is a finite input alphabet, and $\delta : Z \times \Xi \rightarrow 2^Z$ is a transition function. We say that \mathcal{N} is a deterministic program automaton (DPA) if $|Z_{Init}| = 1$ and $|\delta(z, u)| \leq 1$ for each $z \in Z$ and $u \in \Xi$.

Definition 4.5. The syntax of asCSL is defined as follows:

$$\Phi ::= p \mid \Phi \wedge \Phi \mid \neg\Phi \mid \mathcal{S}_{\sim\lambda}(\Phi) \mid \mathcal{P}_{\sim\lambda}(\mathcal{N}(\Xi)^I),$$

where $a \in AP$ is an atomic proposition, $I \subseteq \mathbb{R}_{\geq 0}$ is a nonempty interval, $\sim \in \{<, \leq, \geq, >\}$ is a comparison operator, $\lambda \in [0, 1]$ is a probability, and $\mathcal{N}(\Xi)$ is an NPA with input alphabet Ξ such that $\Xi \subseteq \{(\Phi, b) \mid \Phi \text{ is an asCSL formula } \wedge b \in Act \cup \{\surd\}\}$.

We write $z \xrightarrow{u} z'$ to denote $z' \in \delta(z, u)$ (that is, to denote a transition of an NPA).

Definition 4.6. A run of an NPA \mathcal{N} is a (finite) sequence $r = z_0 \xrightarrow{u_0} z_1 \xrightarrow{u_1} \dots \xrightarrow{u_{n-1}} z_n$ of transitions of \mathcal{N} . If $z_0 \in Z_{Init}$ and $z_n \in Z_F$, then we say that the run r is accepting. We use $last(r)$ to denote the last state z_n of r . For a state z of \mathcal{N} , let $\text{Runs}^{\mathcal{N}}(z)$ be the set of runs of \mathcal{N} with the first state z and, for $Z' \subseteq Z$, let $\text{Runs}^{\mathcal{N}}(Z') = \bigcup_{z \in Z'} \text{Runs}^{\mathcal{N}}(z)$.

Definition 4.7. For $\mathcal{M} = \langle S, Act, AP, lab, \mathbf{R} \rangle$ and state $s \in S$, the satisfaction relation \models_{asCSL} is defined as follows:

$$\begin{aligned} \mathcal{M}, s \models_{\text{asCSL}} p &\iff p \in lab(s), \\ \mathcal{M}, s \models_{\text{asCSL}} \neg\Phi &\iff \mathcal{M}, s \not\models_{\text{asCSL}} \Phi, \\ \mathcal{M}, s \models_{\text{asCSL}} \Phi_1 \wedge \Phi_2 &\iff \mathcal{M}, s \models_{\text{asCSL}} \Phi_1 \text{ and } \\ &\quad \mathcal{M}, s \models_{\text{asCSL}} \Phi_2, \\ \mathcal{M}, s \models_{\text{asCSL}} \mathcal{S}_{\sim\lambda}(\Phi) &\iff \sum_{s' \in Sat_{\text{asCSL}}^{\mathcal{M}}(\Phi)} \pi(s, s') \sim \lambda, \\ \mathcal{M}, s \models_{\text{asCSL}} \mathcal{P}_{\sim\lambda}(\mathcal{N}(\Xi)^I) &\iff \text{Pr}_s^{\mathcal{M}}(\text{AccPath}^{\mathcal{M}}(s, (\mathcal{N}(\Xi)^I)) \sim \lambda). \end{aligned}$$

where $\text{AccPath}^{\mathcal{M}}(s, \mathcal{N}(\Xi)^I)$ is defined in the following way: Let z be a state of \mathcal{N} and σ be a finite path of \mathcal{M} . We define $\text{Runs}^{\mathcal{N}}(z, \sigma)$ as the greatest set of runs $z \xrightarrow{\Phi_0 b_0} z_1 \xrightarrow{\Phi_1 b_1} \dots \xrightarrow{\Phi_{n-1} b_{n-1}} z_n$ such that:

1. $z \in \text{Runs}^{\mathcal{N}}(z, \sigma)$, if and only if $|\sigma| = 0$;
2. if $z \xrightarrow{\Phi_0 b_0} z_1 \xrightarrow{\Phi_1 b_1} \dots \xrightarrow{\Phi_{n-1} b_{n-1}} z_n \in \text{Runs}^{\mathcal{N}}(z, \sigma)$ and $n \geq 1$, then

- $\mathcal{M}, \sigma(0) \models_{\text{asCSL}} \Phi_0$;
- if $b_0 \in Act$, then $\sigma = s \xrightarrow{b_0, \tau} \sigma'$ with $z_1 \xrightarrow{\Phi_1 b_1} \dots \xrightarrow{\Phi_{n-1} b_{n-1}} z_n \in \text{Runs}^{\mathcal{N}}(z_1, \sigma')$;
- if $b_0 = \surd$, then $z_1 \xrightarrow{\Phi_1 b_1} \dots \xrightarrow{\Phi_{n-1} b_{n-1}} z_n \in \text{Runs}^{\mathcal{N}}(z_1, \sigma)$.

For the set $Z' \subseteq Z$ of states of \mathcal{N} , we let $\text{Runs}^{\mathcal{N}}(Z', \sigma) = \bigcup_{z \in Z'} \text{Runs}^{\mathcal{N}}(z, \sigma)$. The set of ASMC paths accepted by \mathcal{N} with time interval $I \subseteq \mathbb{R}_{\geq 0}$, denoted by $\text{AccPath}^{\mathcal{M}}(s, \mathcal{N}^I)$, is defined as $\{\sigma \in \text{Path}^{\mathcal{M}}(s) \mid \exists \text{ finite prefix } \sigma_{\text{fin}} \text{ of } \sigma \text{ s.t. } \exists \text{ accepting } r \in \text{Runs}^{\mathcal{N}}(Z_{Init}, \sigma_{\text{fin}}) \text{ and } \tau(\sigma_{\text{fin}}) \in I\}$.

When clear from the context, we write $\text{Runs}(Z', \sigma)$ instead of $\text{Runs}^{\mathcal{N}}(Z', \sigma)$.

4.2.2 From NPA to DTA

To show that every asCSL formula can be expressed as a CSL^{TA} formula, it suffices to show how a formula $\mathcal{P}_{\sim\lambda}(\mathcal{N}(\Xi)^I)$ can be encoded as a CSL^{TA} formula $\mathcal{P}_{\sim\lambda}(\mathcal{A})$, similarly to the translation from CSL to CSL^{TA} of Section 4.1. In order to obtain the required DTA \mathcal{A} from $\mathcal{N}(\Xi)^I$, two steps are required: First, it is necessary to construct a DPA, which will be used as the graph of \mathcal{A} , from the NPA $\mathcal{N}(\Xi)$, using a standard subset construction; then, it is necessary to represent the time interval I within \mathcal{A} using clock guards, in particular to constrain the global time of entry to final locations to those times in the interval I .

We first present the determinization of an NPA. Note that \surd -transitions are eliminated from the NPA to obtain the determinized automaton; unlike the case of \sharp -transitions in DTA, \surd -transitions do not have priority over other transitions. Hence, if a state z has an outgoing a -transition (for $a \in Act$) and an outgoing \surd -transition leading to a state with an outgoing a -transition, when reading an a -transition of an ASMC, there will be a nondeterministic choice between the a -transition and the \surd -transition from z . We choose to eliminate sequences of \surd -transitions in order to avoid such situations, noting that we also have to consider the case in which sequences of \surd -transitions which reach final states are taken after an action.

We require the following notation: Let $\mathcal{N} = \langle Z, \Xi, \delta, Z_{Init}, Z_F \rangle$ be an NPA, $Z' \subseteq Z$, Φ be an asCSL formula, and $a \in Act$. Then, we let $\text{Runs}(Z', (\Phi, a))$ equal:

$$\{z_0 \xrightarrow{\Phi_0} \sqrt{} \dots \xrightarrow{\Phi_{n-2}} \sqrt{} z_{n-1} \xrightarrow{\Phi_{n-1}^a} z_n \in \text{Runs}(Z') \mid z_0 \in Z' \wedge \Phi \Rightarrow \bigwedge_{i < n} \Phi_i\}.$$

Therefore, $\text{Runs}(Z', (\Phi, a))$ is the set of runs of \mathcal{N} starting from a state in Z' which perform $\sqrt{}$ -transitions before performing a single a -transition, where the conjunction of the asCSL formulae labeling transitions along the path is implied by Φ . We also let $\text{Runs}(Z', (\Phi, \sqrt{}), Z_F)$ equal:

$$\{z_0 \xrightarrow{\Phi_0} \sqrt{} \dots \xrightarrow{\Phi_{n-2}} \sqrt{} z_n \in \text{Runs}(Z') \mid z_0 \in Z' \wedge \Phi \Rightarrow \bigwedge_{i < n} \Phi_i \wedge z_n \in Z_F\}.$$

Hence, $\text{Runs}(Z', (\Phi, \sqrt{}), Z_F)$ is the set of runs from a state in Z' to a state in Z_F which perform $\sqrt{}$ -transitions only, where the conjunction of formulas along the path is implied by Φ (it is possible to have a run of length 0 featuring a single state, which must be both in Z' and Z_F).

Before defining the DPA corresponding to an NPA, we must identify the set of formulas which can appear in transition labels of the DPA, and which will be used subsequently to define the set of state propositions of the DTA. This set of formulas will be constructed such that an ASMC path is accepted by at most one run of the DTA obtained from an NPA. In particular, we construct a set of disjoint state propositions. Formally, let Σ be the smallest set of asCSL formulas such that:

1. $\Phi_1 \wedge \Phi_2 \Leftrightarrow \text{false}$ for all $\Phi_1, \Phi_2 \in \Sigma$;
2. for each Φ such that $\text{Runs}(Z, (\Phi, a)) \neq \emptyset$, for some $a \in \text{Act}$ or $\text{Runs}(Z, (\Phi, \sqrt{}), Z_F) \neq \emptyset$, there exists $\{\Phi_1, \dots, \Phi_n\} \subseteq \Sigma$ such that $\Phi \equiv \bigvee_{1 \leq i \leq n} \Phi_i$.

Definition 4.8. Let $\mathcal{N} = \langle Z, \Xi, \delta, Z_{\text{Init}}, Z_F \rangle$ be an NPA. The determinization of \mathcal{N} is the DPA $\text{det}(\mathcal{N}) = \langle Q, \Xi', \Delta, q_{\text{Init}}, Q_F \rangle$, where:

- $Q = 2^Z$, $q_{\text{Init}} = Z_{\text{Init}}$, and $Q_F = \{q \in 2^Z \mid \exists \Phi \in \Sigma \text{ s.t. } \text{Runs}(q, (\Phi, \sqrt{}), Z_F) \neq \emptyset\}$;
- $\Xi' = \{(\Phi, a) \mid \Phi \in \Sigma \wedge a \in \text{Act}\}$;
- $\Delta : Q \times \Xi' \rightarrow 2^Q$ is the transition function defined by

$$\Delta(q, (\Phi, a)) = \{z \in Z \mid \exists r \in \text{Runs}(Z, (\Phi, a)) \wedge \text{last}(r) = z\}$$

for all $q \in Q$ and $(\Phi, a) \in \Xi'$.

It can be verified that $|\Delta(q, (\Phi, a))| \leq 1$ for all $q \in Q$ and $(\Phi, a) \in \Xi'$.

We now define the DTA $\mathcal{A}(\mathcal{N}^I)$ by pushing the asCSL state formulas featured in transition labels in $\text{det}(\mathcal{N})$ into location labels of $\mathcal{A}(\mathcal{N}^I)$ and using the interval I in guards of edges leading directly to final locations (which correspond to the set E'' of DTA edges in the following definition). We also have to consider the case in which a final state of the NPA is reached at a time before the interval I (considered in the set E' of DTA edges below), which does not correspond to acceptance.

Definition 4.9. Let \mathcal{N} be an NPA, $[\alpha, \beta] \subseteq \mathbb{R}_{\geq 0}$ such that $\alpha > 0$, and $\text{det}(\mathcal{N}) = \langle Q, \Xi', \Delta, q_{\text{Init}}, Q_F \rangle$ be the determinization of \mathcal{N} . We let $\mathcal{A}(\mathcal{N}^{[\alpha, \beta]}) = \langle \Sigma, \text{Act}, (Q \times \Sigma) \cup (\overline{Q}_F \times \Sigma), \Lambda, \text{Init}, \text{Final}, \rightarrow \rangle$ be such that:

- $\overline{Q}_F = \{\overline{q} \mid q \in Q_F\}$;
- $\Lambda(q, \Phi) = \Lambda(\overline{q}, \Phi) = \Phi$ for each $q \in Q$ and $\Phi \in \Sigma$;
- $\text{Init} = \{(q_{\text{Init}}, \Phi) \mid \Phi \in \Sigma\}$;
-

$\text{Final} =$

$$\{(\overline{q}, \Phi) \in \overline{Q}_F \times \Sigma \mid q \in Q_F \wedge \text{Runs}(q, (\Phi, \sqrt{}), Z_F) \neq \emptyset\};$$

- \rightarrow is equal to the set:

$$\bigcup_{(q, \Phi) \in Q \times \Sigma, a \in \text{Act}, \Phi' \in \Sigma} \left\{ \begin{array}{l} E((q, \Phi), a, \Phi') \cup \\ E'((q, \Phi), a, \Phi') \cup \\ E''((q, \Phi), a, \Phi') \end{array} \right\}$$

where:

$$E((q, \Phi), a, \Phi') = \{(q, \Phi) \xrightarrow{\text{true}, a, \emptyset} (q', \Phi') \mid q' = \Delta(q, (\Phi, a)) \wedge q' \notin Q_F\};$$

$$E'((q, \Phi), a, \Phi') = \{(q, \Phi) \xrightarrow{x, < a, \emptyset} (q', \Phi') \mid q' = \Delta(q, (\Phi, a)) \wedge q' \in Q_F\};$$

$$E''((q, \Phi), a, \Phi') = \{(q, \Phi) \xrightarrow{a, < x, \leq \beta, a, \emptyset} (\overline{q'}, \Phi') \mid q' = \Delta(q, (\Phi, a)) \wedge q' \in Q_F\}.$$

Definition 4.9 considers the case in which $\alpha > 0$; the case for $\alpha = 0$ is similar, but the set E' of DTA edges is empty. It can be verified that $\mathcal{A}(\mathcal{N}^{[\alpha, \beta]})$ satisfies the conditions of initial determinism and determinism on actions of Definition 2.3. The other two conditions of Definition 2.3 are irrelevant because $\mathcal{A}(\mathcal{N}^{[\alpha, \beta]})$ does not have any boundary edges.

Proposition 4.10. Let \mathcal{M} be an ASMC, s be a state of \mathcal{M} , \mathcal{N} be an NPA, and $[\alpha, \beta] \subseteq \mathbb{R}_{\geq 0}$. Then, $\text{AccPath}^{\mathcal{M}}(s, \mathcal{N}^{[\alpha, \beta]}) = \text{AccPath}^{\mathcal{M}}(s, \mathcal{A}(\mathcal{N}^{[\alpha, \beta]}))$.

The proof of the proposition can be found in the appendix, which can be found in the Computer Society Digital Library at <http://doi.ieeecomputersociety.org/10.1109/TSE.2008.108>. From Proposition 4.10 and the observation that

$$\begin{aligned} \text{Pr}_s^{\mathcal{M}}(\text{AccPath}^{\mathcal{M}}(s, \mathcal{N}^{[\alpha, \beta]})) &= \text{Pr}_s^{\mathcal{M}}(\text{AccPath}^{\mathcal{M}}(s, \mathcal{N}^{(\alpha, \beta]})) \\ &= \text{Pr}_s^{\mathcal{M}}(\text{AccPath}^{\mathcal{M}}(s, \mathcal{N}^{[\alpha, \beta)})) = \text{Pr}_s^{\mathcal{M}}(\text{AccPath}^{\mathcal{M}}(s, \mathcal{N}^{(\alpha, \beta)})), \end{aligned}$$

the subsequent corollary then follows.

Corollary 4.11. Let \mathcal{M} be an ASMC, s be a state of \mathcal{M} , \mathcal{N} be an NPA, $I \subseteq \mathbb{R}_{\geq 0}$, $\sim \in \{<, \leq, \geq, >\}$, and $\lambda \in [0, 1]$. Then, $\mathcal{M}, s \models_{\text{asCSL}} \mathcal{P}_{\sim \lambda}(\mathcal{N}^I)$ if and only if $\mathcal{M}, s \models \mathcal{P}_{\sim \lambda}(\mathcal{A}(\mathcal{N}^I))$.

4.3 CSL^{TA} Is Strictly More Expressive than CSL

In this section, we give an example of a CSL^{TA} formula for which no equivalent CSL formula exists. We note that the




Fig. 10. A family of Markov chains $\mathcal{M}[\mu, \mu']$ and a DTA \mathcal{A} of CSL^{TA} .

DTA of the CSL^{TA} formula considered does not feature time constraints and, therefore, we can obtain an NPA which is equivalent to the DTA in a straightforward manner; hence, our result suffices also to show that there exists an asCSL formula for which no equivalent CSL formula exists.

Proposition 4.12. *There is a formula of CSL^{TA} for which there is no equivalent CSL formula.*

The proof of Proposition 4.12 follows a scheme that is different from proofs of similar results on expressiveness of temporal logics for transition systems. We first define the left-hand delimiter \langle for intervals, where \langle denotes either $[$ or $($. Similarly, the right-hand delimiter \rangle denotes either $]$ or $)$. Consider the family of ASMCs $\mathcal{M}[\mu, \mu']$ of Fig. 10 (left), for $0 < \mu, \mu' < 1$. Let Φ be a formula of CSL or CSL^{TA} (for simplicity, we write \models to denote the satisfaction relation of both CSL and CSL^{TA}). Then, $[\Phi](s) = \{(\mu, \mu') \in (0, 1)^2 \mid \mathcal{M}[\mu, \mu'], s \models \Phi\}$. For any $0 < \zeta < 1$, let $\Phi^\zeta = \mathcal{P}_{\geq \zeta}(\mathcal{A})$, where \mathcal{A} is the DTA depicted in Fig. 10 (right). It follows that $[\Phi^\zeta](s_0) = \{(\mu, \mu') \in (0, 1)^2 \mid \mu \cdot \mu' \geq \zeta\}$.

The following lemma specifies that, for any CSL formula and any state $s \in \{s_0, s_1, s_2, s_3\}$, the sets of parameters μ, μ' which result in the satisfaction of the CSL formula in state s of $\mathcal{M}[\mu, \mu']$ will be of a particular form. The lemma will then be used as the basis of our expressiveness result: Sets of parameters with the form described in the lemma cannot be used to obtain the set $[\Phi^\zeta](s_0) = \{(\mu, \mu') \in (0, 1)^2 \mid \mu \cdot \mu' \geq \zeta\}$ of parameters which result in the satisfaction of Φ^ζ in s_0 and, hence, no CSL formula is equivalent to Φ^ζ .

Lemma 4.13. *Let Φ be a formula of CSL. Then:*

1. for $i \in \{2, 3\}$, $[\Phi](s_i)$ is either $(0, 1)^2$ or \emptyset ;
2. $[\Phi](s_1)$ is a finite union of rectangles of the form $(0, 1) \times \langle a, b \rangle$;
3. $[\Phi](s_0)$ is a finite union of (open, closed, or mixed) rectangles of $(0, 1)^2$.

Proof. Assertion 1. When starting from s_2 or s_3 , the satisfaction of Φ does not depend on μ or μ' . Therefore, assertion 1 is satisfied trivially.

Assertion 2. We prove assertion 2 by induction on the size of the formula, taking into account all CSL operators one by one. Let Φ be a formula of CSL. If Φ is an atomic proposition, then $[\Phi](s_1)$ is either $(0, 1)^2$ or \emptyset . If $\Phi = \neg\Phi'$, then $[\Phi](s_1) = (0, 1)^2 \setminus [\Phi'](s_1)$ and, thus, $[\Phi](s_1)$ is a finite union of rectangles of the form $(0, 1) \times \langle a, b \rangle$. If

$\Phi = \Phi' \wedge \Phi''$, then $[\Phi](s_1) = [\Phi'](s_1) \cap [\Phi''](s_1)$ and, thus, $[\Phi](s_1)$ is a finite union of rectangles of the form $(0, 1) \times \langle a, b \rangle$.

If $\Phi = \mathcal{S}_{\geq \lambda}(\Phi')$, then first we observe that the steady-state distribution π of $\mathcal{M}[\mu, \mu']$ starting from s_1 is such that $\pi(s_1, s_2) = \mu'$ and $\pi(s_1, s_3) = 1 - \mu'$. Now, we distinguish different cases depending on whether $s_2 \models_{\text{CSL}} \Phi'$ and $s_3 \models_{\text{CSL}} \Phi'$. If both states satisfy Φ' , then $[\Phi](s_1) = (0, 1)^2$; if neither satisfies Φ' , then $[\Phi](s_1) = \emptyset$; if $s_2 \models_{\text{CSL}} \Phi'$ and $s_3 \not\models_{\text{CSL}} \Phi'$. Hence, $[\Phi](s_1) = (0, 1) \times [\lambda, 1]$; if $s_2 \not\models_{\text{CSL}} \Phi'$ and $s_3 \models_{\text{CSL}} \Phi'$, then $[\Phi](s_1) = (0, 1) \times (0, 1 - \lambda]$. The cases of $\Phi = \mathcal{S}_{< \lambda}(\Phi')$, with $\sim \in \{\leq, <, >\}$, follow similarly.

If $\Phi = \mathcal{P}_{\geq \lambda}(\mathcal{X}^{[\alpha, \beta]} \Phi')$, we distinguish different cases depending on whether $s_2 \models_{\text{CSL}} \Phi'$ and $s_3 \models_{\text{CSL}} \Phi'$. All of the cases are handled similarly and we only consider that in which $s_2 \models_{\text{CSL}} \Phi'$ and $s_3 \not\models_{\text{CSL}} \Phi'$. Then, $[\Phi](s_1) = \{(\mu, \mu') \in (0, 1)^2 \mid (e^{-\alpha} - e^{-\beta}) \cdot \mu' \geq \lambda\} = (0, 1) \times [\frac{\lambda - e^{-\beta}}{e^{-\alpha} - e^{-\beta}}, 1]$, which is of the required form. The cases of $\Phi = \mathcal{P}_{< \lambda}(\mathcal{X}^{[\alpha, \beta]} \Phi')$, with $\sim \in \{\leq, <, >\}$, follow similarly.

If $\Phi = \mathcal{P}_{\geq \lambda}(\Phi' \mathcal{U}^{[\alpha, \beta]} \Phi'')$, we make a case analysis with respect to the rectangles where the satisfaction of Φ' and Φ'' by s_1 is invariant (that is, we consider rectangles of the partition of $(0, 1)^2$ induced by the rectangles of $[\Phi'](s_1)$ and $[\Phi''](s_1)$). Our aim is to obtain $[\Phi](s_1)$ by replacing each such rectangle with a set of rectangles in which Φ is satisfied.

Given such a rectangle $\mathcal{R} \subseteq (0, 1)^2$ for which $\mathcal{M}[\mu, \mu'], s_1 \models_{\text{CSL}} \Phi''$ for all $(\mu, \mu') \in \mathcal{R}$, then $\mathcal{M}[\mu, \mu'], s_1 \models_{\text{CSL}} \Phi$ for all $(\mu, \mu') \in \mathcal{R}$. Hence, \mathcal{R} is included in $[\Phi](s_1)$. Conversely, if $\mathcal{M}[\mu, \mu'], s_1 \not\models_{\text{CSL}} \neg\Phi' \wedge \neg\Phi''$ for all $(\mu, \mu') \in \mathcal{R}$, then $\mathcal{M}[\mu, \mu'], s_1 \not\models_{\text{CSL}} \Phi$ for all $(\mu, \mu') \in \mathcal{R}$. Hence, no rectangle contained in \mathcal{R} is included in $[\Phi](s_1)$.

Now, consider a rectangle \mathcal{R} for which, for all $(\mu, \mu') \in \mathcal{R}$, we have $\mathcal{M}[\mu, \mu'], s_1 \models_{\text{CSL}} \Phi' \wedge \neg\Phi''$. Assume that $\mathcal{M}[\mu, \mu'], s_2 \models_{\text{CSL}} \Phi''$ and $\mathcal{M}[\mu, \mu'], s_3 \not\models_{\text{CSL}} \Phi''$ (the other cases are handled similarly). Then, we obtain $\{(\mu, \mu') \in \mathcal{R} \mid \mathcal{M}[\mu, \mu'], s_1 \models_{\text{CSL}} \Phi\} = \{(\mu, \mu') \in \mathcal{R} \mid (1 - e^{-\beta}) \cdot \mu' + e^{-\beta} \geq \lambda\} = \{(\mu, \mu') \in \mathcal{R} \mid \mu' \geq \frac{\lambda - e^{-\beta}}{1 - e^{-\beta}}\}$. Thus, we include in $[\Phi](s_1)$, the rectangle $\mathcal{R} \cap ((0, 1) \times [\frac{\lambda - e^{-\beta}}{1 - e^{-\beta}}, 1])$.

The cases of $\Phi = \mathcal{P}_{< \lambda}(\Phi' \mathcal{U}^{[\alpha, \beta]} \Phi'')$, $\sim \in \{\leq, <, >\}$, follow similarly.

Assertion 3. We now prove assertion 3 by induction on the size of the formulas. Let Φ be a formula of CSL. The cases in which Φ is an atomic proposition, $\Phi = \neg\Phi'$, $\Phi = \Phi' \wedge \Phi''$, and $\Phi = \mathcal{S}_{< \lambda}(\Phi')$ are proven exactly as for assertion 2 (the steady-state distribution of $\mathcal{M}[\mu, \mu']$ starting from s_0 is the same as that starting from s_1).

If $\Phi = \mathcal{P}_{\geq \lambda}(\mathcal{X}^{[\alpha, \beta]} \Phi')$, we make a case analysis with respect to the rectangles in which the satisfaction of Φ' by s_0 and s_1 is invariant (that is, we consider rectangles of the partition of $(0, 1)^2$ induced by the rectangles of $[\Phi'](s_0)$ and $[\Phi'](s_1)$). As above, we obtain $[\Phi](s_0)$ by replacing each such rectangle with a set of rectangles in which Φ is satisfied. We only consider one such case (the other cases are handled similarly). Consider a rectangle $\mathcal{R} \subseteq (0, 1)^2$ for which, for all $(\mu, \mu') \in \mathcal{R}$, we have $\mathcal{M}[\mu, \mu'], s_0 \models_{\text{CSL}} \Phi'$ and $\mathcal{M}[\mu, \mu'], s_1 \not\models_{\text{CSL}} \Phi'$. Then, $\{(\mu, \mu') \in \mathcal{R} \mid \mathcal{M}[\mu, \mu'], s_0 \models_{\text{CSL}} \Phi\} = \{(\mu, \mu') \in \mathcal{R} \mid (e^{-\alpha} / \mu - e^{-\beta/\mu}) \cdot \mu \geq \lambda\}$. Let $f(\mu) = (e^{-\alpha/\mu} - e^{-\beta/\mu}) \cdot \mu$. Note that the derivative of f




Fig. 11. A family of ASMCs $\mathcal{M}[\mu, \mu', \mu'']$ and a DTA \mathcal{A} of CSL^{TA}.

changes its sign only a finite number of times inside $(0, 1)$ (in fact in \mathbb{R}). Therefore, $(0, 1)$ may be decomposed into a finite number of consecutive intervals where inside an interval f is monotonic. As a consequence, $(0, 1)$ may be partitioned into a finite number of consecutive intervals (different from the previous ones) where alternatively f is greater than or equal to λ or strictly smaller than λ . The intervals for which f is greater than or equal to λ induce a finite number of rectangles of the form $\langle a, b \rangle \times (0, 1)$, which are included $[\Phi](s_0)$. The cases of $\Phi = \mathcal{P}_{\sim\lambda}(\mathcal{X}^{[\alpha, \beta]}\Phi')$, with $\sim \in \{\leq, <, >\}$, follow similarly.

If $\Phi = \mathcal{P}_{\geq\lambda}(\Phi' \mathcal{U}^{[\alpha, \beta]}\Phi'')$, we make a case analysis with respect to the rectangles where the satisfaction of Φ' and Φ'' by s_0 and by s_1 is invariant (that is, we consider rectangles of the partition of $(0, 1)^2$ induced by the rectangles of $[\Phi'](s_0)$, $[\Phi'](s_1)$, $[\Phi''](s_0)$, and $[\Phi''](s_1)$). Again, we obtain $[\Phi](s_0)$ by replacing each such rectangle with a set of rectangles in which Φ is satisfied.

We handle only one case, noting that the other cases are handled similarly. Consider the rectangle $\mathcal{R} \subseteq (0, 1)^2$ such that, for $i \in \{0, 1\}$, we have $\mathcal{M}[\mu, \mu', s_i] \models_{\text{CSL}} \Phi' \wedge \neg\Phi''$, $\mathcal{M}[\mu, \mu', s_2] \models_{\text{CSL}} \neg\Phi' \wedge \Phi''$ and $\mathcal{M}[\mu, \mu', s_3] \models_{\text{CSL}} \neg\Phi' \wedge \neg\Phi''$. The key observation here is that, inside any such rectangle, the loop around s_0 is irrelevant due to the nature of the Until operator. Then, we have

$$\begin{aligned} & \{(\mu, \mu') \in \mathcal{R} \mid \mathcal{M}[\mu, \mu', s_0] \models_{\text{CSL}} \Phi\} \\ &= \{(\mu, \mu') \in \mathcal{R} \mid (e^{-2\alpha}(1+2\alpha) - e^{-2\beta}(1+2\beta)) \cdot \mu' \geq \lambda\} \\ &= \left\{ (\mu, \mu') \in \mathcal{R} \mid \mu' \geq \frac{\lambda}{e^{-2\alpha}(1+2\alpha) - e^{-2\beta}(1+2\beta)} \right\} \end{aligned}$$

(the first formula has been obtained by applying an Erlang distribution). Then, we include the rectangle $\mathcal{R} \cap ((0, 1) \times [\frac{\lambda}{e^{-2\alpha}(1+2\alpha) - e^{-2\beta}(1+2\beta)}, 1])$ in $[\Phi](s_0)$.

The cases of $\Phi = \mathcal{P}_{\sim\lambda}(\Phi' \mathcal{U}^{[\alpha, \beta]}\Phi'')$, $\sim \in \{\leq, <, >\}$, follow similarly. \square

Because $[\Phi^\zeta](s_0) = \{(\mu, \mu') \mid \mu \cdot \mu' \geq \zeta\}$ cannot be expressed as a finite union of rectangles, Lemma 4.13 establishes that Φ^ζ is not equivalent to any formula of CSL. Lemma 4.14 then gives a direct proof of Proposition 4.12.

Lemma 4.14. *For each $0 < \zeta < 1$, the CSL^{TA} formula Φ^ζ is not equivalent to any formula of CSL.*

We also conjecture that there exists a CSL^{TA} formula for which no equivalent asCSL formula exists. This conjecture is based on the result shown in the next subsection, which shows that there exists a CSL^{TA} formula which does not use nesting for which there exists no single equivalent asCSL formula which does not use nesting.

4.4 CSL^{TA} without Nesting Is Strictly More Expressive than asCSL without Nesting

Consider the ASMC $\mathcal{M}[\mu, \mu', \mu'']$ shown on the left of Fig. 11. Given a CSL^{TA} formula Φ , let $[\Phi] = \{(\mu, \mu', \mu'') \in (0, 1)^3 \mid \mathcal{M}[\mu, \mu', \mu''] \models s_0 \models \Phi\}$; similarly, given an asCSL formula Φ' , let $[\Phi'] = \{(\mu, \mu', \mu'') \in (0, 1)^3 \mid \mathcal{M}[\mu, \mu', \mu''] \models_{\text{asCSL}} \Phi'\}$. (Note that, in contrast to the case of CSL in Section 4.3, we consider only the state s_0 of \mathcal{M} .) The DTA on the right of Fig. 11 will be used to define the CSL^{TA} formula for which there exists no equivalent asCSL formula without nesting.

Given that the action set Act of $\mathcal{M}[\mu, \mu', \mu'']$ is equal to $\{a\}$, the single state proposition of $\mathcal{M}[\mu, \mu', \mu'']$ is true, and that, for the purposes of the following results, we do not allow nesting within asCSL formulas, we have that the input alphabet Ξ of any considered \mathcal{N} is equal to the set $\{\langle \text{true}, a \rangle, \langle \text{true}, \sqrt{} \rangle\}$: For this reason, we write $z \xrightarrow{a} z'$ and $z \xrightarrow{\sqrt{}} z'$ for $z \xrightarrow{\text{true } a} z'$ and $z \xrightarrow{\text{true } \sqrt{}} z'$, respectively.

We identify the set of NPA referring to sequences of a -transitions of at most length 2: let 2NPA refer to the set of NPA for which the longest path from the set of initial states to the set of final states features at most two a -transitions. Let $\mathcal{N} = \langle Z, \Xi, \delta, Z_{\text{Init}}, Z_F \rangle$ be an NPA, let $\text{Runs}^{[M]} = \text{Runs}^{\mathcal{N}}(Z_{\text{Init}})$, and let $\text{Runs}_{\leq 2}^{[M]} = \{r \in \text{Runs}^{[M]} \mid r \text{ has at most two } a\text{-transitions}\}$ and let $\text{Runs}_{\geq 3}^{[M]} = \text{Runs}^{[M]} \setminus \text{Runs}_{\leq 2}^{[M]}$.

We can obtain an NPA $\mathcal{N}_{\leq 2}$ such that we have $\text{Runs}^{[M_{\leq 2}]} = \text{Runs}_{\leq 2}^{[M]}$. Informally, to obtain $\mathcal{N}_{\leq 2}$, we represent each state $z \in Z$ of \mathcal{N} by three copies z^0, z^1 , and z^2 in $\mathcal{N}_{\leq 2}$. The transition relation of $\mathcal{N}_{\leq 2}$ is defined such that a transition $z_1 \xrightarrow{\sqrt{}} z_2$ of \mathcal{N} is represented by $z_1^i \xrightarrow{\sqrt{}} z_2^i$ in $\mathcal{N}_{\leq 2}$, for $i \in \{0, 1, 2\}$. A transition $z_1 \xrightarrow{a} z_2$ of \mathcal{N} is represented in $\mathcal{N}_{\leq 2}$ by the transitions $z_1^i \xrightarrow{a} z_2^{i+1}$ for $i \in \{0, 1\}$ and by $z_1^2 \xrightarrow{a} z_{\text{sink}}$, where z_{sink} is an additional state without outgoing transitions. The set of initial states of $\mathcal{N}_{\leq 2}$ is $\{z^0 \mid z \in Z_{\text{Init}}\}$, and the set of final states of $\mathcal{N}_{\leq 2}$ is $\cup_{i \in \{0, 1, 2\}} \{z^i \mid z \in Z_F\}$. It can be verified that $\mathcal{N}_{\leq 2}$ is a 2NPA.

Similarly, we can obtain an NPA $\mathcal{N}_{\geq 3}$ such that, for any finite path σ of $\mathcal{M}[\mu, \mu', \mu'']$, the set $\text{Runs}^{[M_{\geq 3}]} = \text{Runs}_{\geq 3}^{[M]}$. To obtain $\mathcal{N}_{\geq 3}$, we represent each state $z \in Z$ of \mathcal{N} by three copies z^0, z^1 , and z^2 in $\mathcal{N}_{\geq 3}$. The transition relation of $\mathcal{N}_{\geq 3}$ is defined such that a transition $z_1 \xrightarrow{\sqrt{}} z_2$ of \mathcal{N} is represented by $z_1^i \xrightarrow{\sqrt{}} z_2^i$ in $\mathcal{N}_{\geq 3}$, for $i \in \{0, 1, 2\}$. A transition $z_1 \xrightarrow{a} z_2$ of \mathcal{N} is represented by $z_1^i \xrightarrow{a} z_2^{i+1}$ for $i \in \{0, 1\}$ and by $z_1^2 \xrightarrow{a} z_2^2$. The set of initial states of $\mathcal{N}_{\geq 3}$ is $\{z^0 \mid z \in Z_{\text{Init}}\}$, and the set of final states of $\mathcal{N}_{\geq 3}$ is $\{z^2 \mid z \in Z_F\}$.

For NPA \mathcal{N} and interval $I \subseteq \mathbb{R}_{\geq 0}$, we write $\text{Pr}_{s_0}^{\mathcal{M}[\mu, \mu', \mu'']}(I)$ for $\text{Pr}_{s_0}^{\mathcal{M}[\mu, \mu', \mu'']}(AccPath^{\mathcal{M}[\mu, \mu', \mu'']}(s_0, \mathcal{N}^I))$ when clear from the context.

Lemma 4.15. *Let $\langle \alpha, \beta \rangle \subseteq \mathbb{R}_{\geq 0}$ such that $\beta < \infty$ and let $\mu, \mu', \mu'' \in \mathbb{R}_{>0}$. Then, for any NPA \mathcal{N} , we have $\text{Pr}_{s_0}^{\mathcal{M}[\mu, \mu', \mu'']}(I_{\leq 2}^{\langle \alpha, \beta \rangle}) = \lim_{\nu \rightarrow 0} \text{Pr}_{s_0}^{\mathcal{M}[\mu, \mu', \mu'']}(I_{\leq 2}^{\langle \alpha, \beta \rangle})$.*

Proof. First, note that $AccPath^{\mathcal{M}[\mu, \mu', \mu'']}(s_0, \mathcal{N}^{\langle \alpha, \beta \rangle})$ equals

$$AccPath^{\mathcal{M}[\mu, \mu', \mu'']}(s_0, \mathcal{N}_{\leq 2}^{(\alpha, \beta)}) \cup AccPath^{\mathcal{M}[\mu, \mu', \mu'']}(s_0, \mathcal{N}_{\geq 3}^{(\alpha, \beta)}).$$

Hence,

$$\begin{aligned} & \Pr_{s_0}^{\mathcal{M}[\mu, \mu', \mu'']}(\mathcal{N}^{(\alpha, \beta)}) \\ & \leq \Pr_{s_0}^{\mathcal{M}[\mu, \mu', \mu'']}(\mathcal{N}_{\leq 2}^{(\alpha, \beta)}) + \Pr_{s_0}^{\mathcal{M}[\mu, \mu', \mu'']}(\mathcal{N}_{\geq 3}^{(\alpha, \beta)}). \end{aligned}$$

We then observe the following fact:

$$\begin{aligned} & \lim_{\nu \rightarrow 0} \Pr_{s_0}^{\mathcal{M}[\mu, \mu', \nu]}(\mathcal{N}_{\leq 2}^{(\alpha, \beta)}) \\ & \leq \lim_{\nu \rightarrow 0} \Pr_{s_0}^{\mathcal{M}[\mu, \mu', \nu]}(\mathcal{N}^{(\alpha, \beta)}), \\ & \leq \lim_{\nu \rightarrow 0} (\Pr_{s_0}^{\mathcal{M}[\mu, \mu', \nu]}(\mathcal{N}_{\leq 2}^{(\alpha, \beta)}) + \Pr_{s_0}^{\mathcal{M}[\mu, \mu', \nu]}(\mathcal{N}_{\geq 3}^{(\alpha, \beta)})). \end{aligned}$$

Now note that $\lim_{\nu \rightarrow 0} \Pr_{s_0}^{\mathcal{M}[\mu, \mu', \nu]}(\mathcal{N}_{\geq 3}^{(\alpha, \beta)}) = 0$ (from the fact that, as ν tends to 0, the probability that there exists a finite prefix of a path of $\mathcal{M}[\mu, \mu', \nu]$ comprising at least three a -transitions and with time duration in $\langle \alpha, \beta \rangle$ tends to 0). Hence, $\lim_{\nu \rightarrow 0} \Pr_{s_0}^{\mathcal{M}[\mu, \mu', \nu]}(\mathcal{N}_{\leq 2}^{(\alpha, \beta)}) = \lim_{\nu \rightarrow 0} \Pr_{s_0}^{\mathcal{M}[\mu, \mu', \nu]}(\mathcal{N}^{(\alpha, \beta)})$. From the fact that the value of μ'' is irrelevant to the probability of satisfying a property specified by a 2NPA, we have $\lim_{\nu \rightarrow 0} \Pr_{s_0}^{\mathcal{M}[\mu, \mu', \nu]}(\mathcal{N}_{\leq 2}^{(\alpha, \beta)}) = \Pr_{s_0}^{\mathcal{M}[\mu, \mu', \mu'']}(\mathcal{N}_{\leq 2}^{(\alpha, \beta)})$ for any $\mu'' \in \mathbb{R}_{>0}$. Hence $\Pr_{s_0}^{\mathcal{M}[\mu, \mu', \mu'']}(\mathcal{N}_{\leq 2}^{(\alpha, \beta)}) = \lim_{\nu \rightarrow 0} \Pr_{s_0}^{\mathcal{M}[\mu, \mu', \nu]}(\mathcal{N}^{(\alpha, \beta)})$. \square

Lemma 4.16. *Let $\langle \alpha, \infty \rangle \subseteq \mathbb{R}_{\geq 0}$ and let $\mu, \mu', \mu'' \in \mathbb{R}_{>0}$. Then, for any NPA \mathcal{N} , we have either $\lim_{\nu \rightarrow 0} \Pr_{s_0}^{\mathcal{M}[\mu, \mu', \nu]}(\mathcal{N}^{(\alpha, \infty)}) = 1$ or $\Pr_{s_0}^{\mathcal{M}[\mu, \mu', \mu'']}(\mathcal{N}_{\leq 2}^{(\alpha, \infty)}) = \Pr_{s_0}^{\mathcal{M}[\mu, \mu', \mu'']}(\mathcal{N}^{(\alpha, \infty)})$.*

Proof. First, note that, if $AccPath^{\mathcal{M}[\mu, \mu', \mu'']}(s_0, \mathcal{N}_{\geq 3}^{(\alpha, \infty)}) \neq \emptyset$, then $\lim_{\nu \rightarrow 0} \Pr_{s_0}^{\mathcal{M}[\mu, \mu', \nu]}(\mathcal{N}_{\geq 3}^{(\alpha, \infty)}) = 1$ from the fact that, as ν tends to 0, the probability that there exists a finite prefix of a path of $\mathcal{M}[\mu, \mu', \nu]$ comprising at least three a -transitions and with time duration in $\langle \alpha, \infty \rangle$ tends to 1. In this case, noting that

$$\begin{aligned} & \lim_{\nu \rightarrow 0} \Pr_{s_0}^{\mathcal{M}[\mu, \mu', \nu]}(\mathcal{N}_{\geq 3}^{(\alpha, \infty)}) \\ & \leq \lim_{\nu \rightarrow 0} \Pr_{s_0}^{\mathcal{M}[\mu, \mu', \nu]}(\mathcal{N}^{(\alpha, \infty)}), \end{aligned}$$

we conclude that $\lim_{\nu \rightarrow 0} \Pr_{s_0}^{\mathcal{M}[\mu, \mu', \nu]}(\mathcal{N}^{(\alpha, \infty)}) = 1$

If $AccPath^{\mathcal{M}[\mu, \mu', \mu'']}(s_0, \mathcal{N}_{\geq 3}^{(\alpha, \infty)}) = \emptyset$, then

$$\begin{aligned} & AccPath^{\mathcal{M}[\mu, \mu', \mu'']}(s_0, \mathcal{N}^{(\alpha, \infty)}) \\ & = AccPath^{\mathcal{M}[\mu, \mu', \mu'']}(s_0, \mathcal{N}_{\leq 2}^{(\alpha, \infty)}). \end{aligned}$$

Therefore, $\Pr_{s_0}^{\mathcal{M}[\mu, \mu', \mu'']}(\mathcal{N}_{\leq 2}^{(\alpha, \infty)}) = \Pr_{s_0}^{\mathcal{M}[\mu, \mu', \mu'']}(\mathcal{N}^{(\alpha, \infty)})$. \square

Hence, by Lemma 4.15 and Lemma 4.16 and by letting the rate of the self-loop labeling s_2 approach 0, it suffices to consider asCSL formulas which use 2NPA, rather than arbitrary NPA (noting that the case in which $\lim_{\nu \rightarrow 0} \Pr_{s_0}^{\mathcal{M}[\mu, \mu', \nu]}(\mathcal{N}^{(\alpha, \infty)}) = 1$ is not of interest).

Let \mathcal{A} be the DTA of Fig. 11. Then, consider the CSL^{TA} formula $\Phi^\zeta = \mathcal{P}_{\geq \zeta}(\mathcal{A})$ for some $\zeta \in (0, 1)$. It follows that $[\Phi^\zeta] = \{(\mu, \mu', \mu'') \in (0, 1)^3 \mid e^{-\mu} \cdot e^{-\mu'} \geq \zeta\}$. Rewriting, we obtain $[\Phi^\zeta] = \{(\mu, \mu', \mu'') \in (0, 1)^3 \mid \mu + \mu' \leq -\ln(\zeta)\}$.

Lemma 4.17. *Let $\Phi = \mathcal{P}_{\sim \lambda}(\mathcal{N}^{(\alpha, \beta)})$ be an asCSL formula for which \mathcal{N} is a 2NPA. Then, $[\Phi] = \{(\mu, \mu', \mu'') \in (0, 1)^3 \mid \mu \neq \mu' \wedge \Upsilon^\neq \sim \lambda\} \cup \{(\mu, \mu', \mu'') \in (0, 1)^3 \mid \mu = \mu' \wedge \Upsilon^= \sim \lambda\}$, where either:*

- Υ^\neq and $\Upsilon^=$ are both 0, or are both 1;
- Υ^\neq and $\Upsilon^=$ are both $e^{-\mu\alpha} - e^{-\mu\beta}$;
- Υ^\neq is of the form $\frac{\mu'(e^{-\mu\alpha} - e^{-\mu\beta}) - \mu(e^{-\mu'\alpha} - e^{-\mu'\beta})}{\mu' - \mu}$, and $\Upsilon^=$ is of the form $\mu(e^{-\mu\alpha} - e^{-\mu\beta})$;
- Υ^\neq is of the form $(e^{-\mu\alpha} - e^{-\mu\beta}) + \frac{\mu'(e^{-\mu\alpha} - e^{-\mu\beta}) - \mu(e^{-\mu'\alpha} - e^{-\mu'\beta})}{\mu' - \mu}(1 - (e^{-\mu\alpha} - e^{-\mu\beta}))$ and $\Upsilon^=$ is of the form $(e^{-\mu\alpha} - e^{-\mu\beta}) + (\mu(e^{-\mu\alpha} - e^{-\mu\beta}))(1 - (e^{-\mu\alpha} - e^{-\mu\beta}))$.

The case in which Υ^\neq or $\Upsilon^=$ is of the form $e^{-\mu\alpha} - e^{-\mu\beta}$ corresponds to the case in which \mathcal{N} accepts paths which have the duration of the first transition in $\langle \alpha, \beta \rangle$. The third point considers the case in which \mathcal{N} accepts paths which have the duration of the prefix consisting of the first two transitions in $\langle \alpha, \beta \rangle$ (where the expression is obtained by applying an Erlang distribution). The fourth point corresponds to the case in which \mathcal{N} accepts paths which have the duration of the first transition in $\langle \alpha, \beta \rangle$ (first summand) and paths which do not have the duration of the first transition in $\langle \alpha, \beta \rangle$ but have the duration of the prefix consisting of the first two transitions in $\langle \alpha, \beta \rangle$ (second summand).

Given the fact that it is not possible to represent the set $[\Phi^\zeta] = \{(\mu, \mu', \mu'') \in (0, 1)^3 \mid \mu + \mu' \leq -\ln(\zeta)\}$ using any of the expressions of Lemma 4.17, we have the following proposition.

Proposition 4.18. *There is a formula of CSL^{TA} without nesting for which there is no equivalent asCSL formula of the form $\mathcal{P}_{\sim \lambda}(\mathcal{N}^I)$ without nesting.*

5 CONCLUSION

In this paper, we have defined a new stochastic temporal logic CSL^{TA}, based on timed automata, which we propose as a good trade-off between adding flexibility to property specification and limiting the explosion of complexity in analysis. With regard to the specification of properties, the most significant extension is the possibility of specifying an arbitrary number of timing constraints along an execution path which may also depend on the history of the process. We have shown that CSL^{TA} is at least as expressive as both CSL and asCSL. Furthermore, the evaluation process is handled in an uniform way via Markov regenerative processes rather than by ad hoc transformations as previously. We note that the two restrictions that we have placed on the timed automata used, namely that they are deterministic and have one clock, allow us to obtain a tractable stochastic process for the joint process of the system and the property, namely, a Markov regenerative process, for which there exists well-known solution methods [26], [29].

Further work can consider an implementation of the proposed method (possibly exploiting existing Deterministic Stochastic Petri Net tools) and the extension of CSL^{TA} to allow for rewards [30]. We would also like to investigate the use of CSL^{TA} for the definition of properties of performance models generated automatically from the sequence diagrams of UML, where the ability of CSL^{TA} to reason about concatenated time intervals could be of use.

ACKNOWLEDGMENTS

The work of S. Haddad was supported in part by the project ANR-06-SETI-002 CheckBound. The work of S. Donatelli and J. Sproston was supported in part by the EEC project 027513 Crucial. The authors would like to thank the anonymous reviewers for the many useful and detailed suggestions and corrections.

REFERENCES

- [1] C. Smith, *Performance Engineering of Software Systems*. Addison-Wesley, 1990.
- [2] *UML Profile for Schedulability, Performance and Time Specification*, Object Management Group, version 1.1, formal/05-01-02, Jan. 2005.
- [3] *A UML Profile for Modeling and Analysis of Real Time Embedded Systems (MARTE)*, OMG, revised submission, 2007.
- [4] V. Cortellessa and R. Mirandola, "Deriving a Queuing Network Based Performance Model from UML Diagrams," *Proc. ACM Workshop Software and Performance*, pp. 58-70, Sept. 2000.
- [5] S. Bernardi and J. Merseguer, "Performance Evaluation of Uml Design with Stochastic Well-Formed Nets," *J. Systems and Software*, vol. 80, no. 11, pp. 1843-1865, Nov. 2007.
- [6] C. Canevet, S. Gilmore, J. Hillston, M. Prowse, and P. Stevens, "Performance Modelling with UML and Stochastic Process Algebras," *IEEE Proc.: Computers and Digital Techniques*, vol. 150, no. 2, pp. 107-120, Mar. 2003.
- [7] S. Balsamo, A. Di Marco, P. Inverardi, and M. Simeoni, "Model-Based Performance Prediction in Software Development: A Survey," *IEEE Trans. Software Eng.*, vol. 30, no. 5, pp. 295-310, May 2004.
- [8] A. Aziz, K. Sanwal, V. Singhal, and R. Brayton, "Model-Checking Continuous Time Markov Chains," *ACM Trans. Computational Logic*, vol. 1, no. 1, pp. 162-170, 2000.
- [9] C. Baier, B. Haverkort, H. Hermanns, and J.-P. Katoen, "Model-Checking Algorithms for Continuous-Time Markov Chains," *IEEE Trans. Software Eng.*, vol. 29, no. 6, pp. 524-541, June 2003.
- [10] C. Baier, L. Cloth, B. Haverkort, M. Kuntz, M. Siegle, "Model Checking Action- and State-Labelled Markov Chains," *IEEE Trans. Software Eng.*, vol. 33, no. 4, pp. 209-224, Apr. 2007.
- [11] M.Y. Vardi and P. Wolper, "Reasoning About Infinite Computations," *Information and Computation*, vol. 115, no. 1, pp. 1-37, 1994.
- [12] E.M. Clarke, O. Grumberg, and R.P. Kurshan, "A Synthesis of Two Approaches for Verifying Finite State Concurrent Systems," *J. Logic Computation*, vol. 2, no. 5, pp. 605-618, 1992.
- [13] R. Alur and D.L. Dill, "A Theory of Timed Automata," *Theoretical Computer Science*, vol. 126, no. 2, pp. 183-235, 1994.
- [14] G. Behrmann, A. David, K.G. Larsen, J. Håkansson, P. Pettersson, W. Yi, and M. Hendriks, "UPPAAL 4.0," *Proc. IEEE Quantitative Evaluation of Systems*, pp. 125-126.
- [15] R. Alur, C. Courcoubetis, and D.L. Dill, "Model-Checking in Dense Real-Time," *Information and Computation*, vol. 104, no. 1, pp. 2-34, 1993.
- [16] R. Alur, T. Feder, and T.A. Henzinger, "The Benefits of Relaxing Punctuality," *J. ACM*, vol. 43, no. 1, pp. 116-146, 1996.
- [17] E.A. Emerson and J.Y. Halpern, "Sometimes" and "Not Never" Revisited: On Branching Versus Linear Time Temporal Logic," *J. ACM*, vol. 33, no. 1, pp. 151-178, 1986.
- [18] G. Clark and J. Hillston, "Towards Automatic Derivation of Performance Measures from PEPA Models," *Proc. UK Performance Eng. Workshop*, 1996.
- [19] W.D. Obal II and W.H. Sanders, "State-Space Support for Path-Based Reward Variables," *Performance Evaluation*, vol. 35, nos. 3/4, pp. 233-251, 1999.
- [20] A. Bouajjani, Y. Lakhnech, and S. Yovine, "Model-Checking for Extended Timed Temporal Logics," *Proc. Formal Techniques in Real-Time and Fault-Tolerant Systems*, pp. 306-325, 1996.
- [21] F. Laroussinie, N. Markey, and P. Schnoebelen, "Model Checking Timed Automata with One or Two Clocks," *Proc. Conf. Concurrency Theory*, pp. 387-401, 2004.
- [22] J. Ouaknine and J. Worrell, "On the Language Inclusion Problem for Timed Automata: Closing a Decidability Gap," *Proc. IEEE Logic in Computer Science*, pp. 54-63, 2004.

- [23] S. Donatelli, S. Haddad, and J. Sproston, "CSL^{TA}: An Expressive Logic for Continuous-Time Markov Chains," *Proc. IEEE Quantitative Evaluation of Systems*, pp. 31-40, 2007.
- [24] S. Yovine, "KRONOS: A Verification Tool for Real-Time Systems," *Software Tools for Technology Transfer*, vol. 1, nos. 1/2, pp. 123-133, 1997.
- [25] E.M. Clarke, O. Grumberg, and D.A. Peled, *Model Checking*. MIT Press, 1999.
- [26] R. German, *Performance Analysis of Communication Systems with Non-Markovian Stochastic Petri Nets*. Wiley, 2000.
- [27] M. Ajmone Marsan and G. Chiola, "On Petri Nets with Deterministic and Exponentially Distributed Firing Times," *Proc. Int'l Conf. Application and Theory of Petri Nets and Other Models of Concurrency*, pp. 132-145, 1986.
- [28] A. Jensen, "Markov Chains as An aid in the Study of Markov Processes," *Skandinavisk Aktuarietidskrift*, vol. 36, pp. 87-91, 1953.
- [29] C. Lindemann, *Performance Modelling with Deterministic and Stochastic Petri Nets*. Wiley, 1998.
- [30] B. Haverkort, L. Cloth, H. Hermanns, J.-P. Katoen, and C. Baier, "Model Checking Performability Properties," *Proc. IEEE Int'l Conf. Dependable Systems and Networks*, pp. 103-112, 2002.



Susanna Donatelli received the master's degree in electrical and computer engineering from the University of Massachusetts, Amherst, in 1987 and the PhD degree in computer science from the University of Turin, Italy, in 1989. She is currently a full professor of computer science at the University of Turin. Her research has focused on performance evaluation and probabilistic verification of discrete event systems based on queuing networks, stochastic Petri nets, high-level Petri nets, and stochastic process algebras. Her current research interests mainly include modeling and evaluation of interdependencies in large critical infrastructures, as well as on stochastic Petri Nets and stochastic model checking. She is a member of the steering committees of the Petri nets community and of the International Conference on Quantitative Evaluation of Systems. She is a member of the IEEE, the IEEE Computer Society, and the ACM and she is currently serving as associate editor for the *IEEE Transactions on Software Engineering*.



Serge Haddad is a former student at the École Normale Supérieure de Cachan. He received the MSc degree in mathematics in 1977 from the University of Orsay and the MSc and PhD degrees in computer science in 1983 and 1987, respectively, from the University of Paris 6. He is currently a full professor at the École Normale Supérieure de Cachan. His research interests include quantitative verification with emphasis on timed and stochastic systems and applications to software engineering.



Jeremy Sproston received the bachelor's degree in mathematical economics in 1995, the master's degree in computer science in 1996, and the PhD degree in computer science in 2001, all from the University of Birmingham, United Kingdom. He has held postdoctoral research positions at the University of Birmingham and the University of Turin, Italy. Since 2002, he has held a researcher position at the University of Turin, during which time he has been an honorary research fellow at the University of Birmingham and a visiting researcher at the Laboratoire Spécification et Vérification, École Normale Supérieure de Cachan, France. His research interests include the use of formal methods to verify the correctness and reliability of computer systems, with emphasis on model-checking techniques for probabilistic and timed systems.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.