

MODEL CHECKING OF CONTINUOUS-TIME MARKOV CHAINS AGAINST TIMED AUTOMATA SPECIFICATIONS

TAOLUE CHEN^a, TINGTING HAN^b, JOOST-PIETER KATOEN^c, AND ALEXANDRU MEREACRE^d

^a Formal Methods and Tools, University of Twente, The Netherlands
e-mail address: chent@ewi.utwente.nl

^{b,d} Software Modelling and Verification, RWTH Aachen University, Germany
e-mail address: {tingting.han,mereacre}@cs.rwth-aachen.de

^c Software Modelling and Verification, RWTH Aachen University, Germany;
Formal Methods and Tools, University of Twente, The Netherlands
e-mail address: katoen@cs.rwth-aachen.de

ABSTRACT. We study the verification of a finite continuous-time Markov chain (CTMC) \mathcal{C} against a linear real-time specification given as a deterministic timed automaton (DTA) \mathcal{A} with finite or Muller acceptance conditions. The central question that we address is: what is the probability of the set of paths of \mathcal{C} that are accepted by \mathcal{A} , i.e., the likelihood that \mathcal{C} satisfies \mathcal{A} ? It is shown that under finite acceptance criteria this equals the reachability probability in a finite piecewise deterministic Markov process (PDP), whereas for Muller acceptance criteria it coincides with the reachability probability of terminal strongly connected components in such a PDP. Qualitative verification is shown to amount to a graph analysis of the PDP. Reachability probabilities in our PDPs are then characterized as the least solution of a system of Volterra integral equations of the second type and are shown to be approximated by the solution of a system of partial differential equations. For single-clock DTA, this integral equation system can be transformed into a system of linear equations where the coefficients are solutions of ordinary differential equations. As the coefficients are in fact transient probabilities in CTMCs, this result implies that standard algorithms for CTMC analysis suffice to verify single-clock DTA specifications.

1998 ACM Subject Classification: D.2.4.

Key words and phrases: continuous-time Markov chains, deterministic timed automata, linear-time specification, model checking, piecewise-deterministic Markov processes.

^a This research is funded by the DFG research training group 1295 AlgoSyn, the SRO DSN project of CTIT, University of Twente, the EU FP7 project QUASIMODO and the DFG-NWO ROCKS project.

1. INTRODUCTION

Continuous-time Markov chains (CTMCs) are one of the most prominent models in performance and dependability analysis. They are exploited in a broad range of applications, and constitute the underlying semantical model of a plethora of modeling formalisms for real-time probabilistic systems such as Markovian queueing networks, stochastic Petri nets, stochastic variants of process algebras, and calculi for systems biology. CTMC model checking has been mainly focused on the branching-time temporal logic CSL (Continuous Stochastic Logic [3, 7]), a variant of timed CTL where the CTL universal and existential path quantifiers are replaced by a probabilistic operator. Like CTL model checking, CSL model checking of finite CTMCs proceeds by a recursive descent over the parse tree of the CSL formula. One of the key ingredients is that time-bounded reachability probabilities can be approximated arbitrarily closely by a reduction to transient analysis in CTMCs [7]. This results in an efficient polynomial-time algorithm that has been realized in model-checking tools such as PRISM [19] and MRMC [20] and has been successfully applied to various case studies from diverse application areas.

Verifying a finite CTMC \mathcal{C} against linear-time (but untimed) specifications in the form of a regular or ω -regular language is rather straightforward and boils down to computing reachability probabilities in discrete-time Markov chains (DTMCs). This can be seen as follows. Assume that the specification is provided as a deterministic automaton \mathcal{A} on finite words, or alternatively as a deterministic Muller automaton \mathcal{A} . The underlying idea is that the evolution of a CTMC is “synchronized” with an accepting run of \mathcal{A} by considering the state labels in a CTMC, i.e., atomic propositions, as letters read by \mathcal{A} . As \mathcal{A} does not constrain the timing of events in the CTMC \mathcal{C} , it suffices to take a synchronous product of \mathcal{A} and \mathcal{C} ’s embedded DTMC, denoted $emb(\mathcal{C})$, which is obtained by just ignoring the random state residence times in \mathcal{C} while keeping all other ingredients, in particular the transition probabilities and state labels. For finite acceptance criteria, the probability that $\mathcal{C} \models \mathcal{A}$, i.e., the probability of the set of paths in \mathcal{C} that are accepted by \mathcal{A} , $\Pr(\mathcal{C} \models \mathcal{A})$ for short, is obtained as the reachability probability in the product $emb(\mathcal{C}) \otimes \mathcal{A}$ of the final states in \mathcal{A} . Since \mathcal{A} is deterministic, $emb(\mathcal{C}) \otimes \mathcal{A}$ is a DTMC. In case of Muller acceptance criteria, $\Pr(\mathcal{C} \models \mathcal{A})$ corresponds to the reachability probability of accepting terminal strongly connected components in $emb(\mathcal{C}) \otimes \mathcal{A}$. This follows directly from results in [14]. The reachability probabilities in a DTMC can be obtained by solving a system of linear equations whose size is linear in the size of the DTMC, see, e.g., [18].

In this paper, we consider the verification of CTMCs against linear *real-time* specifications that are given as deterministic *timed* automata (DTA) [1]. That is to say, we explore the following problem: given a CTMC \mathcal{C} , and a linear real-time specification provided as a *deterministic timed automaton* \mathcal{A} , what is the probability of the set of paths of \mathcal{C} that are accepted by \mathcal{A} , i.e., what is $\Pr(\mathcal{C} \models \mathcal{A})$?

Example 1.1. Let us illustrate the usage of DTA specifications by means of a small example. Consider a robot randomly moving in some area. It starts in some zone (A , say) and has to reach zone B within 10 time units, cf. Figure 1(a). (For simplicity, all zones on the map are equally-sized, but this is not a restriction.) The robot randomly moves through the zones, and resides in a zone for an exponentially distributed amount of time. The robot may pass through all zones to reach B , but should not stay longer than 2 time units in any gray zone. The specification “reach B from A within 10 time units while residing in any gray zone for at most 2 time units” is modeled by a simple DTA which accepts once location

q_2 is reached, cf. Figure 1(b). Clock x controls the timing constraint on the residence times of the gray zones (assumed to be labeled with g), while clock y controls the global time constraint to reach zone B . In state q_0 , the robot traverses non-gray zones, in q_1 gray zones, and in q_2 it has reached the goal zone B .

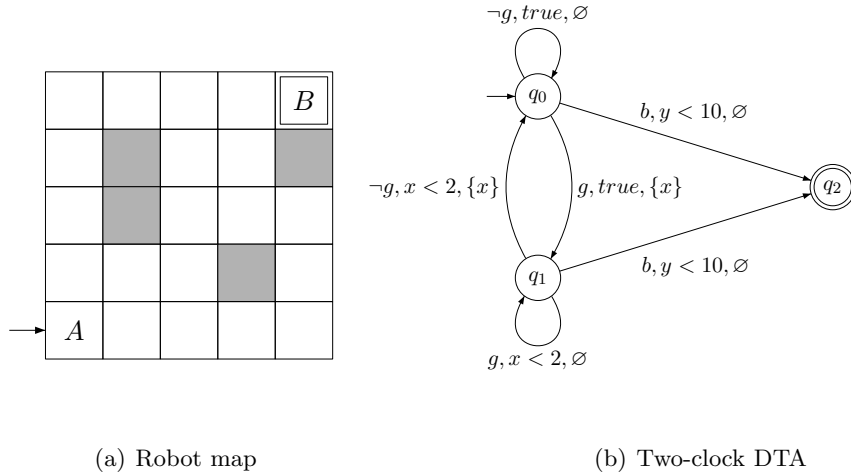


Figure 1: A robot example

Like in the untimed setting discussed before, we consider two variants: DTA that accept finite timed words, and DTA that accept infinite timed words according to a Muller acceptance condition. (Note that DTA with Muller acceptance condition are strictly more expressive than DTA with Büchi acceptance conditions [1].) The considered verification problem is substantially harder than the case for untimed linear specifications, e.g., as the DTA may constrain the timing of events in \mathcal{C} , it does not suffice to take the embedded DTMC $emb(\mathcal{C})$ as starting-point. In addition, the product of a CTMC and a DTA is neither a CTMC nor a DTA, and has an infinite state space. It is unclear which (and whether a) stochastic process is obtained from such infinite product, and if so, how to analyze it.

We tackle the verification of a finite CTMC against a DTA specification as follows:

- (1) We first show that the problem $\mathcal{C} \models \mathcal{A}$ is well-defined in the sense that the set of paths of \mathcal{C} that are accepted by \mathcal{A} is measurable.
- (2) We define the product $\mathcal{C} \otimes \mathcal{A}$ for CTMC \mathcal{C} and DTA \mathcal{A} as a variant of DTA in which, besides the usual ingredients of timed automata like guards and clock resets, the location residence time is exponentially distributed, and define a probability space over sets of timed paths in this model. In particular, we show that the probability of $\mathcal{C} \models \mathcal{A}$ coincides with the reachability probability of accepting paths in $\mathcal{C} \otimes \mathcal{A}$.
- (3) We adapt the standard region construction for timed automata [1] to this variant of DTA, and show that the thus obtained region automata are in fact *piecewise deterministic Markov processes* (PDPs) [16], a model that is frequently used in, e.g., stochastic control theory and financial mathematics. The characterization of region automata as PDPs sets the ground for obtaining the following results concerning qualitative and quantitative verification of CTMCs against DTA.
- (4) For finite acceptance criteria, we show that $\Pr(\mathcal{C} \models \mathcal{A})$ equals the reachability probability in the embedded PDP of $\mathcal{C} \otimes \mathcal{A}$. Under Muller acceptance criteria, $\Pr(\mathcal{C} \models \mathcal{A})$

equals the reachability probability of accepting terminal strongly connected components in this embedded PDP. In case of qualitative verification —does CTMC \mathcal{C} satisfy \mathcal{A} with probability larger than zero, or equal to one?— a graph traversal of the (embedded) PDP suffices.

- (5) We then show that reachability probabilities in our PDPs can be characterized as the least solution of a system of *Volterra integral equations* of the second type [2]. This probability is shown to be approximated by the solution of a system of partial differential equations (PDEs).
- (6) For the case of single-clock DTA, we show that the system of integral equations can be transformed into a system of *linear equations*, whose coefficients are solutions of some ordinary differential equations (ODEs). For these coefficients either an analytical solution (for small state space) can be obtained or an arbitrarily closely approximated solution can be determined efficiently.

Related work. Model checking CTMCs against linear real-time specifications has received scant attention so far. To our knowledge, this issue has only been (partly) addressed in [17, 6]. Baier et al. [6] define the logic asCSL where path properties are characterized by (time-bounded) regular expressions over actions and state formulas. The truth value of path formulas depends not only on the available actions in a given time interval, but also on the validity of certain state formulas in intermediate states. asCSL is strictly more expressive than CSL [6]. Model checking asCSL is performed by representing the regular expressions as finite-state automata, followed by computing time-bounded reachability probabilities in the product of CTMC \mathcal{C} and this automaton. In CSL^{TA} [17], time constraints of until modalities are specified by single-clock DTA; the resulting logic is at least as expressive as asCSL [17]. The combined behavior of \mathcal{C} and DTA \mathcal{A} is interpreted as a Markov renewal process and model checking CSL^{TA} is reduced to computing reachability probabilities in a DTMC whose transition probabilities are given by subordinate CTMCs. This paper takes a completely different approach. The technique of [17] cannot be generalized to multiple clocks, whereas our approach does not restrict the number of clocks and thus supports more specifications than CSL^{TA}. The DTA specification of our robot example, for instance, can neither be expressed in CSL^{TA} nor in asCSL. For the single-clock case, our approach produces the same result as [17], but yields a (in our opinion) conceptually simpler formulation whose correctness can be derived by simplifying the system of integral equations obtained for the general case. Moreover, measurability has not been addressed in [17]. Other related work [4, 5, 10] provides a quantitative interpretation to timed automata where delays and discrete choices are interpreted probabilistically. In this approach, delays of unbounded clocks are governed by exponential distributions like in CTMCs. Decidability results have been obtained for almost-sure properties [5] and quantitative verification [10] for (a subclass of) single-clock timed automata.

Organization of the paper. Section 2 defines the three models that are central to this paper: CTMCs, DTA, and PDPs. Section 3 shows that the set of paths in CTMC \mathcal{C} accepted by DTA \mathcal{A} is measurable and coincides with reachability probabilities in the product $\mathcal{C} \otimes \mathcal{A}$. It also shows that the underlying region graph of $\mathcal{C} \otimes \mathcal{A}$ is a (simple instance of a) PDP. Section 4 constitutes the main part of the paper and deals with the verification of DTA with finite acceptance conditions, and analyzes the quantitative reachability problem in our

PDPs, for both the general case and single-clock DTA. Section 5 considers DTA with Muller acceptance criteria, as well as qualitative verification. Finally, section 6 concludes.

This paper extends the conference paper [11] with complete proofs, illustrative examples, and by considering Muller acceptance criteria.

2. PRELIMINARIES

Given a set H , let $\Pr : \mathcal{F}(H) \rightarrow [0, 1]$ be a probability measure on the measurable space $(H, \mathcal{F}(H))$, where $\mathcal{F}(H)$ is a σ -algebra over H . Let $\text{Distr}(H)$ denote the set of probability measures on this measurable space.

2.1. Continuous-time Markov chains.

Definition 2.1 (CTMC). A (labeled) *continuous-time Markov chain* (CTMC) is a tuple $\mathcal{C} = (S, \text{AP}, L, \alpha, \mathbf{P}, E)$ where S is a *finite* set of *states*; AP is a finite set of *atomic propositions*; $L : S \rightarrow 2^{\text{AP}}$ is the *labeling function*; $\alpha \in \text{Distr}(S)$ is the *initial distribution*; $\mathbf{P} : S \times S \rightarrow [0, 1]$ is a stochastic *transition probability matrix*; and $E : S \rightarrow \mathbb{R}_{\geq 0}$ is the *exit rate function*.

The probability to exit state s in t time units is given by $\int_0^t E(s) \cdot e^{-E(s)\tau} d\tau$; the probability to take the transition $s \rightarrow s'$ in t time units equals $\mathbf{P}(s, s') \cdot \int_0^t E(s) e^{-E(s)\tau} d\tau$. A state s is *absorbing* if $\mathbf{P}(s, s) = 1$. The *embedded* discrete-time Markov chain (DTMC) of CTMC \mathcal{C} is obtained by deleting the exit rate function E , i.e., $\text{emb}(\mathcal{C}) = (S, \text{AP}, L, \alpha, \mathbf{P})$.

Definition 2.2 (Timed paths). Let \mathcal{C} be a CTMC. $\text{Paths}_n^{\mathcal{C}} := S \times (\mathbb{R}_{>0} \times S)^n$ is the set of paths of length n in \mathcal{C} ; the set of finite paths in \mathcal{C} is defined by $\text{Paths}_*^{\mathcal{C}} = \bigcup_{n \in \mathbb{N}} \text{Paths}_n^{\mathcal{C}}$ and $\text{Paths}_\omega^{\mathcal{C}} := (S \times \mathbb{R}_{>0})^\omega$ is the set of infinite paths in \mathcal{C} . $\text{Paths}^{\mathcal{C}} = \text{Paths}_*^{\mathcal{C}} \cup \text{Paths}_\omega^{\mathcal{C}}$ denotes the set of all paths in \mathcal{C} .

We denote a path $\rho \in \text{Paths}^{\mathcal{C}}(s_0)$ ($\rho \in \text{Paths}(s_0)$ for short) as the sequence $\rho = s_0 \xrightarrow{t_0} s_1 \xrightarrow{t_1} s_2 \cdots$ starting in state s_0 such that for $n \leq |\rho|$ ($|\rho|$ is the number of transitions in ρ if ρ is finite); $\rho[n] := s_n$ is the n -th state of ρ and $\rho\langle n \rangle := t_n$ is the time spent in state s_n . Let $\rho@t$ be the state occupied in ρ at time $t \in \mathbb{R}_{\geq 0}$, i.e. $\rho@t := \rho[n]$ where n is the smallest index such that $\sum_{i=0}^n \rho\langle i \rangle > t$. We assume w.l.o.g. $t_i > 0$ for any i .

The definition of a Borel space on paths through CTMCs follows [25, 7]. A CTMC \mathcal{C} yields a probability measure $\Pr^{\mathcal{C}}$ on paths as follows. Let $s_0, \dots, s_k \in S$ with $\mathbf{P}(s_i, s_{i+1}) > 0$ for $0 \leq i < k$ and I_0, \dots, I_{k-1} nonempty intervals in $\mathbb{R}_{\geq 0}$. Let $C(s_0, I_0, \dots, I_{k-1}, s_k)$ denote the *cylinder set* consisting of all paths $\rho \in \text{Paths}(s_0)$ such that $\rho[i] = s_i$ ($i \leq k$), and $\rho\langle i \rangle \in I_i$ ($i < k$). $\mathcal{F}(\text{Paths}(s_0))$ is the smallest σ -algebra on $\text{Paths}(s_0)$ which contains all sets $C(s_0, I_0, \dots, I_{k-1}, s_k)$ for all state sequences $(s_0, \dots, s_k) \in S^{k+1}$ with $\mathbf{P}(s_i, s_{i+1}) > 0$ ($0 \leq i < k$) and I_0, \dots, I_{k-1} range over all sequences of nonempty intervals in $\mathbb{R}_{\geq 0}$. The probability measure $\Pr^{\mathcal{C}}$ on $\mathcal{F}(\text{Paths}(s_0))$ is the unique measure defined by induction on k by $\Pr^{\mathcal{C}}(C(s_0)) = \alpha(s_0)$ and for $k > 0$:

$$\begin{aligned} \Pr^{\mathcal{C}}(C(s_0, I_0, \dots, I_{k-1}, s_k)) &= \Pr^{\mathcal{C}}(C(s_0, I_0, \dots, I_{k-2}, s_{k-1})) \\ &\cdot \int_{I_{k-1}} \mathbf{P}(s_{k-1}, s_k) E(s_{k-1}) \cdot e^{-E(s_{k-1})\tau} d\tau. \end{aligned} \quad (2.1)$$




Figure 2: An example CTMC

Example 2.3. An example CTMC is illustrated in Figure 2, where $\text{AP} = \{a, b, c\}$ and s_0 is the initial state, i.e., $\alpha(s_0) = 1$ and $\alpha(s) = 0$ for any $s \neq s_0$. The exit rates are indicated at the states, whereas the transition probabilities are attached to the transitions. An example timed path is $\rho = s_0 \xrightarrow{2.5} s_1 \xrightarrow{1.4} s_0 \xrightarrow{2} s_1 \xrightarrow{2\pi} s_2 \cdots$ with $\rho[2] = s_0$ and $\rho@6 = \rho[3] = s_1$.

2.2. Deterministic timed automata. Let $\mathcal{X} = \{x_1, \dots, x_n\}$ be a set of *nonnegative* real-valued variables, called *clocks*. An \mathcal{X} -valuation is a function $\eta : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ assigning to each variable x a nonnegative real value $\eta(x)$. Let $\mathcal{V}(\mathcal{X})$ denote the set of all valuations over \mathcal{X} . A *clock constraint* on \mathcal{X} , denoted by g , is a conjunction of expressions of the form $x \bowtie c$ for clock $x \in \mathcal{X}$, comparison operator $\bowtie \in \{<, \leq, >, \geq\}$ and $c \in \mathbb{N}$. Let $\mathcal{CC}(\mathcal{X})$ denote the set of clock constraints over \mathcal{X} . An \mathcal{X} -valuation η *satisfies* constraint $x \bowtie c$, denoted $\eta \models x \bowtie c$, if and only if $\eta(x) \bowtie c$; it satisfies a conjunction of such expressions if and only if η satisfies all of them. Let $\vec{0}$ denote the valuation that assigns 0 to all clocks. For a subset $X \subseteq \mathcal{X}$, the *reset* of X , denoted $\eta[X := 0]$, is the valuation η' such that $\forall x \in X. \eta'(x) := 0$ and $\forall x \notin X. \eta'(x) := \eta(x)$. For $\delta \in \mathbb{R}_{\geq 0}$ and \mathcal{X} -valuation η , $\eta + \delta$ is the \mathcal{X} -valuation η'' such that $\forall x \in \mathcal{X}. \eta''(x) := \eta(x) + \delta$, which implies that all clocks proceed at the same speed.

Definition 2.4 (DTA). A *deterministic timed automaton* (or DTA for short) is a tuple $\mathcal{A} = (\Sigma, \mathcal{X}, Q, q_0, Q_{\mathbf{F}}, \rightarrow)$ where Σ is a finite *alphabet*; \mathcal{X} is a finite set of *clocks*; Q is a nonempty, finite set of *locations* with *initial location* $q_0 \in Q$; $Q_{\mathbf{F}}$ is the *acceptance condition*, which is either:

- $Q_{\mathbf{F}} \subseteq Q$, a set of *accepting locations* (reachability or finite acceptance), or
- $Q_{\mathbf{F}} \subseteq 2^Q$, an *acceptance family* (Muller acceptance).

The relation $\rightarrow \subseteq Q \times \Sigma \times \mathcal{CC}(\mathcal{X}) \times 2^{\mathcal{X}} \times Q$ is the *edge relation* satisfying:

$$(q \xrightarrow{a, g, X} q' \text{ and } q \xrightarrow{a, g', X'} q'' \text{ with } g \neq g') \text{ implies } g \cap g' = \emptyset.$$

We refer to $q \xrightarrow{a, g, X} q'$ as an *edge*, where $a \in \Sigma$ is an input symbol, the *guard* g is a clock constraint on the clocks of \mathcal{A} , X is the set of clocks that are to be reset and q' is the successor location. Intuitively, the edge $q \xrightarrow{a, g, X} q'$ asserts that the DTA \mathcal{A} can move from location q to q' when the input symbol is a and the guard g holds, while the clocks in X should be reset when entering q' . DTA are deterministic as they have a single initial location, and outgoing edges of a location labeled with the same input symbol are required to have disjoint guards. In this way, the next location is uniquely determined for a given location and a given clock valuation. In case no guard is satisfied in a location for a given clock valuation, time can progress. If the advance of time will never reach a situation in which a guard holds, the DTA will stay in that location ad infinitum. Note that DTA do not have location invariants, as in safety timed automata. For the sake of simplicity, diagonal




Figure 3: DTA with (a) reachability and (b) Muller acceptance conditions

constraints like $x - y \bowtie c$ are not considered. This restriction does, however, not harm the expressiveness [9].

An (infinite) *timed path* of DTA \mathcal{A} is of the form $\theta = q_0 \xrightarrow{a_0, t_0} q_1 \xrightarrow{a_1, t_1} \dots$ such that $\eta_0 = \vec{0}$, and for all $j \geq 0$, it holds $t_j > 0$, $\eta_j + t_j \models g_j$, $\eta_{j+1} = (\eta_j + t_j)[X_j := 0]$, where η_j is the clock evaluation when *entering* q_j . The definitions on timed paths (such as $\theta[i]$, $\theta@t$, and so forth) for CTMCS can readily be adapted for DTA. We consider DTA with two types of acceptance criteria. Let DTA^\diamond and DTA^ω denote the set of DTA with reachability and Muller acceptance conditions, respectively. DTA denotes the general case covering both DTA^\diamond and DTA^ω .

Definition 2.5 (DTA accepting paths). An infinite timed path θ is *accepted* by a DTA^\diamond if $\theta[i] \in Q_F$ for some $i \geq 0$; θ is accepted by a DTA^ω if $\text{inf}(\theta) \in Q_{\mathcal{F}}$, where $\text{inf}(\theta)$ is the set of states $q \in Q$ such that $q = q_i$ for infinitely many $i \geq 0$.

The timed path θ is accepted according to a reachability criterion if it reaches some final location, whereas it is accepted according to a Muller acceptance condition if the set of infinitely visited locations equals some set in $Q_{\mathcal{F}}$. As a convention, we assume each location $q \in Q_F$ in DTA^\diamond to be a sink.

Example 2.6. Figure 3(a) depicts an example DTA^\diamond over the alphabet $\{a, b\}$ with initial location q_0 . The timed automaton is deterministic as q_0 is the only initial location and both a -labeled edges have disjoint guards. Any timed path ending in $Q_F = \{q_1\}$ is accepting.

Figure 3(b) depicts an example DTA^ω over the alphabet $\{a, b, c\}$. Its initial location is q_0 ; its Muller acceptance family equals $Q_{\mathcal{F}} = \{\{q_0, q_2\}\}$. Any accepting path should cycle between the locations q_0 and q_1 *finitely* often, and between q_0 and q_2 *infinitely* often.

Remark 2.7. [Expressive power of DTA^ω] DTA^ω is the set of (deterministic) timed Muller automata, (*D*)*MTA*, for short. A (deterministic) timed Büchi automaton, (*D*)*TBA* for short, has a set Q_F of accepting locations, and accepts an infinite timed path θ if θ visits some location in Q_F infinitely often, i.e., $\text{inf}(\theta) \cap Q_F \neq \emptyset$. The expressive power of (*D*)*TMA* and (*D*)*TBA* is related as follows [1]:

$$TMA = TBA > DTMA > DTBA.$$

Note that in nondeterministic *TMA* and *TBA*, guards on edges emanating from a location may overlap. *DTMA* are closed under all Boolean operators (union, intersection, and complement), while *DTBA* are *not* closed under complement.

Remark 2.8. [Successor location] Since DTA are deterministic, the edge relation \rightarrow can be replaced by a (partial) function $\text{succ} : Q \times \Sigma \times \mathcal{CC}(\mathcal{X}) \mapsto 2^{\mathcal{X}} \times Q$. If only the successor location is of interest, we simply use the function $\widetilde{\text{succ}} : Q \times \Sigma \times \mathcal{CC}(\mathcal{X}) \mapsto Q$, i.e., $q' = \widetilde{\text{succ}}(q, a, g)$.

2.3. Piecewise-deterministic Markov processes. PDPs [15] constitute a general model for stochastic systems without diffusions [16] and has been applied to a variety of problems in engineering, operations research, management science, and economics. Powerful analysis and control techniques for PDPs have been developed [23, 24, 13]. A PDP is a hybrid stochastic process involving discrete control (i.e., locations) and continuous variables.

Let us introduce some auxiliary notions. Let $\mathcal{X} = \{x_1, \dots, x_n\}$ be a set of variables in \mathbb{R} . Note that clock variables are a special case of these variables. A *constraint* over \mathcal{X} , denoted by g , is a subset of \mathbb{R}^n . Let $\mathcal{B}(\mathcal{X})$ denote the set of constraints over \mathcal{X} . An \mathcal{X} -valuation η satisfies constraint g , denoted $\eta \models g$, if and only if $(\eta(x_1), \dots, \eta(x_n)) \in g$. For $g \in \mathcal{B}(\mathcal{X})$, a constraint over $\mathcal{X} = \{x_1, \dots, x_n\}$, let \bar{g} be the closure of g , \mathring{g} the interior of g , and $\partial g = \bar{g} \setminus \mathring{g}$ the boundary of g . For instance, for $g = x_1^2 - 2x_2 \leq 1.5 \wedge x_3 > 2$, we have $\mathring{g} = x_1^2 - 2x_2 < 1.5 \wedge x_3 > 2$, $\bar{g} = x_1^2 - 2x_2 \leq 1.5 \wedge x_3 \geq 2$, and ∂g equals $x_1^2 - 2x_2 = 1.5 \wedge x_3 = 2$.

To each control location z of a PDP, an *invariant* $Inv(z)$ is associated, a constraint over \mathcal{X} which constrains the variable values in z . The state of a PDP is a pair (z, η) with control location z and η a variable valuation. Let $\mathbb{S} = \{(z, \eta) \mid z \in Z, \eta \models Inv(z)\}$, where Z is the set of locations. The notions of closure, interior and boundary can be lifted to \mathbb{S} in a straightforward manner, e.g., $\partial \mathbb{S} = \bigcup_{z \in Z} \{z\} \times \partial Inv(z)$ is the boundary of \mathbb{S} ; $\mathring{\mathbb{S}}$ and $\bar{\mathbb{S}}$ are defined in a similar way.

Definition 2.9 (PDP [16]). A *piecewise-deterministic (Markov) process* (PDP) is a tuple $\mathcal{Z} = (Z, \mathcal{X}, Inv, \phi, \Lambda, \mu)$ where Z is a finite set of *locations*, \mathcal{X} is a finite set of *variables*, $Inv : Z \rightarrow \mathcal{B}(\mathcal{X})$ is an *invariant function*, and

- $\phi : Z \times \mathcal{V}(\mathcal{X}) \times \mathbb{R} \rightarrow \mathcal{V}(\mathcal{X})$ is a *flow function*, which is the solution of a system of ODEs with a Lipschitz continuous vector field,
- $\Lambda : \mathbb{S} \rightarrow \mathbb{R}_{\geq 0}$ is an *exit rate function* satisfying for any $\xi \in \mathbb{S}$:

$$\exists \epsilon(\xi) > 0. \text{function } t \mapsto \Lambda(\xi \oplus t) \text{ is integrable on } [0, \epsilon(\xi)), \quad (\Delta)$$

where $(z, \eta) \oplus t = (z, \phi(z, \eta, t))$, and

- $\mu : \mathring{\mathbb{S}} \rightarrow \text{Distr}(\mathbb{S})$ is the *transition probability function* satisfying:

$$\mu(\xi, \{\xi\}) = 0 \quad \text{and} \quad \xi \mapsto \mu(\xi, A) \text{ is measurable for any } A \in \mathcal{F}(\mathbb{S}),$$

where $\mu(\xi, A)$ denotes $(\mu(\xi))(A)$, $\mathcal{F}(\mathbb{S})$ is a σ -algebra generated by $\bigcup_{z \in Z} \{z\} \times A_z$ with $A_z \subseteq \mathcal{F}(Inv(z))$, and $\mathcal{F}(Inv(z))$ is a σ -algebra generated by $Inv(z)$.

Let us explain the behavior of a PDP. A PDP can reside in a state $\xi = (z, \eta) \in \mathring{\mathbb{S}}$ as long as $Inv(z)$ holds. In state $\xi = (z, \eta)$, the PDP can either *delay* or take a *Markovian jump*. Delaying by t time units yields the next state $\xi' = \xi \oplus t$, i.e., the PDP remains in location z while all its continuous variables are updated according to $\phi(z, \eta, t)$. The flow function ϕ defines the time-dependent behavior in a single location, in particular, it specifies how the variable valuations change when time elapses. In case of a Markovian jump in state ξ , the next state $\xi'' = (z'', \eta'') \in \mathbb{S}$ is reached with probability $\mu(\xi, \{\xi''\})$. The residence time of a state is exponentially distributed; this is defined by the function Λ . A third possibility for a PDP to evolve is by taking *forced transitions*. When the variable valuation η satisfies the boundary of the invariant, i.e., $\eta \models \partial Inv(z)$, the PDP is forced to take a boundary jump, i.e., it has to leave state ξ . With probability $\mu(\xi, \{\xi''\})$ it then moves to state ξ'' . For any $T \in \mathbb{R}_{\geq 0}$, the function Λ is integrable as the interval $[0, T]$ can be divided into finitely many small intervals, on which by equation (Δ) , the function Λ is integrable.




Figure 4: An example PDP with constant exit rate 5 and boundary measure $\mu((z_0, 2), \{(z_1, 2)\}) := 1$

A PDP is named piecewise-deterministic because in each location (one piece) the behavior is deterministically determined by the flow function ϕ . The PDP is Markovian as the current state contains all the information to determine the future progress of the PDP.

2.4. Embedded PDP. The embedded *discrete-time Markov process* (DTMP) $emb(\mathcal{Z})$ of the PDP \mathcal{Z} has the same state space \mathbb{S} as \mathcal{Z} and is equipped with a transition probability function $\hat{\mu}$. The one-jump *transition probability* from a state ξ to a set $A \subseteq \mathbb{S}$ of states (with different location as ξ), denoted $\hat{\mu}(\xi, A)$, is given by [16]:

$$\hat{\mu}(\xi, A) = \int_0^{b(\xi)} (\mathcal{Q}\mathbf{1}_A)(\xi \oplus t) \cdot \Lambda(\xi \oplus t) e^{-\int_0^t \Lambda(\xi \oplus \tau) d\tau} dt \quad (2.2)$$

$$+ (\mathcal{Q}\mathbf{1}_A)(\xi \oplus b(\xi)) \cdot e^{-\int_0^{b(\xi)} \Lambda(\xi \oplus \tau) d\tau} \quad (2.3)$$

where $b(\xi) = \inf\{t > 0 \mid \xi \oplus t \in \partial\mathbb{S}\}$ is the minimal time to hit the boundary if such time exists; $b(\xi) = \infty$ otherwise. $(\mathcal{Q}\mathbf{1}_A)(\xi) = \int_{\mathbb{S}} \mathbf{1}_A(\xi') \mu(\xi, d\xi')$ is the accumulative (one-jump) transition probability from ξ to A and $\mathbf{1}_A(\xi)$ is the characteristic function such that $\mathbf{1}_A(\xi) = 1$ when $\xi \in A$ and $\mathbf{1}_A(\xi) = 0$ otherwise. Term (2.2) specifies the probability to delay to state $\xi \oplus t$ (on the same location) and take a Markovian jump from $\xi \oplus t$ to A . Note the delay t can take a value from $[0, b(\xi))$. Term (2.3) is the probability to stay in the same location for $b(\xi)$ time units and then it is forced to take a boundary jump from $\xi \oplus b(\xi)$ to A since $Inv(z)$ will be by any delay invalid.

Example 2.10. Figure 4 depicts a 3-location PDP \mathcal{Z} with $\mathcal{X} = x$, where $Inv(z_0) = x < 2$ and $Inv(z_1) = Inv(z_2) = x \in \mathbb{R}_{\geq 0}$. Solving $\dot{x} = 1$ yields the flow function $\phi(z_i, \eta(x), t) = \eta(x) + t$ for $i = 0, 1, 2$. The state space of \mathcal{Z} is $\mathbb{S} = \{(z_0, \eta) \mid \eta(x) < 2\} \cup \{(z_1, \mathbb{R}_{\geq 0})\} \cup \{(z_2, \mathbb{R}_{\geq 0})\}$. Let exit rate $\Lambda(\xi) = 5$ for any $\xi \in \mathbb{S}$. For $\eta \models Inv(z_0)$, let $\mu((z_0, \eta), \{(z_1, \eta)\}) := \frac{1}{3}$, $\mu((z_0, \eta), \{(z_2, \eta)\}) := \frac{2}{3}$ and the boundary measure be given as $\mu((z_0, 2), \{(z_1, 2)\}) := 1$. The time for $\xi_0 = (z_0, 0)$ to hit the boundary is $b(\xi_0) = 2$. For set of states $A = \{(z_1, \mathbb{R})\}$ and state ξ_0 , $(\mathcal{Q}\mathbf{1}_A)(\xi_0 \oplus t) = \frac{1}{3}$ if $t < 2$, and $(\mathcal{Q}\mathbf{1}_A)(\xi_0 \oplus t) = 1$ if $t = 2$. This yields for the transition probability from state ξ_0 to A in $emb(\mathcal{Z})$ is:

$$\hat{\mu}(\xi_0, A) = \int_0^2 \frac{1}{3} \cdot 5 \cdot e^{-\int_0^t 5 d\tau} dt + 1 \cdot e^{-\int_0^2 5 d\tau} = \frac{1}{3} + \frac{2}{3} e^{-10}.$$

3. THE PRODUCT OF A CTMC AND A DTA

In this section, we will make the first steps towards the quantitative and qualitative verification of CTMCs against linear real-time properties specified by DTA. The aim is to computing the probability of the set of paths in CTMC \mathcal{C} accepted by a DTA \mathcal{A} , i.e.,

$\Pr(\mathcal{C} \models \mathcal{A})$. We first prove that this question is well-defined, i.e., that this set of paths is measurable. The next step is to define the product of a CTMC \mathcal{C} and a DTA \mathcal{A} . As we will see, this is neither a CTMC nor a DTA, but a mixture of the two. We define the semantics of such products and define a probability space on their paths. The central result of this section is that $\Pr(\mathcal{C} \models \mathcal{A})$ equals the reachability probability in the product of \mathcal{C} and \mathcal{A} , cf. Theorem 3.10. In order to facilitate the effective computation of these reachability probabilities, we adapt the region construction of timed automata to the product $\mathcal{C} \otimes \mathcal{A}$, and show that this yields a PDP. The analysis of these PDPs will be the subject of the next two sections.

To simplify the notations, we assume w.l.o.g. that a CTMC has a single initial state s_0 , i.e., $\alpha(s_0) = 1$, and $\alpha(s) = 0$ for $s \neq s_0$. The state labels of the CTMC will act as input symbols of the DTA. Thus, the alphabet of DTA equals the powerset of the atomic propositions, i.e., 2^{AP} . A timed path in a CTMC is accepted by a DTA \mathcal{A} if there exists a corresponding accepting path in \mathcal{A} .

Definition 3.1 (CTMC paths accepted by a DTA). Let CTMC $\mathcal{C} = (S, \text{AP}, L, s_0, \mathbf{P}, E)$ and DTA $\mathcal{A} = (2^{\text{AP}}, \mathcal{X}, Q, q_0, Q_{\mathbf{F}}, \rightarrow)$. The CTMC path $s_0 \xrightarrow{t_0} s_1 \xrightarrow{t_1} s_2 \cdots$ is *accepted by* \mathcal{A} if there exists a corresponding DTA path

$$q_0 \xrightarrow{L(s_0), t_0} \underbrace{\widetilde{\text{succ}}(q_0, L(s_0), g_0)}_{=q_1} \xrightarrow{L(s_1), t_1} \underbrace{\widetilde{\text{succ}}(q_1, L(s_1), g_1)}_{=q_2} \cdots$$

which is accepted by \mathcal{A} , where $\eta_0 = \vec{0}$, g_i is the (unique) guard in q_i such that $\eta_i + t_i \models g_i$ and $\eta_{i+1} = (\eta_i + t_i)[X_i := 0]$, and η_i is the clock evaluation when entering q_i , for all i .

3.1. Measurability. The quantitative verification of CTMC \mathcal{C} against DTA \mathcal{A} amounts to compute the probability of the set of paths in \mathcal{C} that is accepted by \mathcal{A} . Formally, let

$$\text{Paths}^{\mathcal{C}}(\mathcal{A}) = \{ \rho \in \text{Paths}^{\mathcal{C}} \mid \rho \text{ is accepted by DTA } \mathcal{A} \}.$$

We first prove its measurability:

Theorem 3.2. *For any CTMC \mathcal{C} and DTA \mathcal{A} , $\text{Paths}^{\mathcal{C}}(\mathcal{A})$ is measurable.*

Proof. It suffices to show that $\text{Paths}^{\mathcal{C}}(\mathcal{A})$ can be written as a finite union or intersection of measurable sets. The proof is split in two parts: DTA with (1) reachability acceptance, and (2) Muller acceptance. The proof of the first case is carried out by (1a) considering DTA that only contain strict inequalities as guards, (1b) equalities, and (1c) non-strict inequalities. (Note that constraint $x = K$ can be obtained by $x > K \wedge x \geq K$).

(1a): Let $\text{DTA}^{\diamond} \mathcal{A}$ only contain strict inequalities as clock constraints. As all accepting paths are finite, $\text{Paths}^{\mathcal{C}}(\mathcal{A}) = \bigcup_{n \in \mathbb{N}} \text{Paths}_n^{\mathcal{C}}(\mathcal{A})$, where $\text{Paths}_n^{\mathcal{C}}(\mathcal{A})$ is the set of paths of length n accepted by \mathcal{A} . Let $\rho = s_0 \xrightarrow{t_0} s_1 \cdots s_{n-1} \xrightarrow{t_{n-1}} s_n \in \text{Paths}_n^{\mathcal{C}}(\mathcal{A})$. Then there exists a corresponding path $\theta = q_0 \xrightarrow{L(s_0), t_0} q_1 \cdots q_{n-1} \xrightarrow{L(s_{n-1}), t_{n-1}} q_n$ of \mathcal{A} which is induced by the sequence:

$$q_0 \xrightarrow{L(s_0), g_0, X_0} q_1 \cdots q_{n-1} \xrightarrow{L(s_{n-1}), g_{n-1}, X_{n-1}} q_n,$$

with $q_n \in Q_{\mathbf{F}}$ such that there exist $\{\eta_i\}_{0 \leq i < n}$ with 1) $\eta_0 = \vec{0}$; 2) $\eta_i + t_i \models g_i$; and 3) $\eta_{i+1} = (\eta_i + t_i)[X_i := 0]$, where η_i is the clock valuation when entering q_i .

We prove the measurability of $Paths_n^C(\mathcal{A})$ by showing that for any path

$$\rho = s_0 \xrightarrow{t_0} \dots \xrightarrow{t_{n-1}} s_n \in Paths_n^C(\mathcal{A}),$$

there exists a cylinder set $C(s_0, I_0, \dots, I_{n-1}, s_n)$ (C_ρ for short) such that:

$$\rho \in C_\rho \quad \text{and} \quad C_\rho \subseteq Paths_n^C(\mathcal{A}) \quad \text{for} \quad |\rho| = n. \quad (3.1)$$

This is proven in two steps:

- a. ($\rho \in C_\rho$.) Let $\rho = s_0 \xrightarrow{t_0} \dots \xrightarrow{t_{n-1}} s_n \in Paths_n^C(\mathcal{A})$. We define C_ρ by considering intervals I_i with rational bounds that are based on t_i . Let $I_i = [t_i^-, t_i^+]$ such that $t_i^- = t_i^+ := t_i$ if $t_i \in \mathbb{Q}$, and $t_i^-, t_i^+ \in \mathbb{Q}$ otherwise, such that:

$$t_i^- \leq t_i \leq t_i^+, \quad [t_i^-] = [t_i], \quad [t_i^+] = [t_i], \quad \text{and} \quad t_i^+ - t_i^- < \frac{\Delta}{2 \cdot n}$$

where $\Delta = \min_{0 \leq j < n, x \in \mathcal{X}} \left\{ \{\eta_j(x) + t_j\}, 1 - \{\eta_j(x) + t_j\} \mid \{\eta_j(x) + t_j\} \neq 0 \right\}$, with $\{\cdot\}$ denoting the fractional part. Since DTA \mathcal{A} only contains strict inequalities, for any i with $\eta_i + t_i \models g_i$, it follows $\{\eta_i(x) + t_i\} \neq 0$.

- b. ($C_\rho \subseteq Paths_n^C(\mathcal{A})$.) Let $\rho' := s_0 \xrightarrow{t'_0} \dots \xrightarrow{t'_{n-1}} s_n \in C_\rho$. Let $\eta'_0 := \vec{0}$ and $\eta'_{i+1} := (\eta'_i + t'_i)[X_i := 0]$. It remains to show that $\eta'_i + t'_i \models g_i$. Observe that $\eta'_0 = \eta_0$, and for any $i > 0$ and clock variable x ,

$$|\eta'_i(x) - \eta_i(x)| \leq \sum_{j=0}^{i-1} |t'_j - t_j| \leq \sum_{j=0}^{i-1} t_j^+ - t_j^- \leq n \cdot (t_j^+ - t_j^-) \leq \frac{\Delta}{2}.$$

Given that guard g_i only contains strict inequalities, it follows $\eta'_i + t'_i \models g_i$. This can be seen as follows. Let $g_i = x > K$ for some natural K . As $|\eta'_i(x) - \eta_i(x)| \leq \frac{\Delta}{2}$ and $|t'_i - t_i| < \frac{\Delta}{2}$, it follows $|(\eta'_i(x) + t'_i) - (\eta_i(x) + t_i)| < \Delta$. Note that $\eta_i(x) + t_i > K$, and thus $\eta_i(x) + t_i - \{\eta_i(x) + t_i\} = \lceil \eta_i(x) + t_i \rceil \geq K$. Hence, $\eta_i(x) + t_i - \Delta \geq K$ since, by definition, $\Delta \leq \{\eta_i(x) + t_i\}$. It follows that $\eta'_i(x) + t'_i > K$. A similar argument applies to the case $x < K$ and extends to conjunctions of strict inequalities. Thus, $\eta'_i + t'_i \models g_i$, and $\rho' \in Paths_n^C(\mathcal{A})$.

By (3.1) and the fact that $Paths_n^C(\mathcal{A}) \subseteq \bigcup_{\rho \in Paths_n^C(\mathcal{A})} C_\rho$, we have:

$$Paths_n^C(\mathcal{A}) = \bigcup_{\rho \in Paths_n^C(\mathcal{A})} C_\rho \quad \text{and} \quad Paths^C(\mathcal{A}) = \bigcup_{n \in \mathbb{N}} \bigcup_{\rho \in Paths_n^C(\mathcal{A})} C_\rho.$$

As each interval in C_ρ has rational bounds, C_ρ is measurable. It follows that $Paths^C(\mathcal{A})$ is a union of *countably many* cylinder sets, and hence is measurable.

- (1b):** Consider DTA $^\diamond$ \mathcal{A} with equalities of the form $x = K$ for natural K . Measurability is shown by induction on the number of equalities in \mathcal{A} . The base case (only strict inequalities) has been shown above. Now suppose there exists an edge $e = q \xrightarrow{a, g, X} q'$ in \mathcal{A} where g contains the constraint $x = K$. Let DTA $^\diamond$ \mathcal{A}_e be obtained from \mathcal{A} by deleting all the outgoing edges from q except e . We then consider the DTA $\bar{\mathcal{A}}_e$, $\mathcal{A}_e^>$, and $\mathcal{A}_e^<$ where $\bar{\mathcal{A}}_e$ is obtained from \mathcal{A}_e by replacing $x = K$ by *true*; $\mathcal{A}_e^>$ is obtained from \mathcal{A}_e by replacing $x = K$ by $x > K$ and $\mathcal{A}_e^<$ is obtained from \mathcal{A}_e by replacing $x = K$ by $x < K$. Since \mathcal{A} is deterministic, it follows that

$$Paths^C(\mathcal{A}_e) = Paths^C(\bar{\mathcal{A}}_e) \setminus (Paths^C(\mathcal{A}_e^>) \cup Paths^C(\mathcal{A}_e^<)).$$

By the induction hypothesis, the sets $Paths^C(\bar{\mathcal{A}}_e)$, $Paths^C(\mathcal{A}_e^>)$ and $Paths^C(\mathcal{A}_e^<)$ are measurable. Hence, $Paths^C(\mathcal{A}_e)$ is measurable. Furthermore, as

$$Paths^C(\mathcal{A}) = \bigcup_{e=q \xrightarrow{a, g, X} q'} Paths^C(\mathcal{A}_e),$$

where all guards g of edge e are equalities, it follows that $Paths^C(\mathcal{A})$ is measurable.

(1c): Let $DTA^\diamond \mathcal{A}$ have clock constraints of the form $x \bowtie K$ where $\bowtie \in \{\geq, \leq\}$. We consider the DTA $\mathcal{A}_=$ and \mathcal{A}_{\bowtie} , where $\mathcal{A}_=$ is obtained from \mathcal{A} by changing all constraints of the form $x \bowtie K$ by $x = K$, and \mathcal{A}_{\bowtie} is obtained from \mathcal{A} by changing any constraint $x \bowtie K$ by $x \bar{\bowtie} K$, with $\bar{\geq} = >$ and $\bar{\leq} = <$ otherwise. Clearly, $Paths^C(\mathcal{A}) = Paths^C(\mathcal{A}_=) \cup Paths^C(\mathcal{A}_{\bowtie})$. As it was shown before that $Paths^C(\mathcal{A}_=)$ and $Paths^C(\mathcal{A}_{\bowtie})$ are measurable, it follows that $Paths^C(\mathcal{A})$ is measurable.

(2): Let $DTA^\omega \mathcal{A}$ with $Q_{\mathcal{F}} = \{F_1, \dots, F_k\}$. $Paths^C(\mathcal{A}) = \bigcap_{0 < i \leq k} Paths^i$ where $Paths^i$ is the set of paths in CTMC \mathcal{C} whose corresponding DTA paths are accepted by $F_i \in Q_{\mathcal{F}}$, i.e., $Paths^i = \{\theta \in Paths^C(\mathcal{A}) \mid \inf(\theta) = F_i\}$. We have:

$$Paths^i = \bigcap_{n \geq 0} \bigcup_{m \geq n} \bigcup_{s_0, \dots, s_n, s_{n+1}, \dots, s_m} C(s_0, I_0, \dots, I_{n-1}, s_n, \dots, I_{m-1}, s_m),$$

where $\{s_{n+1}, \dots, s_m\} = L_{F_i}$ with L_{F_i} the set of CTMC states whose corresponding DTA states are F_i , and $C(s_0, I_0, \dots, I_{n-1}, s_n, \dots, I_{m-1}, s_m)$ is the cylinder set such that each timed path of the cylinder set of the form $s_0 \xrightarrow{t_0} \dots \xrightarrow{t_{n-1}} s_n \dots \xrightarrow{t_{m-1}} s_m$ is a prefix of an accepting path of \mathcal{A} . It follows that $Paths^i$ is measurable. Thus, $Paths^C(\mathcal{A})$ is measurable. \square

3.2. The product of a CTMC and a DTA. A central step in the verification of a CTMC \mathcal{C} against a DTA \mathcal{A} is to construct its synchronous product $\mathcal{C} \otimes \mathcal{A}$. The resulting object is neither a CTMC nor a DTA, but a mixture of the two. We first define this model, called deterministic Markovian timed automata, and define a measurable space over its paths. In Section 4, we consider the computation of $\Pr(\mathcal{C} \models \mathcal{A}) = \Pr(Paths^C(\mathcal{A}))$ which is based on this product.

Definition 3.3 (DMTA). A *deterministic Markovian timed automaton* (DMTA) is a tuple $\mathcal{M} = (Loc, \mathcal{X}, \ell_0, Loc_{\mathbf{F}}, E, \rightsquigarrow)$, where Loc is a nonempty finite set of *locations*; \mathcal{X} is a finite set of *clocks*; $\ell_0 \in Loc$ is the *initial location*; $Loc_{\mathbf{F}}$ is the *acceptance condition* with $Loc_{\mathbf{F}} = Loc_F \subseteq Loc$ the reachability condition and $Loc_{\mathbf{F}} = Loc_{\mathcal{F}} \subseteq 2^{Loc}$ the Muller condition; $E : Loc \rightarrow \mathbb{R}_{\geq 0}$ is the *exit rate function*; and $\rightsquigarrow \subseteq Loc \times \mathcal{CC}(\mathcal{X}) \times 2^{\mathcal{X}} \times Distr(Loc)$ is an *edge relation* such that:

$$\left(\ell \xrightarrow{g, X} \zeta \text{ and } \ell \xrightarrow{g', X'} \zeta' \text{ with } g \neq g' \right) \text{ implies } g \cap g' = \emptyset.$$

DMTA closely resemble DTA, but have in addition to DTA an exit rate function that determines the random residence time in a location, and an edge relation where the target of an edge is a probability distribution over the locations. Concepts such as clock valuation, clock constraints and so forth are defined as for DTA. We refer to $\ell \xrightarrow{g, X} \zeta$ for distribution $\zeta \in Distr(Loc)$ as an *edge* and to $\ell \xrightarrow[p]{g, X} \ell'$ with $p = \zeta(\ell')$ as a *transition* of this edge.

The intuition is that when entering location ℓ , the DMTA chooses a residence time which is governed by an exponential distribution with rate $E(\ell)$. Thus, the probability to leave ℓ within t time units is $1 - e^{-E(\ell)t}$. Due to the determinism of the edge relation, at most one edge, say $\ell \xrightarrow{g, X} \zeta$, is enabled. The probability to jump to ℓ' via this edge equals $\zeta(\ell')$. Similar as for DTAs, DMTA^\diamond and DMTA^ω are defined and DMTA refers to both classes.

Definition 3.4 (DMTA paths). An (infinite) *symbolic path* of DMTA \mathcal{M} is of the form:

$$\ell_0 \xrightarrow[p_0]{g_0, X_0} \ell_1 \xrightarrow[p_1]{g_1, X_1} \ell_2 \cdots \quad \text{where } \ell_i \xrightarrow{\zeta_i, X_i} \zeta_i \text{ and } p_i = \zeta_i(\ell_{i+1}), \text{ for all } i \in \mathbb{N}.$$

A symbolic path induces *infinite paths* of the form $\tau = \ell_0 \xrightarrow{t_0} \ell_1 \xrightarrow{t_1} \ell_2 \cdots$ such that $\eta_0 = \vec{0}$, $(\eta_i + t_i) \models g_i$, and $\eta_{i+1} = (\eta_i + t_i)[X_i := 0]$ where $i \geq 0$ and η_i is the clock valuation of \mathcal{X} in \mathcal{M} when *entering* location ℓ_i . The path τ is *accepted* by a DMTA^\diamond if there exists $n \geq 0$, such that $\tau[n] \in \text{Loc}_F$. It is *accepted* by DMTA^ω if and only if $\text{inf}(\tau) \in \text{Loc}_F$.

DMTA semantics. Consider clock valuation η in location ℓ . As the DMTA is deterministic, at most one guard is enabled in state (ℓ, η) . The *one-jump* probability of taking the transition $\ell \xrightarrow[p]{g, X} \ell'$ within time interval I starting at clock valuation η in location ℓ , denoted $p_\eta(\ell, \ell', I)$, is defined as follows:

$$p_\eta(\ell, \ell', I) = \int_I \underbrace{E(\ell) \cdot e^{-E(\ell)\tau}}_{\text{(i) density to leave } \ell \text{ at } \tau} \cdot \underbrace{\mathbf{1}_g(\eta + \tau)}_{\text{(ii) } \eta + \tau \models g?} \cdot \underbrace{p}_{\text{(iii) probabilistic jump}} d\tau \quad (3.2)$$

Note the resemblance with (2.1). Actually, part (i) characterizes the delay τ at location ℓ which is exponentially distributed with rate $E(\ell)$; (ii) is the *characteristic function*, where $\mathbf{1}_g(\eta + \tau) = 1$ if and only if $\eta + \tau \models g$. It compares the current valuation $\eta + \tau$ with guard g and rules out those violating g . Part (iii) indicates the probability of the transition under consideration. Note that (i) and (iii) are features from CTMCS while (ii) stems from DTA. The characteristic function $\mathbf{1}_g$ is Riemann integrable as it is bounded and its support is an interval; therefore, $p_\eta(\ell, \ell', I)$ is well-defined. The one-jump probability can be uniquely defined in this way because it relates to a fixed clock evaluation η .




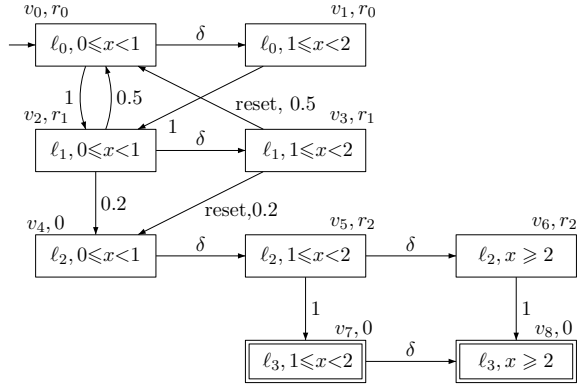
The above characterisation of the one-jump probability provides the basis for defining the probability of a set of DMTA paths. Let $C(\ell_0, I_0, \dots, I_{n-1}, \ell_n)$ be the cylinder set with $(\ell_0, \dots, \ell_n) \in \text{Loc}^{n+1}$ and $I_i \subseteq \mathbb{R}_{\geq 0}$. It denotes a set of paths in DMTA \mathcal{M} such that for any such path τ , $\tau[i] = \ell_i$ and $\tau[i] \in I_i$. Let $\text{Pr}_{\eta_0}^{\mathcal{M}}(C(\ell_0, I_0, \dots, I_{n-1}, \ell_n))$ denote the probability of $C(\ell_0, I_0, \dots, I_{n-1}, \ell_n)$ such that η_0 is the initial clock valuation in location ℓ_0 . Let $\text{Pr}_{\eta_0}^{\mathcal{M}}(C(\ell_0, I_0, \dots, I_{n-1}, \ell_n)) = \mathbb{P}_0^{\mathcal{M}}(\eta_0)$, where $\mathbb{P}_i^{\mathcal{M}}(\eta)$ is inductively defined as follows:

$$\mathbb{P}_i^{\mathcal{M}}(\eta) = \begin{cases} 1 & \text{if } i = n \\ \int_{I_i} \underbrace{E(\ell_i) \cdot e^{-E(\ell_i)\tau} \cdot \mathbf{1}_{g_i}(\eta + \tau) \cdot p_i}_{(\star)} \cdot \underbrace{\mathbb{P}_{i+1}^{\mathcal{M}}(\eta')}_{(\star\star)} d\tau & \text{if } 0 \leq i < n, \end{cases} \quad (3.3)$$

where $\eta' := (\eta + \tau)[X_i := 0]$. Intuitively, $\mathbb{P}_i^{\mathcal{M}}(\eta_i)$ is the probability of the suffix cylinder set starting from ℓ_i and η_i to ℓ_n . It is recursively defined by the product of the probability of taking a transition from ℓ_i to ℓ_{i+1} within time interval I_i (cf. (\star) and (3.2)) and the

probability of the suffix cylinder set from ℓ_{i+1} and η_{i+1} on (cf. $(\star\star)$). For the same reason as $p_\eta(\ell, \ell', I)$ is well-defined, $\mathbb{P}_i^{\mathcal{M}}(\eta)$ is well-defined.

Example 3.5. The DMTA $^\diamond$ in Figure 5(a) has initial location ℓ_0 with two outgoing edges, with guards $x < 1$ and $1 < x < 2$. We use the small black dots to indicate distributions. Assume t time units elapse in ℓ_0 . If the current clock evaluation η satisfies $\eta(x) < 1$, then the upper edge is enabled and the probability to go to ℓ_1 within time t is $p_{\bar{0}}(\ell_0, \ell_1, [0, t]) = (1 - e^{-r_0 t}) \cdot 1$, where $E(\ell_0) = r_0$; no clock is reset. It is similar when $1 < \eta(x) < 2$, except that x will be reset (cf. the lower edge emanating from location ℓ_0). If $\eta(x) \geq 2$, no outgoing edge is enabled, and the DMTA stays in ℓ_0 ad infinitum.

(a) DMTA $^\diamond \mathcal{C} \otimes \mathcal{A}$ (b) CTMC \mathcal{C} (c) DTA $^\diamond \mathcal{A}$ (d) Reachable region graph of $\mathcal{C} \otimes \mathcal{A}$ Figure 5: Example product DMTA $^\diamond$ of CTMC \mathcal{C} and DTA $^\diamond \mathcal{A}$

3.3. Product DMTA. The product $\mathcal{C} \otimes \mathcal{A}$ for CTMC \mathcal{C} and DTA \mathcal{A} , is a DMTA.

Definition 3.6 (Product of CTMC and DTA). Let $\mathcal{C} = (S, \text{AP}, L, s_0, \mathbf{P}, E)$ be a CTMC and $\mathcal{A} = (2^{\text{AP}}, \mathcal{X}, Q, q_0, Q_{\mathbf{F}}, \rightarrow)$ be a DTA. Let $\mathcal{C} \otimes \mathcal{A} = (Loc, \mathcal{X}, \ell_0, Loc_{\mathbf{F}}, E, \rightsquigarrow)$ be the product DMTA, where $Loc = S \times Q$; $\ell_0 = \langle s_0, q_0 \rangle$; $E(\langle s, q \rangle) = E(s)$; and

- $Loc_{\mathbf{F}} = Loc_F := S \times Q_F$, if $Q_{\mathbf{F}} = Q_F$ (reachability condition)
- $Loc_{\mathbf{F}} = Loc_{\mathcal{F}} := \bigcup_{F \in Q_{\mathcal{F}}} S \times F$, if $Q_{\mathbf{F}} = Q_{\mathcal{F}}$ (Muller condition)

and \rightsquigarrow is defined as the smallest relation defined by the rule:

$$\frac{\mathbf{P}(s, s') > 0 \wedge q \xrightarrow{L(s), g, X} q'}{\langle s, q \rangle \rightsquigarrow^{g, X} \zeta} \text{ such that } \zeta(\langle s', q' \rangle) = \mathbf{P}(s, s').$$

The DMTA $\mathcal{C} \otimes \mathcal{A}$ is basically the synchronous product of CTMC \mathcal{C} and DTA \mathcal{A} such that transition $s \rightarrow s'$ in \mathcal{C} is matched with the edge $q \xrightarrow{L(s), g, X} q'$, i.e., the set of atomic propositions of s acts as input symbol for the edge from location q to q' in \mathcal{A} . The probability of the joint evolvment of \mathcal{C} and \mathcal{A} is given by $\mathbf{P}(s, s')$, the discrete probability of $s \rightarrow s'$ in \mathcal{C} , whereas the residence time in the location $\langle s, q \rangle$ is given by $E(s)$, the exit rate of s in \mathcal{C} . It is easy to see from the construction that $\mathcal{C} \otimes \mathcal{A}$ is indeed a DMTA. The determinism of the DTA \mathcal{A} guarantees that the induced product is also deterministic. In $\mathcal{C} \otimes \mathcal{A}$, from each location there is at most one “input symbol” possible, viz. $L(s)$. For the sake of convenience, input symbols can be omitted from $\mathcal{C} \otimes \mathcal{A}$.

Example 3.7. Let CTMC \mathcal{C} and DTA $^\diamond$ \mathcal{A} be given in Figure 5(b) and 5(c), respectively. The product DMTA $^\diamond$ $\mathcal{C} \otimes \mathcal{A}$ is depicted in Figure 5(a). Since $Q_F = \{q_1\}$ in \mathcal{A} , the set of accepting locations in DMTA $^\diamond$ is $Loc_F = \{\langle s_2, q_1 \rangle\} = \{\ell_3\}$.

Example 3.8. For the CTMC \mathcal{C} in Figure 6(a) and the DTA $^\omega$ \mathcal{A} in Figure 6(b) with acceptance family $Q_{\mathcal{F}} = \{\{q_1, q_2\}, \{q_3, q_4\}\}$, the product DMTA $^\omega$ $\mathcal{C} \otimes \mathcal{A}$ is shown in Figure 6(c). $Loc_{\mathcal{F}} = \{\{\langle s_i, q_1 \rangle, \langle s_j, q_2 \rangle\}, \{\langle s'_i, q_3 \rangle, \langle s'_j, q_4 \rangle\}\}$, for any $s_i, s'_i, s_j, s'_j \in S$, i.e., $Loc_{\mathcal{F}} = \{\{\ell_1, \ell_2, \ell_3\}, \{\ell_4, \ell_5, \ell_6\}\}$.

The set of accepted paths in DMTA $\mathcal{C} \otimes \mathcal{A}$ is defined by:

$$AccPaths^{\mathcal{C} \otimes \mathcal{A}} := \{\tau \in Paths^{\mathcal{C} \otimes \mathcal{A}} \mid \tau \text{ is accepted by } \mathcal{C} \otimes \mathcal{A}\}.$$

For n -ary tuple J , let $J|_i$ denote the i -th entry in J , for $1 \leq i \leq n$. For a $(\mathcal{C} \otimes \mathcal{A})$ -path $\tau = \langle s_0, q_0 \rangle \xrightarrow{t_0} \langle s_1, q_1 \rangle \xrightarrow{t_1} \dots$, let $\tau|_1 := s_0 \xrightarrow{t_0} s_1 \xrightarrow{t_1} \dots$, and for any set Π of $(\mathcal{C} \otimes \mathcal{A})$ -paths, let $\Pi|_1 = \bigcup_{\tau \in \Pi} \tau|_1$. The following lemma asserts that there is a one-to-one relationship between paths in CTMC \mathcal{C} accepted by DTA \mathcal{A} and accepting paths in $\mathcal{C} \otimes \mathcal{A}$.

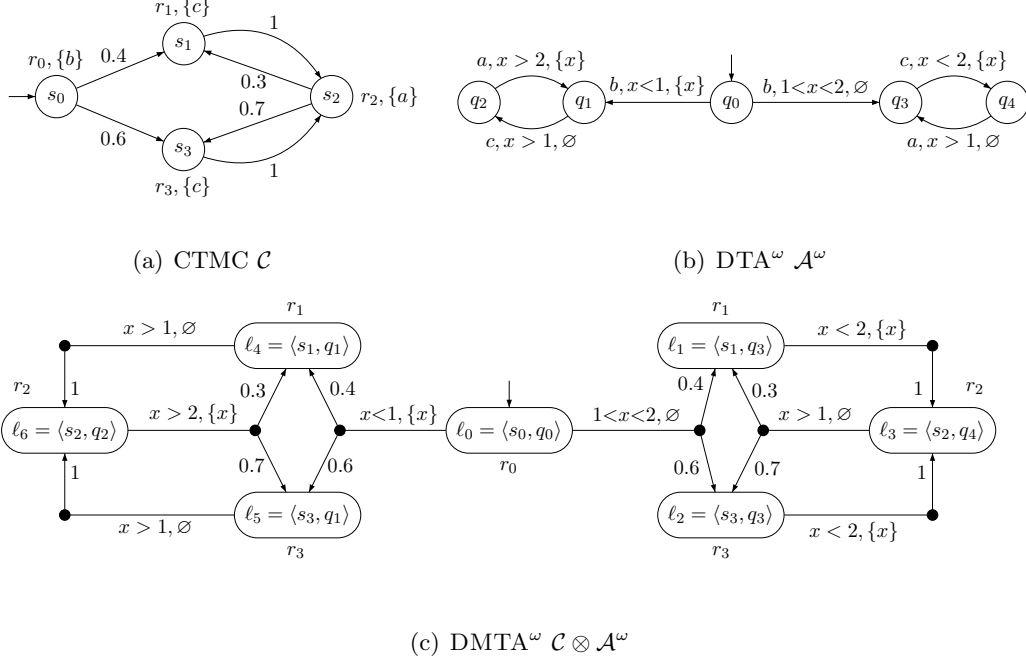
Lemma 3.9. *For any CTMC \mathcal{C} and DTA \mathcal{A} , $Paths^{\mathcal{C}}(\mathcal{A}) = AccPaths^{\mathcal{C} \otimes \mathcal{A}}|_1$.*

Proof. We provide the proof for DTA $^\diamond$ \mathcal{A} ; the proof for DTA $^\omega$ \mathcal{A} is similar.

(\subseteq) Let $\rho \in Paths^{\mathcal{C}}(\mathcal{A})$. We prove that there exists a path $\tau \in AccPaths^{\mathcal{C} \otimes \mathcal{A}}$ with $\rho = \tau|_1$. Assume w.l.o.g. that $\rho = s_0 \xrightarrow{t_0} s_1 \dots s_{n-1} \xrightarrow{t_{n-1}} s_n \in Paths^{\mathcal{C}}(\mathcal{A})$, i.e., $s_n \in Q_F$, $\eta_0 \models \vec{0}$, and for $0 \leq i < n$, $\eta_i + t_i \models g_i$ and $\eta_{i+1} = (\eta_i + t_i)[X_i := 0]$, where η_i is the clock valuation in \mathcal{A} when entering state s_i in \mathcal{C} . We construct a timed path $\theta \in Paths^{\mathcal{A}}$ from ρ such that $\theta = q_0 \xrightarrow{L(s_0), t_0} q_1 \dots q_{n-1} \xrightarrow{L(s_{n-1}), t_{n-1}} q_n$, where the clock valuation on entering s_i and q_i coincides. From ρ and θ , we can now construct the path

$$\tau = \langle s_0, q_0 \rangle \xrightarrow{t_0} \langle s_1, q_1 \rangle \dots \langle s_{n-1}, q_{n-1} \rangle \xrightarrow{t_{n-1}} \langle s_n, q_n \rangle,$$

where $\langle s_n, q_n \rangle \in Loc_F$. It follows that $\tau \in AccPaths^{\mathcal{C} \otimes \mathcal{A}}$ and $\rho = \tau|_1$.

Figure 6: Example product DMTA^ω of CTMC \mathcal{C} and $\text{DTA}^\omega \mathcal{A}^\omega$

(\supseteq) Let $\tau \in \text{AccPaths}^{\mathcal{C} \otimes \mathcal{A}}$. We prove that $\tau|_1 \in \text{Paths}^{\mathcal{C}}(\mathcal{A})$. Assume w.l.o.g. that

$$\tau = \langle s_0, q_0 \rangle \xrightarrow{t_0} \dots \xrightarrow{t_{n-1}} \langle s_n, q_n \rangle \in \text{AccPaths}^{\mathcal{C} \otimes \mathcal{A}},$$

with $\langle s_n, q_n \rangle \in \text{Loc}_F$, $\eta_0 \models \vec{0}$, and for $0 \leq i < n$, $\eta_i + t_i \models g_i$ and $\eta_{i+1} = (\eta_i + t_i)[X_i := 0]$, where η_i is the clock valuation when entering location $\langle s_i, q_i \rangle$. It then directly follows that $q_n \in Q_F$ and $\tau|_1 \in \text{Paths}^{\mathcal{C}}(\mathcal{A})$, given the entering clock valuation η_i of state s_i . \square

Theorem 3.10. For any CTMC \mathcal{C} and DTA \mathcal{A} ,

$$\Pr^{\mathcal{C}}(\text{Paths}^{\mathcal{C}}(\mathcal{A})) = \Pr_{\vec{0}}^{\mathcal{C} \otimes \mathcal{A}}(\text{AccPaths}^{\mathcal{C} \otimes \mathcal{A}}).$$

Proof. We provide the proof for $\text{DTA}^\diamond \mathcal{A}$; the proof for $\text{DTA}^\omega \mathcal{A}$ goes along similar lines as in the proof of Theorem 3.2.

According to Theorem 3.2, $\text{Paths}^{\mathcal{C}}(\mathcal{A})$ can be rewritten as the combination of cylinder sets of the form $C(s_0, I_0, \dots, I_{n-1}, s_n)$ which are all accepted by $\text{DTA}^\diamond \mathcal{A}$. Note that this means that each path in the cylinder set is accepted by \mathcal{A} . By Lemma 3.9, namely by path lifting, we can establish exactly the same combination of cylinder sets $C(\ell_0, I_0, \dots, I_{n-1}, \ell_n)$ for $\text{AccPaths}^{\mathcal{C} \otimes \mathcal{A}}$, where $s_i = \ell_i|_1$. It then suffices to show that for each cylinder set $C(s_0, I_0, \dots, I_{n-1}, s_n)$ which is accepted by \mathcal{A} , $\Pr^{\mathcal{C}}$ and $\Pr^{\mathcal{C} \otimes \mathcal{A}}$ yield the same probabilities.

For the measure $\Pr^{\mathcal{C}}$, according to Eq. (2.1) (cf. page 5),

$$\Pr^{\mathcal{C}}(C(s_0, I_0, \dots, I_{n-1}, s_n)) = \prod_{0 \leq i < n} \int_{I_i} \mathbf{P}(s_i, s_{i+1}) \cdot E(s_i) \cdot e^{-E(s_i)\tau} d\tau.$$

The measure $\Pr_{\vec{0}}^{\mathcal{C} \otimes \mathcal{A}}$, according to Section 3.2, is given by $\mathbb{P}_0^{\mathcal{C} \otimes \mathcal{A}}(\vec{0})$, where $\mathbb{P}_n^{\mathcal{C} \otimes \mathcal{A}}(\eta) = 1$ for any clock valuation η and for any $0 \leq i < n$:

$$\mathbb{P}_i^{\mathcal{C} \otimes \mathcal{A}}(\eta_i) = \int_{I_i} \mathbf{1}_{g_i}(\eta_i + \tau_i) \cdot p_i \cdot E(\ell_i) \cdot e^{-E(\ell_i)\tau_i} \cdot \mathbb{P}_{i+1}^{\mathcal{C} \otimes \mathcal{A}}(\eta_{i+1}) \, d\tau_i,$$

where $\eta_{i+1} = (\eta_i + \tau_i)[X_i := 0]$ and $\mathbf{1}_{g_i}(\eta_i + \tau_i) = 1$, if $\eta_i + \tau_i \models g_i$; 0, otherwise.

We will show, by induction, that $\mathbb{P}_i^{\mathcal{C} \otimes \mathcal{A}}(\eta_i)$ is a constant, i.e., is independent of η_i , if the cylinder set $C(\ell_0, I_0, \dots, I_{n-1}, \ell_n)$ is accepted by $\mathcal{C} \otimes \mathcal{A}$. First note that for this cylinder set there must exist some sequence of transitions

$$\ell_0 \xrightarrow[p_0]{g_0, X_0} \ell_1 \cdots \ell_{n-1} \xrightarrow[p_{n-1}]{g_{n-1}, X_{n-1}} \ell_n$$

with $\eta_0 = \vec{0}$ and $\forall t_i \in I_i$ with $0 \leq i < n$, $\eta_i + t_i \models g_i$ and $\eta_{i+1} := (\eta_i + t_i)[X_i := 0]$. Moreover, according to Definition 3.6, we have:

$$p_i = \mathbf{P}(s_i, s_{i+1}) \quad \text{and} \quad E(\ell_i) = E(s_i). \quad (3.4)$$

We apply a backward induction on n down to 0. The base case is trivial since $\mathbb{P}_n^{\mathcal{C} \otimes \mathcal{A}}(\eta_n) = 1$. By the induction hypothesis, $\mathbb{P}_{i+1}^{\mathcal{C} \otimes \mathcal{A}}(\eta_{i+1})$ is a constant. For the induction step, consider $i < n$. For any $\tau_i \in I_i$, since $\eta_i + \tau_i \models g_i$, $\mathbf{1}_{g_i}(\eta_i + \tau_i) = 1$, it follows that

$$\begin{aligned} \mathbb{P}_i^{\mathcal{C} \otimes \mathcal{A}}(\eta_i) &= \int_{I_i} \mathbf{1}_{g_i}(\eta_i + \tau_i) \cdot p_i \cdot E(\ell_i) \cdot e^{-E(\ell_i)\tau_i} \cdot \mathbb{P}_{i+1}^{\mathcal{C} \otimes \mathcal{A}}(\eta_{i+1}) \, d\tau_i \\ &\stackrel{\text{I.H.}}{=} \int_{I_i} p_i \cdot E(\ell_i) \cdot e^{-E(\ell_i)\tau_i} \, d\tau_i \cdot \mathbb{P}_{i+1}^{\mathcal{C} \otimes \mathcal{A}}(\eta_{i+1}) \\ &\stackrel{\text{Eq.(3.4)}}{=} \int_{I_i} \mathbf{P}(s_i, s_{i+1}) \cdot E(s_i) \cdot e^{-E(s_i)\tau_i} \, d\tau_i \cdot \mathbb{P}_{i+1}^{\mathcal{C} \otimes \mathcal{A}}(\eta_{i+1}). \end{aligned}$$

Clearly, this is a constant. It is thus easy to see that

$$\Pr_{\vec{0}}^{\mathcal{C} \otimes \mathcal{A}}(C(\ell_0, I_0, \dots, I_{n-1}, \ell_n)) := \mathbb{P}_0^{\mathcal{C} \otimes \mathcal{A}}(\vec{0}) = \prod_{0 \leq i < n} \int_{I_i} \mathbf{P}(s_i, s_{i+1}) \cdot E(s_i) \cdot e^{-E(s_i)\tau} \, d\tau,$$

which completes the proof. \square

3.4. Region graph construction. Theorem 3.10 asserts that the probability of CTMC \mathcal{C} satisfying the DTA specification \mathcal{A} equals the reachability probability of some accepting location in $\mathcal{C} \otimes \mathcal{A}$. The state space of $\mathcal{C} \otimes \mathcal{A}$, however, is infinite. As a next step towards obtaining an effective procedure for computing reachability probabilities in $\mathcal{C} \otimes \mathcal{A}$ we adopt the standard region construction of timed automata [1] to DMTA. This yields a stochastic process, namely a PDP. Here, we consider the region construction for finite acceptance conditions, i.e, DMTA $^\diamond$. The details for DMTA $^\omega$ are slightly different (only the acceptance set differs) and are provided in Section 5.

Let us briefly recall the concept of a region. Formally, a region is an equivalence under \cong , an equivalence relation on clock valuations. A region is characterized by a specific form of a clock constraint. Let c_{x_i} be the largest constant with which $x_i \in \mathcal{X}$ is compared in some guard in the (DM)TA. Clock evaluations $\eta, \eta' \in \mathcal{V}(\mathcal{X})$ are *clock-equivalent*, denoted $\eta \cong \eta'$, if and only if either

- (1) for any $x \in \mathcal{X}$ it holds that $\eta(x) > c_x$ and $\eta'(x) > c_x$, or

(2) for any $x_i, x_j \in \mathcal{X}$ with $\eta(x_i), \eta'(x_i) \leq c_{x_i}$ and $\eta(x_j), \eta'(x_j) \leq c_{x_j}$ it holds:

$$\lfloor \eta(x_i) \rfloor = \lfloor \eta'(x_i) \rfloor \quad \text{and} \quad \{\eta(x_i)\} \leq \{\eta'(x_i)\} \text{ iff } \eta(x_j) \leq \eta'(x_j),$$

where $\lfloor d \rfloor$ and $\{d\}$ are the integral and fractional part of $d \in \mathbb{R}$, respectively.

This clock equivalence is coarser than the traditional definition by merging the “boundary” regions (those with point constraints like “ $x = 0$ ”) into the “non-boundary” regions (those only with interval constraints like “ $0 < y < 1$ ”). For instance, for $\mathcal{X} = \{x_1, x_2\}$, the boundary regions $(x_1 = 0, x_2 = 0)$, $(0 < x_1 < 1, x_2 = 0)$ and $(x_1 = 0, 0 < x_2 < 1)$ are merged with the non-boundary region $(0 < x_1 < 1, 0 < x_2 < 1)$ yielding $(0 \leq x_1 < 1, 0 \leq x_2 < 1)$. The reason for this slight change will become clear later.

Let $\mathcal{R}e(\mathcal{X})$ be the set of regions over the set \mathcal{X} of clocks. For $\Theta, \Theta' \in \mathcal{R}e(\mathcal{X})$, Θ' is the *successor region* of Θ if for all $\eta \models \Theta$ there exists $\delta \in \mathbb{R}_{>0}$ such that $\eta + \delta \models \Theta'$ and $\forall \delta' < \delta. \eta + \delta' \models \Theta \vee \Theta'$. The region Θ *satisfies* the guard g , denoted $\Theta \models g$, iff $\forall \eta \models \Theta. \eta \models g$. The *reset operation* on region Θ is defined as $\Theta[X := 0] := \{\eta[X := 0] \mid \eta \models \Theta\}$.

Definition 3.11 (Region graph of DMTA $^\diamond$). The *region graph* of DMTA $^\diamond$ $\mathcal{M} = (Loc, \mathcal{X}, \ell_0, Loc_F, E, \rightsquigarrow)$ is $\mathcal{G}(\mathcal{M}) = (V, v_0, V_F, \Lambda, \hookrightarrow)$, where

- $V = Loc \times \mathcal{R}e(\mathcal{X})$ is a finite set of *vertices* with *initial vertex* $v_0 = (\ell_0, \vec{0})$;
- $V_F = \{v \in V \mid v|_1 \in Loc_F\}$ is the set of *accepting vertices*;
- $\Lambda : V \rightarrow \mathbb{R}_{\geq 0}$ is the *exit rate function* where:

$$\Lambda(v) = \begin{cases} E(v|_1) & \text{if } v \xrightarrow{p, X} v' \text{ for some } v' \in V \\ 0 & \text{otherwise.} \end{cases}$$

- $\hookrightarrow \subseteq V \times (([0, 1] \times 2^{\mathcal{X}}) \cup \{\delta\}) \times V$ is the *transition (edge) relation*, such that:
 - ▶ $v \xrightarrow{\delta} v'$ if $v|_1 = v'|_1$, and $v'|_2$ is the successor region of $v|_2$;
 - ▶ $v \xrightarrow{p, X} v'$ if $v|_1 \xrightarrow{g, X} v'|_1$ with $v|_2 \models g$, and $v|_2[X := 0] \subseteq v'|_2$.

Any vertex in the region graph is a pair consisting of a location and a region. Edges of the form $v \xrightarrow{\delta} v'$ are called *delay edges*, whereas those of the form $v \xrightarrow{p, X} v'$ are called *Markovian edges*. Note that Markovian edges emanating from a boundary region do *not* contribute to the reachability probability as the time to hit the boundary is always zero (i.e., $b(v, \eta) = 0$ in Eq. (4.3), page 20). Therefore, we can safely remove all the Markovian edges emanating from boundary regions and combine each such boundary region with its unique non-boundary (direct) successor. In the sequel, by slight abuse of notation, we refer to this *simplified region graph* as $\mathcal{G}(\mathcal{M})$. Note that then $v|_2[X := 0] \subseteq v'|_2$ in the last item of Definition 3.11.

Remark 3.12. [Exit rates] The exit rate $\Lambda(v)$ equals 0 if only delay transitions emanate from v . The probability to take the delay edge within time t is $e^{-\Lambda(v)t} = 1$, while the probability to take Markovian edges is 0.

Example 3.13. For the DMTA $^\diamond$ $\mathcal{C} \otimes \mathcal{A}$ in Figure 5(a), the reachable part (forward reachable from the initial vertex and backward reachable from the accepting vertices) of the simplified region graph $\mathcal{G}(\mathcal{C} \otimes \mathcal{A})$ is shown in Figure 5(d). Note that the exit rates on v_4 and v_7 are 0, as only a delay edge is enabled in these vertices.

The following result asserts that the region graph obtained from a DMTA is in fact a PDP. This is an important observation, as verification now reduces to analyzing this PDP.

Lemma 3.14. *The region graph of any DMTA induces a PDP.*

Proof. Let $\text{DMTA}^\diamond \mathcal{M} = (\text{Loc}, \mathcal{X}, \ell_0, \text{Loc}_F, E, \rightsquigarrow)$ with region graph $\mathcal{G}(\mathcal{M}) = (V, v_0, V_F, \Lambda, \hookrightarrow)$. Define $\mathcal{Z}(\mathcal{M}) = (V, \mathcal{X}, \text{Inv}, \phi, \Lambda, \mu)$ where for any $v \in V$:

- $\text{Inv}(v) := v|_2$ and the state space $\mathbb{S} := \{(v, \eta) \mid v \in V, \eta \models \text{Inv}(v)\}$;
- $\phi(v, \eta, t) := \eta + t$;
- $\Lambda(v, \eta) := \Lambda(v)$;
- if $v \xrightarrow{\delta} v'$ in $\mathcal{G}(\mathcal{M})$, then $\mu((v, \eta), \{(v', \eta)\}) := 1$, provided $\eta \models \partial \text{Inv}(v)$;
- if $v \xrightarrow{p, X} v'$ in $\mathcal{G}(\mathcal{M})$, then $\mu((v, \eta), \{(v', \eta[X := 0])\}) := p$, provided $\eta \models \text{Inv}(v)$.

It follows directly that $\mathcal{Z}(\mathcal{M})$ is a PDP. \square

Note that the acceptance conditions play no role in the definition of a PDP, thus this lemma applies to both DMTA^\diamond and DMTA^ω .

4. VERIFYING CTMCS AGAINST FINITE DTA SPECIFICATIONS

The characterization of the region graph of $\mathcal{C} \otimes \mathcal{A}$ as a PDP paves the way to the verification of CTMC \mathcal{C} against DTA^\diamond specification \mathcal{A} . This section concentrates on the quantitative verification problem and deals with single-clock DTA separately.

4.1. Quantitative verification with arbitrarily many clocks. The central issue in quantitative verification is to compute the probability of the set of paths in \mathcal{C} accepted by \mathcal{A} . By Theorem 3.10, this is equal to computing reachability probabilities in DTMA $\mathcal{C} \otimes \mathcal{A}$. The remaining question is how to determine these probabilities. To that end, we show that this amounts to determine reachability probabilities of untimed events in the embedded PDP of $\mathcal{Z}(\mathcal{C} \otimes \mathcal{A})$ (cf. Theorem 4.3 below). These probabilities are characterized by a Volterra integral equation system of second type. As solving this integral equation system is typically hard, we present an effective approximation algorithm.

Characterizing reachability probabilities. We first consider determining unbounded reachability probabilities in the PDP $\mathcal{Z} = \mathcal{Z}(\mathcal{C} \otimes \mathcal{A})$. This is done by considering its embedded PDP, the DTMP $\text{emb}(\mathcal{Z})$, as for unbounded reachability probabilities, the timing aspects are not important. Note that the set of locations of PDP \mathcal{Z} and $\text{emb}(\mathcal{Z})$ are equal. Besides, the discrete probabilistic evolution of \mathcal{Z} and $\text{emb}(\mathcal{Z})$ coincide. The main difference is that $\text{emb}(\mathcal{Z})$ is time-abstract whereas \mathcal{Z} is not.

Let initial state $(v_0, \vec{0})$ and $T \subseteq V$ be the set of goal locations. For state (v, η) , let $\text{Prob}^{\text{emb}(\mathcal{Z})}((v, \eta), T)$, $\text{Prob}_v(\eta, T)$ for short, denote the probability to reach some state in (T, \cdot) from state (v, η) in $\text{emb}(\mathcal{Z})$. These probabilities are recursively defined as follows. For vertex $v \in V$, we have:

$$\text{Prob}_v(\eta, T) = \begin{cases} 1 & \text{if } v \in T \\ \text{Prob}_{v, \delta}(\eta, T) + \sum_{v \xrightarrow{p, X} v'} \text{Prob}_{v, v'}(\eta, T) & \text{otherwise} \end{cases} \quad (4.1)$$

The case $v \in T$ is evident. In case $v \notin T$, then either a delay can take place (first summand), or a Markovian edge is taken to vertex v' (second summand).

For a delay transition $v \xrightarrow{\delta} v'$ we have:

$$Prob_{v,\delta}(\eta, T) = e^{-\Lambda(v) \cdot b(v,\eta)} \cdot Prob_{v'}(\eta + b(v,\eta), T), \quad (4.2)$$

where $e^{-\Lambda(v) \cdot b(v,\eta)}$ is the probability to stay in v for at most $b(v,\eta)$ time units. Recall that $b(v,\eta)$ is the minimal time for state (v,η) to hit the boundary $\partial Inv(v)$. Stated in other words, $e^{-\Lambda(v) \cdot b(v,\eta)}$ is the probability to reside in v without violating the invariant. The reachability probability from the resulting state $\eta + b(v,\eta)$ is then given by the second multiplicand in Eq. (4.2). This equation is based on Eq. (2.3) by determining the multi-step reachability probability using a sequence of one-step transition probabilities.

For the Markovian transition $v \xrightarrow{p,X} v'$, we have:

$$Prob_{v,v'}(\eta, T) = \int_0^{b(v,\eta)} p \cdot \Lambda(v) \cdot e^{-\Lambda(v) \cdot \tau} \cdot Prob_{v'}((\eta + \tau)[X := 0], T) d\tau. \quad (4.3)$$

Here, $\Lambda(v) \cdot e^{-\Lambda(v) \cdot \tau}$ denotes the density to stay for exactly τ time units in v . As any delay up to $b(v,\eta)$ does not violate the invariant, τ ranges over the dense interval $[0, b(v,\eta)]$. The state after first delaying τ time units and then taking the edge $v \xrightarrow{p,X} v'$ is $(\eta + \tau)[X := 0]$. Eq. (4.3) is derived from Eq. (2.2).

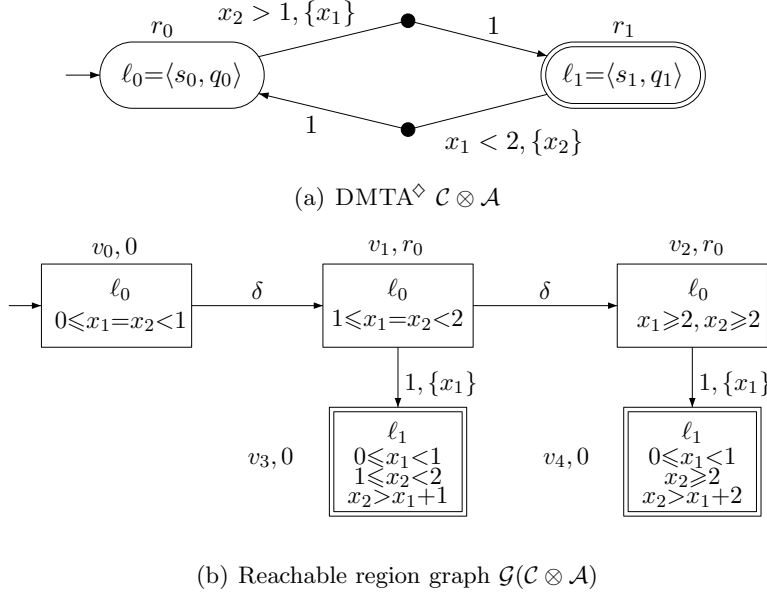


Figure 7: Reachable fragment of its region graph

Example 4.1. Consider the $DMTA^\diamond$ in Figure 7(a) and its region graph in Figure 7(b). Let $T = V_F$ be the set of goal locations, i.e., the set of target states $\{(v,\eta) \mid v \in V_F, \eta \models Inv(v)\}$. The system of integral equations for v_1 in location ℓ_0 is as follows. For $1 \leq x_1 = x_2 < 2$:

$$Prob_{v_1}(x_1, x_2) = Prob_{v_1,\delta}(x_1, x_2) + Prob_{v_1,v_3}(x_1, x_2),$$

where

$$Prob_{v_1,\delta}(x_1, x_2) = e^{-(2-x_1)r_0} \cdot Prob_{v_2}(2, 2)$$

and

$$Prob_{v_1, v_3}(x_1, x_2) = \int_0^{2-x_1} r_0 \cdot e^{-r_0 \tau} \cdot Prob_{v_3}(0, x_2 + \tau) d\tau$$

where $Prob_{v_3}(0, x_2 + \tau) = 1$. The integral equations for vertices v_2, v_4 are similar.

Remark 4.2. Clock valuations η and η' in region Θ may induce different reachability probabilities. This is due to the fact that η and η' may have different periods of time to hit the boundary, Thus, the probability for η and η' to either delay or take a Markovian transition may differ. This is in contrast with timed automata, as well as probabilistic extensions thereof [22], where clock valuations in the same region are not distinguished.

Hence, reachability probabilities in the embedded PDP of $\mathcal{Z}(\mathcal{C} \otimes \mathcal{A})$ are characterized by a system of *Volterra integral equations* (4.1). One can read (4.1) either in the form $f(\xi) = \int_{Dom(\xi)} K(\xi, \xi') f(d\xi')$, where K is the kernel and $Dom(\xi)$ is the domain of integration depending on the continuous state space \mathbb{S} ; or in the operator form $f(\xi) = (\mathcal{J}f)(\xi)$, where \mathcal{J} is the integration operator. Generally, (4.1) does *not* necessarily have a unique solution. It turns out that the reachability probability $Prob_{v_0}(\vec{0})$ coincides with the least fixpoint of the operator \mathcal{J} (denoted by $\text{lfp}\mathcal{J}$) i.e., $Prob_{v_0}(\vec{0}) = (\text{lfp}\mathcal{J})(v_0, \vec{0})$.

Theorem 4.3. For any CTMC \mathcal{C} and DTA $^\diamond$ \mathcal{A} ,

$$\text{Pr}_{\vec{0}}^{\mathcal{C} \otimes \mathcal{A}}(\text{AccPaths}^{\mathcal{C} \otimes \mathcal{A}}) \text{ is the least solution of } Prob_{v_0}^{\mathcal{D}}(\vec{0}, V_F),$$

where DTMP $\mathcal{D} = \text{emb}(\mathcal{Z}(\mathcal{C} \otimes \mathcal{A}))$.

Proof. Let $\text{Pr}_{\vec{0}}^{\mathcal{C} \otimes \mathcal{A}}(\text{AccPaths}^{\mathcal{C} \otimes \mathcal{A}})$ be the least solution of the system of integral equations:

$$\text{Pr}(\ell, \eta) = \begin{cases} 1 & \text{if } \ell \in \text{Loc}_F \\ \int_0^\infty E(\ell) \cdot e^{-E(\ell)\tau} \cdot \sum_{\ell \xrightarrow[p]{g, X} \ell'} \mathbf{1}_g(\eta+\tau) \cdot p \cdot \text{Pr}(\ell', (\eta+\tau)[X := 0]) d\tau & \text{otherwise,} \end{cases}$$

Informally, $\text{Pr}(\ell, \eta)$ is the probability to reach the set of locations Loc_F from location ℓ and clock valuation η . The above integral can be simplified as follows. W.l.o.g. assume clock constraints to be of the form $x \sqsubseteq c$, where $c \in \mathbb{N}$ and $\sqsubseteq \in \{\leq, <, \geq, >\}$. Then we have:

$$\text{Pr}(\ell, \eta) = \int_{t_1}^{t_2} E(\ell) \cdot e^{-E(\ell)\tau} \cdot \sum_{\ell \xrightarrow[p]{g, X} \ell'} p \cdot \text{Pr}(\ell', (\eta+\tau)[X := 0]) d\tau,$$

where $t_1, t_2 \in \mathbb{Q}_{\geq 0} \cup \{\infty\}$ and $\eta + \tau \models g$ for any $t_1 < \tau < t_2$.

If $\ell \in \text{Loc}_F$, the theorem follows directly. In the remainder of the proof, assume $\ell \notin \text{Loc}_F$. Our proof is based on showing that for any $\ell \notin \text{Loc}_F$ and clock valuation η ,

$$\text{Pr}(\ell, \eta) = Prob_{v_0}(\eta, V_F), \quad (4.4)$$

where v_0 is the initial vertex in the region graph $\mathcal{Z}(\mathcal{C} \otimes \mathcal{A})$ with $v_0|_1 = \ell$, and $V_F = \{v \in V \mid v|_1 \in \text{Loc}_F\}$. This is done as follows. For natural n , let $\text{Pr}^n(\ell, \eta)$ be the probability

to reach Loc_F in n steps in $\mathcal{C} \otimes \mathcal{A}$. For $n = 0$, we have $\Pr^n(\ell, \eta) = 1$ if $\ell \in Loc_F$ and 0, otherwise. For $n > 0$, we define inductively:

$$\Pr^n(\ell, \eta) = \int_{t_1}^{t_2} E(\ell) \cdot e^{-E(\ell)\tau} \cdot \sum_{\ell \xrightarrow[p]{g, X} \ell'} p \cdot \Pr^{n-1}(\ell', \eta') \, d\tau.$$

Similarly, let $Prob_v^n(\eta, V_F)$ be the probability to reach the set of goal states V_F in $n > 0$ steps:

$$Prob_v^n(\eta, V_F) = \begin{cases} Prob_{v, \delta}^n(\eta, V_F) + Prob_v^{s, n}(\eta, V_F), & \text{if } v \notin V_F \\ 1, & \text{otherwise} \end{cases} \quad (4.5)$$

$$Prob_v^{s, n}(\eta, V_F) = \int_0^{b(v, \eta)} \Lambda(v) \cdot e^{-\Lambda(v)\tau} \cdot \sum_{v \xrightarrow[p, X]{g, X} v'} p \cdot Prob_{v'}^{n-1}((\eta + \tau)[X := 0], V_F) \, d\tau, \quad (4.6)$$

$$Prob_{v, \delta}^n(\eta, V_F) = e^{-\Lambda(v)b(v, \eta)} \cdot Prob_{v'}^n(\eta + b(v, \eta), V_F). \quad (4.7)$$

In the sequel, we show that for any $n \in \mathbb{N}$, it holds:

$$\Pr^n(\ell, \eta) = Prob_{v_0}^n(\eta, V_F). \quad (4.8)$$

The theorem then follows from the fact that $\lim_{n \rightarrow \infty} \Pr^n(\ell, \eta) = \Pr(\ell, \eta)$ and, similarly, $\lim_{n \rightarrow \infty} Prob_v^n(\eta, V_F) = Prob_v(\eta, V_F)$.




Figure 8: The sub-region graph $\mathcal{Z}(\mathcal{C} \otimes \mathcal{A})$ for the transition from ℓ to ℓ' .

The proof of $\Pr^n(\ell, \eta) = Prob_{v_0}^n(\eta, V_F)$ is by induction on n .

- (1) (Base case.) For $n = 0$, $\Pr^0(\ell, \eta) = 0 = Prob_{v_0}^0(\eta, V_F)$ if $\ell \notin Loc_F$, and 1 otherwise.
- (2) (Induction step.) Consider $n+1$. Let edge $\ell \xrightarrow[g, X]{g, X} \zeta$ in $\mathcal{C} \otimes \mathcal{A}$. Assume the fragment of the region graph $\mathcal{Z}(\mathcal{C} \otimes \mathcal{A})$ that corresponds to this edge with $\zeta(\ell, \ell') > 0$ is as shown in Fig. 8. Location ℓ induces the vertices $\{v_i = (\ell, \Theta_i) \mid 0 \leq i \leq k\}$. Intuitively speaking, the transition from location ℓ to ℓ' is enabled in region Θ_i for $m \leq i \leq k$, whereas only a delay can take place in all regions Θ_i with $i < m$ (while staying in location ℓ).

Let $\hat{\eta}_i$ be the clock valuation when entering vertex v_i , i.e., $\hat{\eta}_0 = \eta$ and $\hat{\eta}_i = \hat{\eta}_{i-1} + b(v_{i-1}, \hat{\eta}_{i-1})$ for $0 < i \leq k$. It is assumed that $\hat{\eta}_i \not\models g$, where g is the guard of the edge at hand, for $i < m$ and $i > k$. Accordingly,

$$t_1 = \sum_{i=0}^{m-1} b(v_i, \hat{\eta}_i) \quad \text{and} \quad t_2 = \sum_{i=0}^k b(v_i, \hat{\eta}_i)$$

are the lower and upper bound, respectively, of the interval during which guard g holds.

For convenience, let $p_v^n(\eta) := Prob_{v,\delta}^n(\eta, V_F) + Prob_v^{s,n}(\eta, V_F)$. Given the fact that only a delay transition can be taken before time t_1 , it holds that

$$\begin{aligned} p_{v_0}^{n+1}(\eta) &= e^{-t_1\Lambda(v_0)} \cdot p_{v_m}^{n+1}(\hat{\eta}_m), \text{ where} \\ p_{v_m}^{n+1}(\hat{\eta}_m) &= Prob_{v_m,\delta}^{n+1}(\hat{\eta}_m, V_F) + Prob_{v_m}^{s,n+1}(\hat{\eta}_m, V_F). \end{aligned}$$

We now derive:

$$\begin{aligned} &e^{-t_1\Lambda(v_0)} \cdot Prob_{v_m}^{s,n+1}(\hat{\eta}_m, V_F) \\ &= e^{-t_1\Lambda(v_0)} \cdot \int_0^{b(v_m, \hat{\eta}_m)} \Lambda(v_m) \cdot e^{-\Lambda(v_m)\tau} \cdot \sum_{v_m \xrightarrow{p,X} v'_m} p \cdot Prob_{v'_m}^n((\hat{\eta}_m + \tau)[X := 0], V_F) d\tau \\ &= \int_{t_1}^{t_1 + b(v_m, \hat{\eta}_m)} \Lambda(v_m) \cdot e^{-\Lambda(v_m)\tau} \cdot \sum_{v_m \xrightarrow{p,X} v'_m} p \cdot Prob_{v'_m}^n((\hat{\eta}_m + \tau - t_1)[X := 0], V_F) d\tau. \end{aligned}$$

Now consider:

$$p_{v_0}^{n+1}(\eta) = e^{-t_1\Lambda(v_0)} \cdot Prob_{v_m,\delta}^{n+1}(\hat{\eta}_m, V_F) + e^{-t_1\Lambda(v_0)} \cdot Prob_{v_m}^{s,n+1}(\hat{\eta}_m, V_F).$$

Using the definition of $Prob_{v_m,\delta}^{n+1}(\hat{\eta}_m, V_F)$ (see Eq. (4.7)), together with the result derived above, yields the following sum of integrals:

$$\begin{aligned} p_{v_0}^{n+1}(\eta) &= \sum_{i=0}^{k-m} \int_{t_1 + \sum_{j=0}^{i-1} b(v_{m+j}, \hat{\eta}_{m+j})}^{t_1 + \sum_{j=0}^i b(v_{m+j}, \hat{\eta}_{m+j})} \Lambda(v_{m+i}) \cdot e^{-\Lambda(v_{m+i})\tau} \\ &\quad \cdot \underbrace{\sum_{v_{m+i} \xrightarrow{p,X} v'_{m+i}} p \cdot Prob_{v'_{m+i}}^n((\hat{\eta}_{m+i} + \tau - t_1 - \sum_{j=0}^{i-1} b(v_{m+j}, \hat{\eta}_{m+j}))[X := 0], V_F) d\tau}_{=F^n(\tau)}. \end{aligned}$$

Using $F^n(t)$ we obtain:

$$p_{v_0}^{n+1}(\eta) = \int_{t_1}^{t_2} \Lambda(v_0) \cdot e^{-\Lambda(v_0)\tau} \cdot F^n(\tau) d\tau. \quad (4.9)$$

Notice that

$$\hat{\eta}_{m+i} = \eta + \underbrace{\sum_{j=0}^{m-1} b(v_j, \hat{\eta}_j)}_{=t_1} + \sum_{j=0}^{i-1} b(v_{m+j}, \hat{\eta}_{m+j}).$$

Therefore, for any $t \in [t_1 + \sum_{j=0}^{i-1} b(v_{m+j}, \hat{\eta}_{m+j}), t_1 + \sum_{j=0}^i b(v_{m+j}, \hat{\eta}_{m+j})]$, $i \leq k - m$ we obtain

$$\hat{\eta}_{m+i} + t - t_1 - \sum_{j=0}^{i-1} b(v_{m+j}, \hat{\eta}_{m+j}) = \eta + t.$$

From the induction hypothesis (for n), it follows that $\Pr^n(\ell, \eta) = Prob_{v_0}^n(\eta, V_F)$ with $v_0|_1 = \ell$. Therefore, for any $t \in [t_1 + \sum_{j=0}^{i-1} b(v_{m+j}, \hat{\eta}_{m+j}), t_1 + \sum_{j=0}^i b(v_{m+j}, \hat{\eta}_{m+j})]$ and

$v'_{m+i}|_1 = \ell'$, $i \leq k - m$, we get

$$\begin{aligned}
F^n(t) &= \sum_{v_{m+i} \xrightarrow{p,X} v'_{m+i}} p \cdot \text{Prob}_{v'_{m+i}}^n \left((\hat{\eta}_{m+i} + t - t_1 - \sum_{j=0}^{i-1} \mathfrak{b}(v_{m+j}, \hat{\eta}_{m+j})) [X := 0], V_F \right) \\
&= \sum_{v_{m+i} \xrightarrow{p,X} v'_{m+i}} p \cdot \text{Prob}_{v'_{m+i}}^n \left((\eta + t) [X := 0], V_F \right) \\
&= \sum_{v_{m+i} \xrightarrow{p,X} v'_{m+i}} p \cdot \text{Pr}^n(\ell', (\eta + t)) [X := 0] \\
&= \sum_{\ell \xrightarrow[p]{g,X} \ell'} p \cdot \text{Pr}^n(\ell', (\eta + t)) [X := 0].
\end{aligned}$$

Substituting this result into equation (4.9) results in

$$p_{v_0}^{n+1}(\eta) = \int_{t_1}^{t_2} \Lambda(\ell) \cdot e^{-\Lambda(\ell)\tau} \cdot \sum_{\ell \xrightarrow[p]{g,X} \ell'} p \cdot \text{Pr}^n(\ell', (\eta + \tau)) [X := 0] d\tau.$$

As for $v_0 \notin V_F$, $\text{Prob}_{v_0}^{n+1}(\eta, V_F) = p_{v_0}^{n+1}(\eta)$ we get that $\text{Prob}_{v_0}^{n+1}(\eta, V_F) = \text{Pr}^{n+1}(\ell, \eta)$. \square

Note that, similar to the computation of reachability probabilities in DTMCs [18], the goal states in $T \subseteq \mathbb{S}$ as well as all states that cannot reach T can be made absorbing, i.e., all outgoing edges can be removed, without affecting the reachability probabilities. This may yield a substantial state-space reduction.

Approximating reachability probabilities. The results so far assert that $\text{Pr}(\mathcal{C} \models \mathcal{A})$ coincides with reachability probabilities in an embedded PDP that is obtained via a region construction applied on the product $\mathcal{C} \otimes \mathcal{A}$. The previous result shows that such reachability probabilities are characterized by Volterra equations of the second type [2]. Such integral equation systems can be solved using techniques explained in standard textbooks, such as [12]. An alternative option —inspired by a formulation of bounded reachability probabilities in arbitrary PDPs [16]— is to *approximate* the probability $\text{Pr}(\text{Paths}^{\mathcal{C}}(\mathcal{A}))$ by a system of partial differential equations (PDEs, for short). The intuition is to consider paths that are accepted within some time bound t_f . Let DTA $\mathcal{A}[t_f]$ be obtained by adding a single fresh clock z , say, to DTA \mathcal{A} which is never reset, and strengthening all guards of incoming edges into $q \in Q_F$ by adding the conjunct $z \leq t_f$. Obviously, $\text{Paths}^{\mathcal{C}}(\mathcal{A}[t_f]) \subseteq \text{Paths}^{\mathcal{C}}(\mathcal{A})$. Note that $\lim_{t_f \rightarrow \infty} \text{Pr}(\text{Paths}^{\mathcal{C}}(\mathcal{A}[t_f])) = \text{Pr}(\text{Paths}^{\mathcal{C}}(\mathcal{A}))$.

Given CTMC \mathcal{C} , DTA $^\diamond$ \mathcal{A} , time bound t_f and PDP $\mathcal{Z}(\mathcal{C} \otimes \mathcal{A}) = (V, \mathcal{X}, \text{Inv}, \phi, \Lambda, \mu)$, we have:

$$\text{Pr}^{\mathcal{C}}(\text{Paths}^{\mathcal{C}}(\mathcal{A}[t_f])) = \sum_{\bar{v} \in V_F} \int_{\text{Inv}(\bar{v})} \mathfrak{h}_{v_0}^{\bar{v}}(t_f, \vec{0}, d\eta),$$

where $\bar{h}_{v_0}^{\bar{v}}(t_f, \vec{0}, \bar{\eta})$ is the probability to reach the state $(\bar{v}, \bar{\eta})$, with $\bar{v} \in V_F$ and $\bar{\eta} \models \text{Inv}(\bar{v})$ at time t_f from state $(v_0, \vec{0})$. The transition probability function $\bar{h}_{v_0}^{\bar{v}}(t_f, \vec{0}, \bar{\eta})$ is described by the following equations:

- for $v \in V \setminus V_F$, $\bar{v} \in V_F$ with $\eta \models \text{Inv}(v)$, $\bar{\eta} \models \text{Inv}(v_f)$ and $y \in (0, t_f)$:

$$\frac{\partial \bar{h}_v^{\bar{v}}(y, \eta, \bar{\eta})}{\partial y} + \sum_{i=1}^{|\mathcal{X}|} \frac{\partial \bar{h}_v^{\bar{v}}(y, \eta, \bar{\eta})}{\partial \eta^{(i)}} + \Lambda(v) \cdot \sum_{\substack{v \xrightarrow{p, X} v'}} p \cdot (\bar{h}_{v'}^{\bar{v}}(y, \eta[X := 0], \bar{\eta}) - \bar{h}_v^{\bar{v}}(y, \eta, \bar{\eta})) = 0, \quad (4.10)$$

where $\eta^{(i)}$ is the i 'th clock variable.

- $\bar{h}_v^{\bar{v}}(0, \eta, \bar{\eta}) = 1$, when $v = \bar{v}$ and $\eta = \bar{\eta}$, $\bar{h}_v^{\bar{v}}(0, \eta, \bar{\eta}) = 0$, otherwise.
- the boundary conditions are: for $v, \bar{v} \in V$, $\eta \models \partial \text{Inv}(v)$, $\bar{\eta} \models \partial \text{Inv}(\bar{v})$ and transition $v \xrightarrow{\delta} v'$ we have $\bar{h}_v^{\bar{v}}(y, \eta, \bar{\eta}) = \bar{h}_{v'}^{\bar{v}}(y, \eta, \bar{\eta})$.

Equation (4.10) is obtained by simplifying a corresponding characterisation in Davis [16], where the author defines the function $\bar{h}_{v_0}^{\bar{v}}(\cdot)$ as an expectation. In our setting, $\bar{h}_{v_0}^{\bar{v}}(t_f, \vec{0}, \bar{\eta}) = \mathbb{E}[\mathbf{1}(X_{t_f}) | X_0 = \xi]$, where X_τ is the underlying stochastic process of the PDP \mathcal{Z} with the state space \mathbb{S} , $\xi = (v, \vec{0})$ and $\mathbf{1}(X_{t_f})$ is the characteristic function such that $\mathbf{1}(X_{t_f}) = 1$ if and only if $X_{t_f} = (\bar{v}, \bar{\eta})$. The PDE (4.10) is a special case of [16] as the flow function in \mathcal{Z} is linear and the probabilistic jumps to the continuous part of the state space \mathbb{S} are non-uniform.

4.2. Single-clock DTA $^\diamond$ specifications. For *single-clock* DTA $^\diamond$ specifications, we can simplify the system of Volterra integral equations (of second type) obtained in the previous section. As we will show in this subsection, the probability that a CTMC satisfies a single-clock DTA is given by a system of linear equations whose coefficients are a solution of a system of ODEs that can be solved efficiently. The key observation is that the region graph corresponding to $\mathcal{C} \otimes \mathcal{A}$ can be naturally divided into a number of subgraphs, each of which is a CTMC.

Let \mathcal{A} be a single-clock DTA with finite acceptance criterion, and $\{c_0, \dots, c_m\}$ be the set of natural numbers that appear in the clock constraints of \mathcal{A} . Assume $0 = c_0 < c_1 < \dots < c_m$, and let $\Delta c_i = c_{i+1} - c_i$ for $0 \leq i < m$. Note that for single-clock DTA, the regions in the region graph of $\mathcal{C} \otimes \mathcal{A}$ can be partitioned by the following intervals: $[c_0, c_1), [c_1, c_2), \dots, [c_m, \infty)$. Using this observation, we partition the region graph $\mathcal{Z}(\mathcal{C} \otimes \mathcal{A})$ as follows.

Definition 4.4 (Partitioning of region graph). Let $\mathcal{G}(\mathcal{C} \otimes \mathcal{A}) = (V, v_0, V_F, \Lambda, \hookrightarrow)$, or \mathcal{G} for short, for single-clock DTA $^\diamond$ \mathcal{A} . The *partitioning* of \mathcal{G} is defined as the collection of subgraphs $\mathcal{G}_i = (V_i, V_{F_i}, \Lambda_i, \hookrightarrow_i)$, for $0 \leq i \leq m$ where:

- $V_i = \{(\ell, \Theta) \in V \mid \Theta \subseteq [c_i, c_{i+1})\}$
- $V_{F_i} = V_i \cap V_F$,
- $\Lambda_i(v) = \Lambda(v)$ if $v \in V_i$, and 0 otherwise, and
- $\hookrightarrow = \bigcup_{0 \leq i \leq m} M_i \cup F_i \cup B_i$, where
 - M_i is the set of Markovian edges (without reset) between vertices in V_i ,
 - F_i is the set of delay edges between V_i and V_{i+1} ,
 - B_i is the set of Markovian edges (with reset) from V_i to V_0 .




Figure 9: Partitioning the region graph of Figure 5(d)

Since the initial vertex of \mathcal{G}_0 is v_0 and the initial vertices of \mathcal{G}_i for $0 < i \leq m$ are implicitly given by the edges in F_{i-1} , we omit them. Note that the subgraph \mathcal{G}_m involves only infinite regions and has no outgoing delay transitions.

Example 4.5. Consider the region graph in Figure 5(a) (page 14). The partitioning of this region graph is depicted in Figure 9. The edges in M_i , F_i and B_i are labeled with probabilities, δ (delay), and “reset” with probabilities, respectively. Observe that if $v = (\ell, [c_i, c_{i+1})) \in V_F$, then $v' = (\ell, [c_j, c_{j+1})) \in V_F$ for $i < j \leq m$. (In this example, this applies to $v = v_7$ and $v' = v_8$.) This is true since $V_F = \{(\ell, \text{true}) \mid \ell \in \text{Loc}_F\}$. Thus, from any final vertex in V_i with $i < m$, there is a delay transition to the next region (if any).

Assume $|V_i| = k_i$. We now define for each type of edge (M , B , or F) a matrix (\mathbf{M} , \mathbf{B} , and \mathbf{F} , respectively). Let $x \in \mathbb{R}$ with $x \in [0, \Delta c_i]$. Then:

- $\mathbf{D}_i(x) \in \mathbb{R}^{k_i \times k_i}$ is the delay probability matrix, where for any $0 \leq j \leq k_i$, $\mathbf{D}_i(x)[j, j] = e^{-E(v_j^i) \cdot x}$ and the off-diagonal elements are zero.
- $\mathbf{M}_i(x) = \mathbf{D}_i(x) \cdot \mathbf{E}_i \cdot \mathbf{P}_i \in \mathbb{R}^{k_i \times k_i}$ is the probability density matrix for M_i -edges, where \mathbf{P}_i and \mathbf{E}_i are the transition probability matrix and exit rate matrix respectively, for vertices in V_i .
- $\mathbf{B}_i(x) \in \mathbb{R}^{k_i \times k_0}$ is the probability density matrix for the B_i -edges, where $\mathbf{B}_i(x)[j, k]$ indicates the probability density function to take a B_i -edge from $v^j \in V_i$ to $v^k \in V_0$.
- $\mathbf{F}_i \in \mathbb{R}^{k_i \times k_{i+1}}$ is the incidence matrix for F_i -edges, i.e., $\mathbf{F}_i[j, k] = 1$ if and only if there is a delay transition between $v^j \in V_i$ and $v^k \in V_{i+1}$.

Due to the fact that in any subgraph \mathcal{G}_i there are only Markovian jumps without resets, and no delay transitions, the subgraph (V_i, Λ_i, M_i) , i.e., \mathcal{G} restricted to Markovian jumps (without resets) forms a CTMC \mathcal{C}_i , say. To take the effect of Markovian jumps with resets into account, we define for each \mathcal{G}_i the *augmented* CTMC \mathcal{C}_i^a with state space $V_i \cup V_0$, where all V_0 -vertices are absorbing, i.e., do not have any outgoing edges. The edges connecting V_i to V_0 are kept. The augmented CTMC is used to calculate the probability to start from a vertex in \mathcal{G}_i and take a reset edge within a certain period of time.

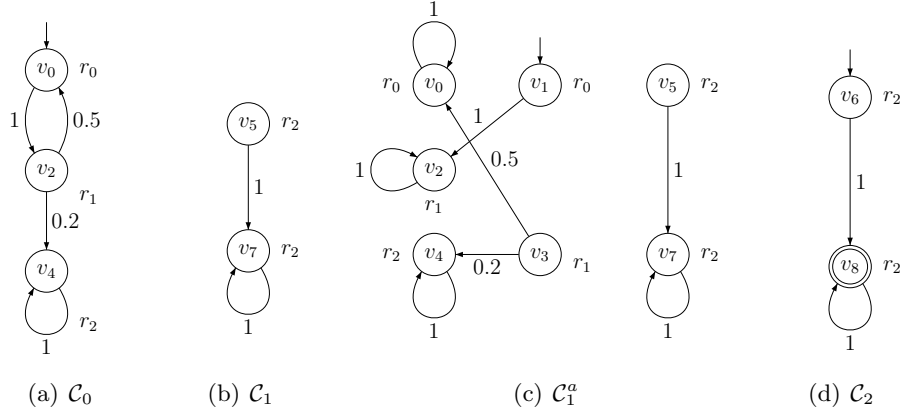


Figure 10: CTMCs corresponding to the (augmented) subgraphs

Example 4.6. Consider the partitioned region graph in Figure 9. The matrices for \mathcal{G}_0 are:

$$\mathbf{M}_0(x) = \begin{pmatrix} 0 & 1 \cdot r_0 \cdot e^{-r_0 x} & 0 \\ 0.5 \cdot r_1 \cdot e^{-r_1 x} & 0 & 0.2 \cdot r_1 \cdot e^{-r_1 x} \\ 0 & 0 & 0 \end{pmatrix} \quad \mathbf{F}_0 = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}$$

The matrices for \mathcal{G}_1 and its augmented version are given by:

$$\mathbf{M}_1(x) = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & r_2 \cdot e^{-r_2 x} \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad \mathbf{B}_1(x) = \begin{pmatrix} 0 & r_0 \cdot e^{-r_0 x} & 0 \\ 0.5 \cdot r_1 \cdot e^{-r_1 x} & 0 & 0.2 \cdot r_1 \cdot e^{-r_1 x} \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

$$\mathbf{F}_1 = \begin{pmatrix} 0 & 0 \\ 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{pmatrix} \quad \mathbf{M}_1^a(x) = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & r_0 \cdot e^{-r_0 x} & 0 \\ 0 & 0 & 0 & 0 & 0.5 \cdot r_1 \cdot e^{-r_1 x} & 0 & 0.2 \cdot r_1 \cdot e^{-r_1 x} \\ 0 & 0 & 0 & r_2 \cdot e^{-r_2 x} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

The corresponding CTMCs and their augmented version are depicted in Figure 10.

For CTMC \mathcal{C} with k states and rate matrix $\mathbf{E} \cdot \mathbf{P}$, let:

$$\mathbf{\Pi}(x) = \int_0^x \mathbf{M}(\tau) \mathbf{\Pi}(x - \tau) d\tau + \mathbf{D}(x). \quad (4.11)$$

Intuitively, $\mathbf{\Pi}(x)[j, m]$ indicates the probability to move from vertex j to m at time x . The following proposition states the close relationship between $\mathbf{\Pi}(x)$ and the transient probability vector of \mathcal{C} . Let $\vec{\varphi}(t)$ be the transient probability vector where $\varphi_s(t)$ is the probability to be in state s at time t given the initial distribution α .

Proposition 4.7. *Given a CTMC \mathcal{C} with initial distribution α , rate matrix $\mathbf{E} \cdot \mathbf{P}$ and $\mathbf{\Pi}(t)$, $\vec{\varphi}(t)$ satisfies the following two equations:*

$$\vec{\varphi}(t) = \alpha \cdot \mathbf{\Pi}(t), \quad (4.12)$$

$$\frac{d\vec{\varphi}(t)}{dt} = \vec{\varphi}(t) \cdot \mathbf{Q}, \quad (4.13)$$

where $\mathbf{Q} = \mathbf{E} \cdot \mathbf{P} - \mathbf{E}$ is the infinitesimal generator.

Equation (4.13) is the well-known forward Chapman-Kolmogorov equation. According to this proposition, solving the integral equation for $\mathbf{\Pi}(t)$ boils down to solving the system of ODEs (4.13) given some initial distribution vector α . This can be done using standard means for CTMCs such as uniformization.

Now let the probability vector $\vec{U}_i(x) = [u_i^1(x), \dots, u_i^{k_i}(x)]^\top \in \mathbb{R}^{k_i \times 1}$ where $u_i^j(x)$ is the probability to move from vertex $v_i^j \in V_i$ to some vertex in V_F (in \mathcal{G}) at time x . Based on the equations (4.1)-(4.3), we provide a set of integral equations for $\vec{U}_i(x)$ which later on is reduced to a system of linear equations. Distinguish two cases:

Case $0 \leq i < m$: for

$$\vec{U}_i(x) = \int_0^{\Delta c_i - x} \mathbf{M}_i(\tau) \vec{U}_i(x+\tau) d\tau + \int_0^{\Delta c_i - x} \mathbf{B}_i(\tau) d\tau \cdot \vec{U}_0(0) + \mathbf{D}_i(\Delta c_i - x) \cdot \mathbf{F}_i \vec{U}_{i+1}(0), \quad (4.14)$$

where $x \in [0, \Delta c_i]$. Let us explain this equation. The last summand is obtained from (4.2) where $\mathbf{D}_i(\Delta c_i - x)$ is the probability to delay until the ‘‘end’’ of region i , and $\mathbf{F}_i \vec{U}_{i+1}(0)$ denotes the probability to continue in \mathcal{G}_{i+1} (at relative time 0). Similarly, the first and second summands are obtained from (4.3); the former reflects the case where clock x is not reset, while the latter considers the reset of x (thus, implying a return to \mathcal{G}_0).

Case $i = m$:

$$\vec{U}_m(x) = \int_0^\infty \hat{\mathbf{M}}_m(\tau) \vec{U}_m(x+\tau) d\tau + \vec{\mathbf{1}}_F + \int_0^\infty \mathbf{B}_m(\tau) d\tau \cdot \vec{U}_0(0) \quad (4.15)$$

where for $x \in [c_m, \infty)$, $\hat{\mathbf{M}}_m(\tau)[v, \cdot] = \mathbf{M}_m(\tau)[v, \cdot]$ for $v \notin V_F$, 0 otherwise, and $\vec{\mathbf{1}}_F$ is the characteristic vector for V_F . Note that $\vec{\mathbf{1}}_F$ stems from the second clause of (4.1), and $\hat{\mathbf{M}}_m$ is obtained by setting the corresponding elements of \mathbf{M}_m to 0.

Example 4.8. The matrices for \mathcal{G}_2 are given as:

$$\hat{\mathbf{M}}_2(x) = \begin{pmatrix} 0 & r_2 \cdot e^{-r_2 x} \\ 0 & 0 \end{pmatrix} \quad \hat{\mathbf{P}}_2 = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$$

For augmented CTMC \mathcal{C}_i^a , let

$$\mathbf{\Pi}_i^a(x) = \left(\begin{array}{c|c} \mathbf{\Pi}_i(x) & \bar{\mathbf{\Pi}}_i^a(x) \\ \mathbf{0} & \mathbf{I} \end{array} \right),$$

where $\mathbf{0} \in \mathbb{R}^{k_0 \times k_i}$ is the zero matrix and $\mathbf{I} \in \mathbb{R}^{k_0 \times k_0}$ is the identity matrix. Matrix $\mathbf{\Pi}_i$ indicates the transient probabilities for the CTMC \mathcal{C}_i . Intuitively speaking, $\bar{\mathbf{\Pi}}_i^a$ contains the probabilities starting from V_i and ending in V_0 .

Theorem 4.9. *For subgraph \mathcal{G}_i (with k_i vertices) of \mathcal{G} , it holds that:*

$$\vec{U}_i(0) = \begin{cases} \mathbf{\Pi}_i(\Delta c_i) \cdot \mathbf{F}_i \cdot \vec{U}_{i+1}(0) + \bar{\mathbf{\Pi}}_i^a(\Delta c_i) \cdot \vec{U}_0(0) & \text{if } i < m \\ \hat{\mathbf{P}}_m \cdot \vec{U}_m(0) + \vec{\mathbf{1}}_F + \hat{\mathbf{B}}_m \cdot \vec{U}_0(0) & \text{if } i = m \end{cases}$$

where $\hat{\mathbf{P}}_i(v, v') = \mathbf{P}_i(v, v')$ if $v \notin V_F$; 0 otherwise and $\hat{\mathbf{B}}_m = \int_0^\infty \mathbf{B}_m(\tau) d\tau$.

Proof. Distinguish two cases: $i < m$ and $i = m$.

- (1) ($i < m$.) Consider the augmented CTMC \mathcal{C}_i^a with $k_i^a = k_i + k_0$ states. From equation (4.14), and the fact that \mathcal{C}_i^a contains reset edges of \mathcal{C}_i , we have:

$$\vec{U}_i^a(x) = \int_0^{\Delta c_i - x} \mathbf{M}_i^a(\tau) \cdot \vec{U}_i^a(x + \tau) d\tau + \mathbf{D}_i^a(\Delta c_i - x) \cdot \mathbf{F}_i^a \cdot \vec{U}_i(0)$$

where $\vec{U}_i^a(x) = \begin{pmatrix} \vec{U}_i(x) \\ \vec{U}_i'(x) \end{pmatrix} \in \mathbb{R}^{k_i^a \times 1}$, $\vec{U}_i'(x) \in \mathbb{R}^{k_0 \times 1}$ is the vector representing the reachability probabilities for the augmented states in \mathcal{G}_i , $\mathbf{F}_i^a = \begin{pmatrix} \mathbf{F}_i' & \mathbf{B}_i' \end{pmatrix} \in \mathbb{R}^{k_i^a \times (k_{i+1} + k_0)}$ such that $\mathbf{F}_i' = \begin{pmatrix} \mathbf{F}_i \\ \mathbf{0} \end{pmatrix} \in \mathbb{R}^{k_i^a \times k_{i+1}}$ is the incidence matrix for delay edges and $\mathbf{B}_i' = \begin{pmatrix} \mathbf{0} \\ \mathbf{I} \end{pmatrix} \in \mathbb{R}^{k_i^a \times k_0}$, and finally $\vec{U}_i(0) = \begin{pmatrix} \vec{U}_{i+1}(0) \\ \vec{U}_0(0) \end{pmatrix} \in \mathbb{R}^{(k_{i+1} + k_0) \times 1}$. The proof of the theorem for $i < m$ proceeds in two steps.

- (a) We first show that:

$$\begin{aligned} \vec{U}_i^a(x) &= \mathbf{\Pi}_i^a(\Delta c_i - x) \cdot \mathbf{F}_i^a \cdot \vec{U}_i(0) \text{ where} \\ \mathbf{\Pi}_i^a(x) &= \int_0^x \mathbf{M}_i^a(\tau) \cdot \mathbf{\Pi}_i^a(x - \tau) d\tau + \mathbf{D}_i^a(x). \end{aligned}$$

Consider the following system of inductively defined integral equations. Let $c_{i,x} = \Delta c_i - x$.

$$\begin{aligned} \vec{U}_i^{a,(0)}(x) &= \vec{0} \\ \vec{U}_i^{a,(j+1)}(x) &= \int_0^{c_{i,x}} \mathbf{M}_i^a(\tau) \cdot \vec{U}_i^{a,(j)}(x + \tau) d\tau + \mathbf{D}_i^a(c_{i,x}) \cdot \mathbf{F}_i^a \cdot \vec{U}_i(0). \end{aligned}$$

and

$$\begin{aligned} \mathbf{\Pi}_i^{a,(0)}(c_{i,x}) &= \mathbf{0} \\ \mathbf{\Pi}_i^{a,(j+1)}(c_{i,x}) &= \int_0^{c_{i,x}} \mathbf{M}_i^a(\tau) \cdot \mathbf{\Pi}_i^{a,(j)}(c_{i,x} - \tau) d\tau + \mathbf{D}_i^a(c_{i,x}). \end{aligned}$$

Clearly, $\mathbf{\Pi}_i^a(c_{i,x}) = \lim_{j \rightarrow \infty} \mathbf{\Pi}_i^{a,(j+1)}(c_{i,x})$ and $\vec{U}_i^a(x) = \lim_{j \rightarrow \infty} \vec{U}_i^{a,(j+1)}(x)$.

By induction on j , we prove the following relation:

$$\vec{U}_i^{a,(j)}(x) = \mathbf{\Pi}_i^{a,(j)}(c_{i,x}) \cdot \mathbf{F}_i^a \cdot \vec{U}_i(0).$$

- (i) (Base case.) $\vec{U}_i^{a,(0)}(x) = \vec{0}$ and $\mathbf{\Pi}_i^{a,(0)}(c_{i,x}) = \mathbf{0}$.

(ii) (Induction step.) By exploiting the induction hypothesis (in the second step), we derive:

$$\begin{aligned}
\vec{U}_i^{a,(j+1)}(x) &= \int_0^{c_{i,x}} \mathbf{M}_i^a(\tau) \vec{U}_i^{a,(j)}(x+\tau) d\tau + \mathbf{D}_i^a(c_{i,x}) \cdot \mathbf{F}_i^a \vec{U}_i^a(0) \\
&= \int_0^{c_{i,x}} \mathbf{M}_i^a(\tau) \mathbf{\Pi}_i^{a,(j)}(c_{i,x}-\tau) \cdot \mathbf{F}_i^a \vec{U}_i^a(0) d\tau + \mathbf{D}_i^a(c_{i,x}) \cdot \mathbf{F}_i^a \vec{U}_i^a(0) \\
&= \left(\int_0^{c_{i,x}} \mathbf{M}_i^a(\tau) \mathbf{\Pi}_i^{a,(j)}(c_{i,x}-\tau) d\tau + \mathbf{D}_i^a(c_{i,x}) \right) \cdot \mathbf{F}_i^a \vec{U}_i^a(0) \\
&= \mathbf{\Pi}_i^{a,(j+1)}(c_{i,x}) \cdot \mathbf{F}_i^a \vec{U}_i^a(0).
\end{aligned}$$

(b) $\mathbf{\Pi}_i^a(\Delta c_i) \cdot \mathbf{F}_i^a \vec{U}_i^a(0) = \left(\frac{\mathbf{\Pi}_i(\Delta c_i) \mathbf{F}_i \vec{U}_{i+1}(0) + \bar{\mathbf{\Pi}}_i^a(\Delta c_i) \vec{U}_0(0)}{\vec{U}_0(0)} \right) = \left(\frac{\vec{U}_i(0)}{\vec{U}_i'(0)} \right)$. Let $x = 0$ and we obtain

$$\vec{U}_i^a(0) = \mathbf{\Pi}_i^a(c_{i,0}) \cdot \mathbf{F}_i^a \vec{U}_i^a(0).$$

We can also write the above relation for $x = 0$ as:

$$\begin{aligned}
\left(\frac{\vec{U}_i(0)}{\vec{U}_i'(0)} \right) &= \mathbf{\Pi}_i^a(\Delta c_i) \left(\mathbf{F}_i' \mid \mathbf{B}_i' \right) \left(\frac{\vec{U}_{i+1}(0)}{\vec{U}_0(0)} \right) \\
&= \left(\frac{\mathbf{\Pi}_i(\Delta c_i) \mid \bar{\mathbf{\Pi}}_i^a(\Delta c_i)}{\mathbf{0} \mid \mathbf{I}} \right) \left(\frac{\mathbf{F}_i \mid \mathbf{0}}{\mathbf{0} \mid \mathbf{I}} \right) \left(\frac{\vec{U}_{i+1}(0)}{\vec{U}_0(0)} \right) \\
&= \left(\frac{\mathbf{\Pi}_i(\Delta c_i) \mathbf{F}_i \mid \bar{\mathbf{\Pi}}_i^a(\Delta c_i)}{\mathbf{0} \mid \mathbf{I}} \right) \left(\frac{\vec{U}_{i+1}(0)}{\vec{U}_0(0)} \right) \\
&= \left(\frac{\mathbf{\Pi}_i(\Delta c_i) \mathbf{F}_i \vec{U}_{i+1}(0) + \bar{\mathbf{\Pi}}_i^a(\Delta c_i) \vec{U}_0(0)}{\vec{U}_0(0)} \right).
\end{aligned}$$

As a result we can represent $\vec{U}_i(0)$ in the following matrix form

$$\vec{U}_i(0) = \mathbf{\Pi}_i(\Delta c_i) \mathbf{F}_i \vec{U}_{i+1}(0) + \bar{\mathbf{\Pi}}_i^a(\Delta c_i) \vec{U}_0(0)$$

by noting that $\mathbf{\Pi}_i$ is formed by the first k_i rows and columns of matrix $\mathbf{\Pi}_i^a$ and $\bar{\mathbf{\Pi}}_i^a$ is formed by the first k_i rows and the last $k_i^a - k_i$ columns of $\mathbf{\Pi}_i^a$.

(2) ($i = m$.) The proof of this case follows almost immediately from equation (4.15). As any region in \mathcal{G}_m is unbounded, delay transitions do not exist. As $\vec{U}_m(x+\tau)$ does not depend on x , the integral $\int_0^\infty \hat{\mathbf{M}}_m(\tau) \vec{U}_m(x+\tau) d\tau$ reduces to $\int_0^\infty \hat{\mathbf{M}}_m(\tau) d\tau \cdot \vec{U}_m(0)$. In addition, $\int_0^\infty \hat{\mathbf{M}}_m(\tau) d\tau$ boils down to $\hat{\mathbf{P}}_m$ and $\int_0^\infty \mathbf{B}_m(\tau) d\tau$ to $\hat{\mathbf{B}}_m$. \square

Since the coefficients of the linear equations are all known, solving the system of linear equations yields $\vec{U}_0(0)$, which contains the probability $Prob_{v_0}(0)$ of reaching V_F from initial vertex v_0 .

Theorem 4.9 is based on the equations (4.14) (for $i < m$), and (4.15) (for $i=m$). The term $\mathbf{\Pi}_i(\Delta c_i) \cdot \mathbf{F}_i \cdot \vec{U}_{i+1}(0)$ stands for the delay transitions, where \mathbf{F}_i specifies how the delay transitions are connected between the sub-graphs \mathcal{G}_i and \mathcal{G}_{i+1} . The term $\bar{\mathbf{\Pi}}_i^a(\Delta c_i) \cdot \vec{U}_0(0)$ stands for Markovian transitions with reset. The term $\bar{\mathbf{\Pi}}_i^a(\Delta c_i)$ in the augmented CTMC \mathcal{C}_i^a specifies the probabilities to first take transitions inside \mathcal{G}_i followed by a one-step Markovian transition back to \mathcal{G}_0 .

Remark 4.10. The approach in this section is focused on single-clock DTA (with finite acceptance criteria). For two-clock DTA $^\diamond$ the approach fails. In case of a single clock x , any reset (of x) from \mathcal{G}_i yields a state in $\mathcal{G}_0(0)$, and any delay (of x) yields some state in $\mathcal{G}_{i+1}(0)$. However, in the setting of two clocks, after a reset generally only one clock has a fixed value while the value of the other one is not determined.

Lemma 4.11. *For CTMC \mathcal{C} and single-clock DTA $^\diamond$ \mathcal{A} , computing $\Pr^{\mathcal{C}}(\text{Paths}^{\mathcal{C}}(\mathcal{A}))$ can be done in time $\mathcal{O}(m^2 \cdot |S| \cdot |Loc| \cdot \lambda \cdot \Delta c + m^3 \cdot |S|^3 \cdot |Loc|^3)$, where m is the number of constants appearing in \mathcal{A} , $|S|$ is the number of states in \mathcal{C} , $|Loc|$ is the number of locations in \mathcal{A} , λ is the maximal exit rate in \mathcal{C} and $\Delta c = \max_{0 \leq i < m} \{\Delta c_i\}$.*

Proof. The DMTA $\mathcal{C} \otimes \mathcal{A}$ has at most $|S| \cdot |Loc|$ locations. The number of vertices in the PDP $\mathcal{Z}(\mathcal{C} \otimes \mathcal{A})$ is at most $m \cdot |S| \cdot |Loc|$, as there are m possible regions. CTMC \mathcal{G}_i and its annotated version \mathcal{G}_i^a thus have at most $\mathcal{O}(m \cdot |S| \cdot |Loc|)$ states. Calculating the transient distribution $\Pi_i(\Delta c_i)$ on CTMC \mathcal{G}_i for any state in \mathcal{G}_i takes at most $\mathcal{O}(m \cdot |S| \cdot |Loc| \cdot \lambda \cdot \Delta c)$ where λ is the maximal exit rate in \mathcal{G}_i (and thus in \mathcal{C}) and $\Delta c = \max_{0 \leq i < m} \{\Delta c_i\}$ is the maximal width of a region. Given that this computation needs to be performed for any subgraph yields the first summand in the time complexity. Subsequently, according to Theorem 4.9, a system of linear equations has to be solved with at most $\mathcal{O}(m \cdot |S| \cdot |Loc|)$ variables. This takes at most $\mathcal{O}(m^3 \cdot |S|^3 \cdot |Loc|^3)$ operations. \square

5. VERIFYING CTMCS AGAINST MULLER DTA SPECIFICATIONS

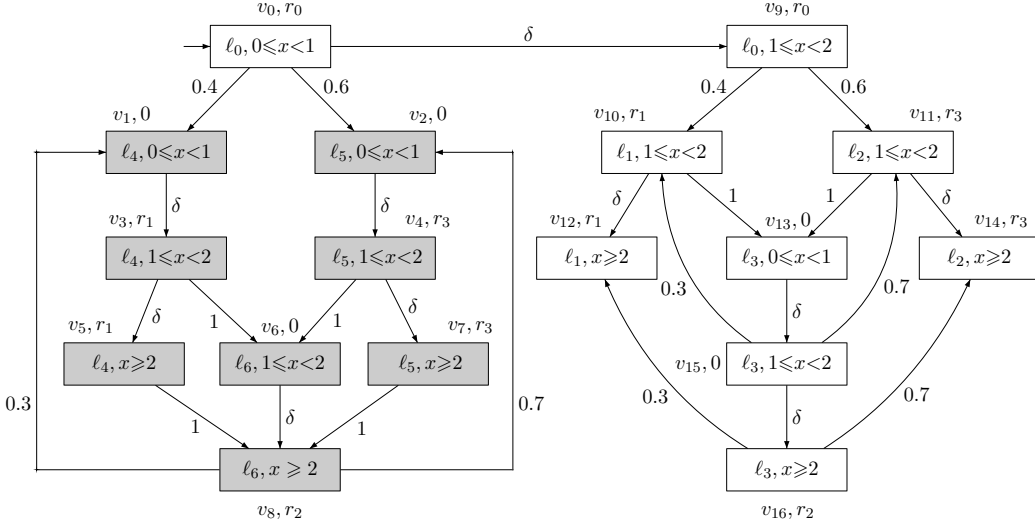
Finally, we deal with the verification of CTMCS against DTA with Muller acceptance conditions. The procedure is very similar to the one for DTA with finite acceptance conditions. Let \mathcal{A} be a DTA $^\omega$, and \mathcal{C} a CTMC. The region graph of the product $\mathcal{C} \otimes \mathcal{A}$ is defined as before (cf. Def. 3.11, page 18), except that the accepting set V_F is defined using bottom (or: terminal) SCCs (BSCCs for short). A strongly connected component (SCC) is terminal if it cannot be left once entered.

Definition 5.1 (Region graph of DMTA $^\omega$). The *region graph* of DMTA $^\omega$ $\mathcal{M} = (Loc, \mathcal{X}, \ell_0, Loc_{\mathcal{F}}, E, \rightsquigarrow)$ is $\mathcal{G}(\mathcal{M}) = (V, v_0, V_F, \Lambda, \hookrightarrow)$, where V, v_0, Λ and \hookrightarrow are defined as in Def. 3.11 (page 18), and $V_F = \{v \in B \mid B \in a\mathcal{B}\}$ where $a\mathcal{B}$ is the set of accepting BSCCs in $\mathcal{G}(\mathcal{M})$. BSCC $B \subseteq V$ is accepting if there exists $L_F \in Loc_{\mathcal{F}}$ such that for any $v \in B, v|_1 \in L_F$.

Example 5.2. Consider the DMTA $^\omega$ in Figure 6(c) with $Loc_{\mathcal{F}} = \{L_{F_1}, L_{F_2}\}$ with $L_{F_1} = \{\ell_1, \ell_2, \ell_3\}$, and $L_{F_2} = \{\ell_4, \ell_5, \ell_6\}$. Its region graph is depicted in Figure 11. There is one accepting BSCC, whose vertices are colored gray, corresponding to the set L_{F_2} . There is no BSCC corresponding to L_{F_1} , due to the presence of the sink vertices v_{12} and v_{14} . These vertices are reachable from locations ℓ_1 and ℓ_2 if $x \geq 2$.

Two remarks are in order. A first observation is that the probability to stay in an accepting BSCC is one, considering both the delay and Markovian transitions. That is to say, there are no outgoing transitions from which some probability can “leak away”. In addition, any pair of accepting BSCCs is disjoint, which allows the addition of, e.g., their reachability probabilities.

Theorem 5.3. *For any CTMC \mathcal{C} , DTA $^\omega$ \mathcal{A} , $\Pr^{\mathcal{C}}(\text{Paths}^{\mathcal{C}}(\mathcal{A}))$ is the least solution of $\text{Prob}_{v_0}^{\mathcal{D}}(\vec{0}, U)$, where DTMP $\mathcal{D} = \text{emb}(\mathcal{Z}(\mathcal{C} \otimes \mathcal{A}))$ and $U = \bigcup_{B \in a\mathcal{B}} B$.*

Figure 11: Region graph of the product DMTA^ω in Figure 6(c)

Proof. To start off, observe that $\text{Pr}^{\mathcal{C}}(\text{Paths}^{\mathcal{C}}(\mathcal{A}))$ is measurable, cf. Theorem 3.2 (page 10). The proof follows from Theorem 3.10 and the following observations. For any DTMP expanded with a finite set of locations—like for finite DTMCs—almost surely the states that are visited infinitely often along a path constitute a BSCC. It thus follows that the probability for visiting a set of states infinitely often equals the reachability probability of some BSCC in the DTMP $\text{emb}(\mathcal{Z}(\mathcal{C} \otimes \mathcal{A}))$. The result now follows from Theorem 4.3. \square

Example 5.4. Consider the region graph in Figure 11. The only BSCC is indicated by the gray shaded states. To determine $\text{Pr}^{\mathcal{C}}(\text{Paths}^{\mathcal{C}}(\mathcal{A}))$, it suffices to consider the reachability probability for $T = \{v_1, v_2\}$. For the delay transition $v_0 \xrightarrow{\delta} v_9$, we have

$$\text{Prob}_{v_0, \delta}(0) = e^{-r_0 \cdot 1} \cdot \text{Prob}_{v_9}(1) = e^{-r_0 \cdot 1} \cdot 0 = 0.$$

For the Markovian transition $v_0 \xrightarrow{0.4, \{x\}} v_1$,

$$\text{Prob}_{v_0, v_1}(0) = \int_0^1 0.4 \cdot r_0 \cdot e^{-r_0 \cdot \tau} \cdot \text{Prob}_{v_1}(\tau) d\tau = \int_0^1 0.4 \cdot r_0 \cdot e^{-r_0 \cdot \tau} d\tau.$$

A similar reasoning applies to $v_0 \xrightarrow{0.6, \{x\}} v_2$. Gathering the results we obtain:

$$\text{Pr}^{\mathcal{C}}(\text{Paths}^{\mathcal{C}}(\mathcal{A})) = \int_0^1 (0.4 + 0.6) \cdot r_0 \cdot e^{-r_0 \cdot \tau} d\tau = \int_0^1 r_0 \cdot e^{-r_0 \cdot \tau} d\tau = 1 - e^{-r_0}.$$

Verifying qualitative specifications. Until now we have investigated the quantitative verification problem, which is to determine the value of $\text{Pr}(\mathcal{C} \models \mathcal{A})$. The qualitative verification problem, on the other hand, is to determine whether the probability that \mathcal{C} satisfies \mathcal{A} exceeds zero, or, dually, equals one. For stochastic processes such as finite CTMCs and finite DTMCs, qualitative verification problems are known to be decidable by means of a simple graph analysis.

Proposition 5.5. *For any CTMC \mathcal{C} and DTA \mathcal{A} ,*

- (1) $\Pr^{\mathcal{C}}(\text{Paths}^{\mathcal{C}}(\mathcal{A})) > 0$ iff $\mathcal{Z}(\mathcal{C} \otimes \mathcal{A}) \models \exists \diamond V_F$,
- (2) $\Pr^{\mathcal{C}}(\text{Paths}^{\mathcal{C}}(\mathcal{A})) = 1$ iff $\mathcal{Z}(\mathcal{C} \otimes \mathcal{A}) \models \forall ((\exists \diamond V_F) \text{W } V_F)$,

where $V_F = \{v \in V \mid v \downarrow_1 \in \text{Loc}_F\}$ for DTA^{\diamond} , $V_F = \{v \in B \mid B \in a\mathcal{B}\}$ for DTA^{ω} , and W denotes the weak until operator.

Proof. Similar to the case for discrete-time Markov chains [8, Chapter 10]. □

From the above theorem, it follows that the qualitative properties can be verified using a standard graph-based CTL model checking algorithm, i.e., by just considering the underlying finite digraph of the PDP $\mathcal{Z}(\mathcal{C} \otimes \mathcal{A})$ —basically the region graph of $\mathcal{C} \otimes \mathcal{A}$ —while ignoring the transition probabilities.

6. CONCLUSION

This paper addressed the quantitative (and qualitative) verification of a finite CTMC \mathcal{C} against a linear real-time specification given as a deterministic timed automaton (DTA). We studied DTA with finite and Muller acceptance criteria. The key result (for finite acceptance) is that the probability of $\mathcal{C} \models \mathcal{A}$ equals the reachability probability in the embedded discrete-time Markov process of a PDP. This PDP is obtained via a standard region construction. Reachability probabilities in the thus obtained PDPs are characterized by a system of Volterra integral equations of the second type and are shown to be approximated by a system of PDEs. For Muller acceptance criteria, the probability of $\mathcal{C} \models \mathcal{A}$ equals the reachability probability of the accepting terminal SCCs in the embedded PDP. These results apply to DTA with arbitrarily (but finitely) many clocks. For single-clock DTA with finite acceptance, $\Pr(\mathcal{C} \models \mathcal{A})$ is obtained by solving a system of linear equations whose coefficients are solutions of a system of ODEs. As the coefficients are in fact transient probabilities in CTMCs, this result implies that standard algorithms for CTMC analysis suffice to verify single-clock DTA specifications.

An interesting future research direction is the verification against non-deterministic timed automata (NTA). NTA are strictly more expressive than DTA, and thus would allow more linear real-time specification. Following the approach in this paper requires a non-deterministic variant of PDP. Another challenging open problem is to consider real-time linear temporal logics as specifications such as metric temporal logic (MTL) [21] or variants thereof.

ACKNOWLEDGEMENT

We thank Jeremy Sproston (University of Turin) for fruitful and insightful discussions and Benoît Barbot (ENS Cachan) for pointing out some flaws in an earlier version of this paper. We are grateful to the reviewers for providing many useful suggestions on improving the presentation of the paper.

REFERENCES

- [1] R. Alur and D. L. Dill. A theory of timed automata. *Theor. Comput. Sci.*, 126(2):183–235, 1994.
- [2] G. B. Arfken and H. J. Weber. *Mathematical Methods for Physicists (4th ed.)*. Academic Press, 1995.
- [3] A. Aziz, K. Sanwal, V. Singhal, and R. K. Brayton. Model-checking continuous-time Markov chains. *ACM Trans. Comput. Log.*, 1(1):162–170, 2000.
- [4] C. Baier, N. Bertrand, P. Bouyer, T. Brihaye, and M. Größer. Probabilistic and topological semantics for timed automata. In *FSTTCS*, LNCS 4855, pages 179–191, 2007.
- [5] C. Baier, N. Bertrand, P. Bouyer, T. Brihaye, and M. Grösser. Almost-sure model checking of infinite paths in one-clock timed automata. In *LICS*, pages 217–226, 2008.
- [6] C. Baier, L. Cloth, B. R. Haverkort, M. Kuntz, and M. Siegle. Model checking Markov chains with actions and state labels. *IEEE Trans. Software Eng.*, 33(4):209–224, 2007.
- [7] C. Baier, B. R. Haverkort, H. Hermanns, and J.-P. Katoen. Model-checking algorithms for continuous-time Markov chains. *IEEE Trans. Software Eng.*, 29(6):524–541, 2003.
- [8] C. Baier and J.-P. Katoen. *Principles of Model Checking*. MIT Press, 2008.
- [9] B. Bérard, A. Petit, V. Diekert, and P. Gastin. Characterization of the expressive power of silent transitions in timed automata. *Fundam. Inform.*, 36(2-3):145–182, 1998.
- [10] N. Bertrand, P. Bouyer, T. Brihaye, and N. Markey. Quantitative model-checking of one-clock timed automata under probabilistic semantics. In *QEST*, pages 55–64, 2008.
- [11] T. Chen, T. Han, J.-P. Katoen, and A. Mereacre. Quantitative model checking of continuous-time Markov chains against timed automata specification. In *LICS*, pages 309–318, 2009.
- [12] C. Corduneanu. *Integral Equations and Applications*. Cambridge University Press, 1991.
- [13] O. L. Costa and M. H. Davis. Approximations for optimal stopping of a piecewise-deterministic process. *Math. Control Signals Systems*, 1(2):123–146, 1988.
- [14] C. Courcoubetis and M. Yannakakis. The complexity of probabilistic verification. *J. ACM*, 42(4):857–907, 1995.
- [15] M. H. A. Davis. Piecewise-deterministic Markov processes: A general class of non-diffusion stochastic models. *Journal of the Royal Statistical Society (B)*, 46(3):353–388, 1984.
- [16] M. H. A. Davis. *Markov Models and Optimization*. Chapman and Hall, 1993.
- [17] S. Donatelli, S. Haddad, and J. Sproston. Model checking timed and stochastic properties with CSL^{TA}. *IEEE Trans. Software Eng.*, 35(2):224–240, 2009.
- [18] H. Hansson and B. Jonsson. A logic for reasoning about time and reliability. *Formal Asp. Comput.*, 6(5):512–535, 1994.
- [19] A. Hinton, M. Z. Kwiatkowska, G. Norman, and D. Parker. PRISM: A tool for automatic verification of probabilistic systems. In *TACAS*, LNCS 3920, pages 441–444, 2006.
- [20] J.-P. Katoen, I. S. Zapreev, E. M. Hahn, H. Hermanns, and D. N. Jansen. The ins and outs of the probabilistic model checker MRMC. In *QEST*, pages 167–176, 2009.
- [21] R. Koymans. Specifying real-time properties with metric temporal logic. *Real-Time Systems*, 2(4):255–299, 1990.
- [22] M. Z. Kwiatkowska, G. Norman, R. Segala, and J. Sproston. Automatic verification of real-time systems with discrete probability distributions. *Theor. Comput. Sci.*, 282(1):101–150, 2002.
- [23] S. M. Lenhart and Y.-C. Liao. Integro-differential equations associated with optimal stopping time of a piecewise-deterministic process. *Stochastics*, 15(3):183–207, 1985.
- [24] S. M. Lenhart and N. Yamada. Perron’s method for viscosity solutions associated with piecewise-deterministic processes. *Funkcialaj Ekvacioj*, 34:173–186, 1991.
- [25] M. Y. Vardi. Automatic verification of probabilistic concurrent finite-state programs. In *FOCS*, pages 327–338, 1985.