2009 24th Annual IEEE Symposium on Logic In Computer Science

Quantitative Model Checking of Continuous-Time Markov Chains Against Timed Automata Specifications

Taolue Chen[1] Tingting Han[2,3] Joost-Pieter Katoen[2,3] Alexandru Mereacre[2]

[1]CWI, NL [2]MOVES, RWTH Aachen University, DE [3] FMT, University of Twente, NL

Abstract

We study the following problem: given a continuous-time Markov chain (CTMC) $\mathcal{C}$*, and a linear real-time property provided as a deterministic timed automaton* (DTA) $\mathcal{A}$*, what is the probability of the set of paths of* $\mathcal{C}$ *that are accepted by* $\mathcal{A}$ ($\mathcal{C}$ *satisfies* $\mathcal{A}$)*? It is shown that this set of paths is measurable and computing its probability can be reduced to computing the reachability probability in a piecewise deterministic Markov process* (PDP)*. The reachability probability is characterized as the least solution of a system of integral equations and is shown to be approximated by solving a system of partial differential equations. For the special case of single-clock* DTA*, the system of integral equations can be transformed into a system of linear equations where the coefficients are solutions of ordinary differential equations.*

1 Introduction

Continuous-time Markov chains (CTMCs) are one of the most important 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, more recently, calculi for system biology. CTMC model checking has been focused on the temporal logic CSL (Continuous Stochastic Logic [3, 7]), a variant of timed CTL where the CTL path quantifiers are replaced by a probabilistic operator. CSL model checking proceeds — like CTL model checking — by a recursive descent over the parse tree of the formula. One of the key ingredients is that reachability probabilities for time-bounded until-formulae can be approximated arbitrarily closely by a reduction to transient analysis in CTMCs. This results in a polynomial-time algorithm that has been realized in model checkers such as PRISM and MRMC.

This paper concerns the problem of verifying CTMCs versus *linear* real-time specifications, which are based on timed automata. Concretely speaking, we explore the following problem: given a CTMC $\mathcal{C}$, and a linear real-time property provided as a *deterministic timed automaton* [1] (DTA) $\mathcal{A}$, what is the probability of the set of paths of $\mathcal{C}$ which are accepted by $\mathcal{A}$ ($\mathcal{C} \models \mathcal{A}$)? We set off to show that this problem is well-defined in the sense that the path set is *measurable*. Computing its probability can then be reduced to computing the reachability probability in a piecewise deterministic Markov process (PDP) [12], a model that is used in, e.g., stochastic control theory and financial mathematics. This result relies on a product construction of CTMC $\mathcal{C}$ and DTA $\mathcal{A}$, denoted $\mathcal{C} \otimes \mathcal{A}$, yielding *deterministic Markov timed automata* (DMTA), 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. We show that the probability of $\mathcal{C} \models \mathcal{A}$ coincides with the reachability probability of accepting paths in $\mathcal{C} \otimes \mathcal{A}$. The underlying PDP of a DMTA is obtained by a slight adaptation of the standard region construction. The desired reachability probability is characterized as the least solution of a system of *integral equations* that is obtained from the PDP. Finally, this probability is shown to be approximated by solving a system of *partial differential equations* (PDEs). For single-clock DTA, we show that the system of integral equations can be transformed into a system of *linear equations*, where the coefficients are solutions of some *ordinary differential equations* (ODEs), which can either have an analytical solution (for small state space) or an arbitrarily closely approximated solution efficiently.

Related work is model checking of asCSL [6] and CSL^{TA} [13]. asCSL allows to impose a time constraint on action sequences described by regular expressions; its model-checking algorithm is based on a determin-

1043-6871/09 $25.00 © 2009 IEEE
DOI 10.1109/LICS.2009.21

istic Rabin automaton construction. In CSL$^{\mathrm{TA}}$, time constraints (of until modalities) are specified by *single-clock* DTA. In [13], $\mathcal{C}\otimes\mathcal{A}$ is interpreted as a Markov renewal processes and model checking CSL$^{\mathrm{TA}}$ is reduced to computing reachability probabilities in a DTMC whose transition probabilities are given by subordinate CTMCs. This technique cannot be generalized to multiple clocks. Our approach does not restrict the number of clocks and supports more specifications than CSL$^{\mathrm{TA}}$. For the single-clock case, our approach produces the same result as [13], but yields a conceptually simpler formulation whose correctness can be derived from the simplification of the system of integral equations obtained in the general case. Moreover, measurability has not been addressed in [13]. Other related work [4, 5, 9] 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 [9] for (a subclass of) single-clock timed automata.

The proofs can be found in the technical report [10].

2 Preliminaries

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

2.1 Continuous-time Markov chains

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

The probability to exit state s as well as to take the transition $s \to s'$ in t time units is $\int_0^t E(s)\cdot e^{-E(s)\tau}d\tau$ and $\mathbf{P}(s,s')\cdot\int_0^t E(s)\cdot e^{-E(s)\tau}d\tau$, respectively. 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., $\mathit{emb}(\mathcal{C}) = (S, \mathrm{AP}, L, \alpha, \mathbf{P})$.

Definition 2 [Timed paths] Let $\mathcal{C}$ be a CTMC. $\mathit{Paths}^{\mathcal{C}}_n := S \times (\mathbb{R}_{>0} \times S)^n$ is the set of paths of length n in $\mathcal{C}$; the set of finite paths in $\mathcal{C}$ is defined by $\mathit{Paths}^{\mathcal{C}}_\star = \bigcup_{n\in\mathbb{N}} \mathit{Paths}^{\mathcal{C}}_n$ and $\mathit{Paths}^{\mathcal{C}}_\omega := (S \times \mathbb{R}_{>0})^\omega$ is the set of infinite paths in $\mathcal{C}$. $\mathit{Paths}^{\mathcal{C}} = \mathit{Paths}^{\mathcal{C}}_\star \cup \mathit{Paths}^{\mathcal{C}}_\omega$ denotes the set of all paths in $\mathcal{C}$.

We denote a path $\rho \in \mathit{Paths}^{\mathcal{C}}(s_0)$ ($\rho \in \mathit{Paths}(s_0)$ 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 $n \leqslant |\rho|$ ($|\rho|$ 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 $t \in \mathbb{R}_{\geqslant 0}$, 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. that the time to stay in any state is strictly greater than 0.

The definition of a Borel space on paths through CTMCs follows [15, 7]. A CTMC $\mathcal{C}$ with initial state s_0 yields a probability measure $\Pr^{\mathcal{C}}$ on paths as follows: Let $s_0, \cdots, s_k \in S$ with $\mathbf{P}(s_i, s_{i+1}) > 0$ for $0 \leqslant i < k$ and $I_0, \cdots, I_{k-1}$ nonempty intervals in $\mathbb{R}_{\geqslant 0}$, $C(s_0, I_0, \cdots, I_{k-1}, s_k)$ denotes the *cylinder set* consisting of all paths $\rho \in \mathit{Paths}(s_0)$ such that $\rho[i] = s_i$ ($i \leqslant k$), and $\rho\langle i\rangle \in I_i$ ($i < k$). $\mathcal{F}(\mathit{Paths}(s_0))$ is the smallest σ-algebra on $\mathit{Paths}(s_0)$ which contains all sets $C(s_0, I_0, \cdots, I_{k-1}, s_k)$ for all state sequences $(s_0, \cdots, s_k) \in S^{k+1}$ with $\mathbf{P}(s_i, s_{i+1}) > 0$ ($0 \leqslant i < k$) and $I_0, \cdots, I_{k-1}$ range over all sequences of nonempty intervals in $\mathbb{R}_{\geqslant 0}$. The probability measure $\Pr^{\mathcal{C}}$ on $\mathcal{F}(\mathit{Paths}(s_0))$ is the unique measure defined by induction on k by $\Pr^{\mathcal{C}}(C(s_0)) = \alpha(s_0)$ and for $k > 0$:

$$\Pr^{\mathcal{C}}\big(C(s_0, I_0, \cdots, I_{k-1}, s_k)\big) = \Pr^{\mathcal{C}}\big(C(s_0, I_0, \cdots, I_{k-2}, s_{k-1})\big) \cdot \int_{I_{k-1}} \mathbf{P}(s_{k-1}, s_k) E(s_{k-1}) \cdot e^{-E(s_{k-1})\tau} d\tau. \quad (1)$$

Example 1 An example CTMC is illustrated in Fig. 2(b) (page 6), where $\mathrm{AP} = \{a, b, c\}$ 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 and transition probabilities are as shown.

2.2 Deterministic timed automata

(Clock) variables and valuations Let $\mathcal{X} = \{x_1, \ldots, x_n\}$ be a set of variables in $\mathbb{R}$. An $\mathcal{X}$-valuation is a function $\eta : \mathcal{X} \to \mathbb{R}$ assigning to each variable x a value $\eta(x)$. Let $\mathcal{V}(\mathcal{X})$ denote the set of all valuations over $\mathcal{X}$. 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 as $\eta \models g$ if $(\eta(x_1), \ldots, \eta(x_n)) \in g$.

Occasionally we use a special case of *nonnegative* variables, called *clocks*. We write $\vec{0}$ for 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 $\delta \in \mathbb{R}_{\geqslant 0}$, $\eta + \delta$ is the valuation η'' such that $\forall x \in \mathcal{X}.\ \eta''(x) := \eta(x) + \delta$, which implies that all clocks proceed at the same speed, or equivalently, $\forall x_i \in \mathcal{X}.\ \dot{x_i} = 1$. A *clock constraint* on $\mathcal{X}$ is an expression of the form $x \bowtie c$, or $x - y \bowtie c$, or the conjunction of any clock constraints, where $x, y \in \mathcal{X}$, $\bowtie \in \{<, \leqslant, >, \geqslant\}$ and $c \in \mathbb{N}$.

Definition 3 [DTA] A *deterministic timed automaton* (DTA) is a tuple $\mathcal{A} = (\Sigma, \mathcal{X}, Q, q_0, Q_F, \rightarrow)$ where Σ is a finite *alphabet*; $\mathcal{X}$ is a finite set of *clocks*; Q is a nonempty finite set of *locations*; $q_0 \in Q$ is the *initial location*; $Q_F \subseteq Q$ is a set of *accepting locations*; and $\rightarrow \in (Q \setminus Q_F) \times \Sigma \times \mathcal{B}(\mathcal{X}) \times 2^{\mathcal{X}} \times Q$ is an *edge relation*[1] satisfying: $q \xrightarrow{a,g,X} q'$ and $q \xrightarrow{a,g',X'} q''$ with $g \neq g'$ implies $g \cap g' = \varnothing$.

We refer to $q \xrightarrow{a,g,X} q'$ as an *edge*, where $a \in \Sigma$ is the input symbol, the *guard* g is a clock constraint on the clocks of $\mathcal{A}$, $X \subseteq \mathcal{X}$ is a set of clocks to be reset and q' is the successor location. The intuition is that the DTA $\mathcal{A}$ can move from location q to location q' when the input symbol is a and the guard g holds, while the clocks in X should be reset when entering q'. As a convention, we assume each location $q \in Q_F$ is a sink. An example DTA is shown in Fig. 2(c).

A finite timed path in $\mathcal{A}$ is of the form $\theta = q_0 \xrightarrow{a_0,t_0} q_1 \cdots q_n \xrightarrow{a_n,t_n} q_{n+1}$, for $t_i > 0$ $(0 \leqslant i \leqslant n)$. All the definitions on paths in CTMCs can be adapted here. A timed path θ is *accepted* by $\mathcal{A}$ if there exists some $0 \leqslant i \leqslant |\theta|$ such that $\theta[i] \in Q_F$ and for all $0 \leqslant j < i$, it holds that $\eta_0 = \vec{0}$, $\eta_j + t_j \models g_j$ and $\eta_{j+1} = (\eta_j + t_j)[X_j := 0]$, where η_j is the clock evaluation on *entering* q_j. We say that an infinite timed path $\rho = s_0 \xrightarrow{t_0} s_1 \xrightarrow{t_1} \cdots$ in CTMC $\mathcal{C}$ is accepted by $\mathcal{A}$ if there exists some $n \in \mathbb{N}$ such that the finite fragment of ρ, i.e. $s_0 \xrightarrow{t_0} s_1 \cdots s_{n-1} \xrightarrow{t_{n-1}} s_n$ gives rise to an augmented timed path $\hat{\rho} = q_0 \xrightarrow{L(s_0),t_0} q_1 \cdots q_{n-1} \xrightarrow{L(s_{n-1}),t_{n-1}} q_n$, which is accepted by $\mathcal{A}$.

2.3 Piecewise-deterministic Markov processes

PDPs constitute a general framework that can model virtually any stochastic system without diffusions [12] and for which powerful analysis and control techniques exist. A PDP consists of a finite set of *locations* each with a *location invariant* over a set of *variables*. A PDP can jump between locations either randomly, in which case the residence time of a location is governed by an exponential distribution, or when the location invariant is violated. While staying in a location, a PDP evolves *deterministically* according to a flow function (which is the solution of a system of ODEs). A *state* of the PDP consists of a location and a valuation of the variables. The target state of the jump is determined by a probability measure depending on the source state. The process is *Markovian* as the current state contains all the information to predict the future progress of the process.

Definition 4 [PDP [12]] A piecewise-deterministic (Markov) process (PDP) is a tuple $\mathcal{Z} = (Z, \mathcal{X}, Inv, \phi, \Lambda, \mu)$ with:

- Z and $\mathcal{X}$, a finite set of *locations* and *variables*, respectively;
- $Inv : Z \rightarrow \mathcal{B}(\mathcal{X})$, an *invariant function*;
- $\phi : Z \times \mathcal{V}(\mathcal{X}) \times \mathbb{R} \rightarrow \mathcal{V}(\mathcal{X})$, a *flow function*[2];
- $\Lambda : \mathbb{S} \rightarrow \mathbb{R}_{\geqslant 0}$, an *exit rate function*;
- $\mu : \mathring{\mathbb{S}} \cup \partial\mathbb{S} \rightarrow Distr(\mathbb{S})$, the *transition probability function*, where:

$\mathbb{S} := \{\xi := (z, \eta) \mid z \in Z, \eta \models Inv(z)\}$ is the state space of the PDP $\mathcal{Z}$, $\mathring{\mathbb{S}}$ is the interior of $\mathbb{S}$ and $\partial\mathbb{S} = \bigcup_{z \in Z} \{z\} \times \partial Inv(z)$ is the boundary of $\mathbb{S}$ with $\partial Inv(z) = \overline{Inv(z)} \setminus \mathring{Inv}(z)$ as the boundary of $Inv(z)$, $\mathring{Inv}(z)$ the interior of $Inv(z)$ and $\overline{Inv(z)}$ the closure of $Inv(z)$. Functions Λ and μ satisfy the following conditions:

- $\forall \xi \in \mathbb{S}.\ \exists \epsilon(\xi) > 0.$ function $t \mapsto \Lambda(\xi \oplus t)$ is integrable on $[0, \epsilon(\xi)[$, where $\xi \oplus t = \big(z, \phi(z, \eta, t)\big)$, for $\xi = (z, \eta)$;
- Function $\xi \mapsto \mu(\xi, A)$[3] is measurable for any $A \in \mathcal{F}(\mathbb{S})$, where $\mathcal{F}(\mathbb{S})$ is a σ-algebra generated by the countable union $\bigcup_{z \in Z} \{z\} \times A_z$ with A_z being a subset of $\mathcal{F}(Inv(z))$ and $\mu(\xi, \{\xi\}) = 0$.

A PDP is only allowed to stay in location z when the constraint $Inv(z)$ is satisfied. If e.g., $Inv(z)$ is $x_1^2 - 2x_2 \leqslant 1.5 \wedge x_3 > 2$, then its interior $\mathring{Inv}(z)$ is $x_1^2 - 2x_2 < 1.5 \wedge x_3 > 2$ and its closure $\overline{Inv(z)}$ is $x_1^2 - 2x_2 \leqslant 1.5 \wedge x_3 \geqslant 2$, and the boundary $\partial Inv(z)$ is $x_1^2 - 2x_2 = 1.5 \wedge x_3 = 2$. When the variable valuation satisfies the boundary $(\eta \models \partial Inv(z))$, the PDP is forced to jump and leave the current location z. The flow function ϕ defines the time-dependent behavior in a single location, in particular, how the variable valuations change when time elapses. State $\xi \oplus t$ is the timed successor of state ξ (on the same location) given

[1] N.B.: We don't consider diagonal constraints like $x - y \bowtie c$ in DTA. However, it is known that this does not harm the expressiveness of a TA [8].

[2] The flow function is assumed to be the solution of a system of ODEs with a Lipschitz continuous vector field.

[3] $\mu(\xi, A)$ is a shorthand for $(\mu(\xi))(A)$.

that t time units have passed. The PDP is piecewise-deterministic because in each location (one piece) the behavior is deterministically determined by ϕ. In summary, when a new state $\xi = (z, \eta)$ is entered and $Inv(z)$ is valid, i.e., $\xi \in \mathbb{S}$, the PDP can either *delay* to state $\xi' = (z, \eta') \in \mathbb{S} \cup \partial\mathbb{S}$ according to both the flow function ϕ and the time delay t (in this case $\xi' = \xi \oplus t$); or take a *Markovian jump* to state $\xi'' = (z'', \eta'') \in \mathbb{S}$ with probability $\mu(\xi, \{\xi''\})$. Note that the residence time of a location is exponentially distributed. When $Inv(z)$ is invalid, i.e., $\xi \in \partial\mathbb{S}$, ξ will be forced to take a *boundary jump* to ξ'' with probability $\mu(\xi, \{\xi''\})$.

The embedded *discrete-time Markov process* (DTMP) $emb(\mathcal{Z})$ of the PDP$\mathcal{Z}$ has the same state space $\mathbb{S}$ as $\mathcal{Z}$. The (one-jump) *transition probability* from a state ξ to a set $A \subseteq \mathbb{S}$ of states (on different locations as ξ), denoted $\hat{\mu}(\xi, A)$, is given by [12]:

$$\hat{\mu}(\xi, A) = \int_0^{\flat(\xi)} (\mathcal{Q}\mathbf{1}_A)(\xi \oplus t) \cdot \Lambda(\xi \oplus t)\, e^{-\int_0^t \Lambda(\xi \oplus \tau) d\tau} dt \quad (2)$$

$$+ \quad (\mathcal{Q}\mathbf{1}_A)(\xi \oplus \flat(\xi)) \cdot e^{-\int_0^{\flat(\xi)} \Lambda(\xi \oplus \tau) d\tau}, \quad (3)$$

where $\flat(\xi) = \inf\{t > 0 \mid \xi \oplus t \in \partial\mathbb{S}\}$ is the minimal time to hit the boundary if such time exists; $\flat(\xi) = \infty$ 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) 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 $[0, \flat(\xi)[$. Term (3) is the probability to stay in the same location for $\flat(\xi)$ time units and then it is forced to take a boundary jump from $\xi \oplus \flat(\xi)$ to A since $Inv(z)$ is invalid.

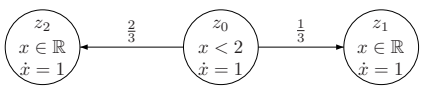

Figure 1. An example PDP

Example 2 Fig. 1 depicts a 3-location PDP $\mathcal{Z}$ with one variable x, where $Inv(z_0)$ is $x < 2$ and $Inv(z_1)$, $Inv(z_2)$ are both $x \in [0, \infty[$. Solving $\dot{x} = 1$ gives the flow function $\phi(z_i, \eta(x), t) = \eta(x) + t$ for $i = 0, 1, 2$. The state space of $\mathcal{Z}$ is $\{(z_0, \eta) \mid 0 < \eta(x) < 2\} \cup \{(z_1, \mathbb{R})\} \cup \{(z_2, \mathbb{R})\}$. Let exit rate $\Lambda(\xi) = 5$ for any $\xi \in \mathbb{S}$. For $\eta \models Inv(z_0)$, let $\mu\big((z_0, \eta), \{(z_1, \eta)\}\big) := \frac{1}{3}$, $\mu\big((z_0, \eta), \{(z_2, \eta)\}\big) := \frac{2}{3}$ and the boundary measure $\mu\big((z_0, 2), \{(z_1, 2)\}\big) := 1$. Given state $\xi_0 = (z_0, 0)$ and the set of states $A = (z_1, \mathbb{R})$, the time for ξ_0 to hit the boundary is $\flat(\xi_0) = 2$. Then $(\mathcal{Q}\mathbf{1}_A)(\xi_0 \oplus t) = \frac{1}{3}$ if $t < 2$, and $(\mathcal{Q}\mathbf{1}_A)(\xi_0 \oplus t) = 1$ if $t = 2$. In $emb(\mathcal{Z})$, the transition probability from state ξ_0 to A 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 Model checking DTA specifications

In this section, we deal with model checking linear real-time properties specified by DTA. The aim of model checking is to compute the probability of the set of paths in CTMC $\mathcal{C}$ accepted by a DTA $\mathcal{A}$. We prove that this can be reduced to computing the reachability probability in the product of $\mathcal{C}$ and $\mathcal{A}$ (Theorem 2), which can be further reduced to computing the reachability probability in a corresponding PDP (Theorem 3). To simplify the notations, we assume w.l.o.g. that a CTMC has only one initial state s_0, i.e., $\alpha(s_0) = 1$, and $\alpha(s) = 0$ for $s \neq s_0$.

3.1 Deterministic Markovian timed automata

To model check a DTA specification, we will exploit the product of a CTMC and a DTA, which is a *deterministic Markovian timed automaton*:

Definition 5 [DMTA] A *deterministic Markovian timed automaton* (DMTA) is a tuple $\mathcal{M} = (Loc, \mathcal{X}, \ell_0, Loc_F, E, \rightsquigarrow)$, where Loc is a finite set of *locations*; $\mathcal{X}$ is a finite set of *clocks*; $\ell_0 \in Loc$ is the *initial location*; $Loc_F \subseteq Loc$ is the set of *accepting locations*; $E : Loc \to \mathbb{R}_{\geqslant 0}$ is the *exit rate function*; and $\rightsquigarrow \subseteq Loc \times \mathcal{B}(\mathcal{X}) \times 2^{\mathcal{X}} \times Distr(Loc)$ is an *edge relation* satisfying $(\ell, g, X, \zeta), (\ell, g', X', \zeta') \in \rightsquigarrow$ with $g \neq g'$ implies $g \cap g' = \varnothing$.

The set of clocks $\mathcal{X}$ and the related concepts, e.g., clock valuation, clock constraints are defined as for DTA. We refer to $\ell \overset{g,X}{\rightsquigarrow} \zeta$ for distribution $\zeta \in Distr(Loc)$ as an *edge* and refer to $\ell \xrightarrow[\zeta(\ell')]{g,X} \ell'$ as a *transition* of this edge. The intuition is that when entering location ℓ, the DMTA chooses a residence time which is governed by the exponential distribution, i.e. the probability to leave ℓ in t time units is $1 - e^{-E(\ell)t}$. When it decides to jump, at most one edge, say $\ell \overset{g,X}{\rightsquigarrow} \zeta$, due to the determinism, is enabled and the probability to jump to ℓ' is given by $\zeta(\ell')$. The DMTA is *deterministic* as it has a unique initial location and disjoint guards for all edges emanating from any location.

Example 3 The DMTA in Fig. 2(a) has initial location ℓ_0 with two edges, with guards $x < 1$ and

$1 < x < 2$. Assume t time units elapsed. If $t < 1$, then the upper edge is enabled and the probability to go to ℓ_1 in time t is $(1 - e^{-r_0 t})\cdot 1$, where $E(\ell_0) = r_0$; no clock is reset. The process is similar for $1 < t < 2$, except that x will be reset. Location ℓ_3 is accepting.

Paths in DMTAs Given a DMTA $\mathcal{M}$ and a *finite symbolic path*

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

where $p_i = \zeta_i(\ell_{i+1})$ is the transition probability of $\ell_i \xrightarrow[\zeta_i(\ell_{i+1})]{g_i, X_i} \ell_{i+1}$, the induced *finite paths* in $\mathcal{M}$ are of the form $\sigma = \ell_0 \xrightarrow{t_0} \ell_1 \cdots \ell_{n-1} \xrightarrow{t_{n-1}} \ell_n$ and have the property that $\eta_0 = \vec{0}$, $(\eta_i + t_i) \models g_i$, and $\eta_{i+1} = (\eta_i + t_i)[X_i := 0]$ where $0 \leqslant i < n$ and η_i is the clock valuation of $\mathcal{X}$ in $\mathcal{M}$ on *entering* location ℓ_i. Finite path σ is *accepting* if $\ell_n \in Loc_F$. All definitions on paths in CTMCs can be carried over to DMTA paths.

Given DMTA $\mathcal{M}$, $C(\ell_0, I_0, \cdots, I_{n-1}, \ell_n)$ is the cylinder set where $(\ell_0, \cdots, \ell_n) \in Loc^{n+1}$ and $I_i \subseteq \mathbb{R}_{\geqslant 0}$. It denotes a set of paths σ in $\mathcal{M}$ such that $\sigma[i] = \ell_i$ and $\sigma\langle i \rangle \in I_i$. Now we define the measure $\Pr^{\mathcal{M}}_{\eta_0}$, which is the probability of $C(\ell_0, I_0, \cdots, I_{n-1}, \ell_n)$ such that the initial clock valuation in location ℓ_0 is η_0 as $\Pr^{\mathcal{M}}_{\eta_0}(C(\ell_0, I_0, \cdots, I_{n-1}, \ell_n)) := \mathbb{P}^{\mathcal{M}}_0(\eta_0)$. Here $\mathbb{P}^{\mathcal{M}}_i(\eta)$ for $0 \leqslant i \leqslant n$ is defined as: $\mathbb{P}^{\mathcal{M}}_n(\eta) = 1$ and for $0 \leqslant i < n$, we note that there exists a transition from ℓ_i to ℓ_{i+1} with $\ell_i \xrightarrow[p_i]{g_i, X_i} \ell_{i+1}$ $(0 \leqslant i < n)$ and thus we define

$$\mathbb{P}^{\mathcal{M}}_i(\eta) = \int_{I_i} \underbrace{\mathbf{1}_{g_i}(\eta + \tau) \cdot p_i \cdot E(\ell_i) \cdot e^{-E(\ell_i)\tau}}_{(\star)} \cdot \underbrace{\mathbb{P}^{\mathcal{M}}_{i+1}(\eta')}_{(\star\star)} \, d\tau,$$

where $\eta' := (\eta + \tau)[X_i := 0]$ and the *characteristic function* $\mathbf{1}_{g_i}(\eta + \tau) = 1$, if $\eta + \tau \models g_i$; 0, otherwise. Intuitively, $\mathbb{P}^{\mathcal{M}}_i(\eta_i)$ is the probability of the suffix cylinder set starting from ℓ_i and η_i to ℓ_n. It is recursively computed by the product of the probability of taking a transition from ℓ_i to ℓ_{i+1} in time interval I_i (cf. $(\star)$) and the probability of the suffix cylinder set from ℓ_{i+1} and η_{i+1} on (cf. $(\star\star)$), where $(\star)$ is computed by first ruling out the paths that do not belong to the cylinder set by $\mathbf{1}_{g_i}(\eta + \tau)$ and then computing the transition probability using the density function $p_i \cdot E(\ell_i) \cdot e^{-E(\ell_i)\tau}$ as in CTMCs. It follows that the characteristic function is Riemann integrable as it is bounded and its support is an interval, and thus $\mathbb{P}^{\mathcal{M}}_i(\eta)$ is well-defined.

3.2 Product DMTAs

Given a CTMC $\mathcal{C}$ and a DTA $\mathcal{A}$, the product $\mathcal{C} \otimes \mathcal{A}$ is a DMTA defined by:

Definition 6 [Product of CTMC and DTA] Let $\mathcal{C} = (S, \mathrm{AP}, L, s_0, \mathbf{P}, E)$ be a CTMC and $\mathcal{A} = (2^{\mathrm{AP}}, \mathcal{X}, Q, q_0, Q_F, \rightarrow)$ be a DTA. We define $\mathcal{C} \otimes \mathcal{A} = (Loc, \mathcal{X}, \ell_0, Loc_F, E, \rightsquigarrow)$ as the product DMTA, where $Loc := S \times Q$; $\ell_0 := \langle s_0, q_0 \rangle$; $Loc_F := S \times Q_F$; $E(\langle s, q \rangle) := E(s)$; and $\rightsquigarrow$ is defined as the smallest relation defined by the rule:

$$\frac{\mathbf{P}(s, s') > 0 \ \wedge \ q \xrightarrow{L(s), g, X} q', \ \text{s.t. } \zeta(\langle s', q' \rangle) = \mathbf{P}(s, s').}{\langle s, q \rangle \overset{g, X}{\rightsquigarrow} \zeta}$$

Example 4 Let CTMC $\mathcal{C}$ and DTA $\mathcal{A}$ be in Fig. 2(b) and 2(c), the product DMTA $\mathcal{C} \otimes \mathcal{A}$ is as in Fig. 2(a).

Remark 1 *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 $\mathcal{C} \otimes \mathcal{A}$, there is at most one "action" possible, viz. $L(s)$, from each location $\ell = \langle s, q \rangle$, probably via different edges, but with disjoint guards. We can thus omit it from the product* DMTA.

We denote $Paths^{\mathcal{C} \otimes \mathcal{A}}(\diamondsuit\, Loc_F) := \{\sigma \in Paths^{\mathcal{C} \otimes \mathcal{A}}_{\star} \mid \sigma \text{ is accepted by } \mathcal{C} \otimes \mathcal{A}\}$ as the set of accepted paths in $\mathcal{C} \otimes \mathcal{A}$, and $Paths^{\mathcal{C}}(\mathcal{A}) := \{\rho \in Paths^{\mathcal{C}}_{\star} \mid \rho \text{ is accepted by DTA } \mathcal{A}\}$ as the set of paths in CTMC $\mathcal{C}$ that are accepted by DTA $\mathcal{A}$. For any n-ary tuple J, let $J\lfloor_i$ denote the i-th entry in J, for $1 \leqslant i \leqslant n$. For a $\mathcal{C} \otimes \mathcal{A}$ path $\sigma = \langle s_0, q_0 \rangle \xrightarrow{t_0} \cdots \xrightarrow{t_{n-1}} \langle s_n, q_n \rangle$, let $\sigma\lfloor_1 := s_0 \xrightarrow{t_0} \cdots \xrightarrow{t_{n-1}} s_n$, and for any set Π of $\mathcal{C} \otimes \mathcal{A}$ paths, let $\Pi\lfloor_1 = \bigcup_{\sigma \in \Pi} \sigma\lfloor_1$.

Lemma 1 For any CTMC $\mathcal{C}$ and DTA $\mathcal{A}$, $Paths^{\mathcal{C}}(\mathcal{A}) = Paths^{\mathcal{C} \otimes \mathcal{A}}(\diamondsuit\, Loc_F)\lfloor_1$.

Theorem 1 For any CTMC $\mathcal{C}$ and DTA $\mathcal{A}$, $Paths^{\mathcal{C}}(\mathcal{A})$ is measurable.

We remark that the set of time-convergent paths in a CTMC has probability measure 0 (see [7]). The following theorem establishes the link between CTMC $\mathcal{C}$ and DMTA $\mathcal{C} \otimes \mathcal{A}$.

Theorem 2 For any CTMC $\mathcal{C}$ and DTA $\mathcal{A}$,

$$\Pr^{\mathcal{C}}\left(Paths^{\mathcal{C}}(\mathcal{A})\right) = \Pr^{\mathcal{C} \otimes \mathcal{A}}_{\vec{0}}\left(Paths^{\mathcal{C} \otimes \mathcal{A}}(\diamondsuit Loc_F)\right).$$

3.3 Region construction for DMTA

In the remainder of this section, we focus on how to compute the probability measure $\Pr^{\mathcal{C} \otimes \mathcal{A}}_{\vec{0}}\left(Paths^{\mathcal{C} \otimes \mathcal{A}}(\diamondsuit\, Loc_F)\right)$ in an effective way.

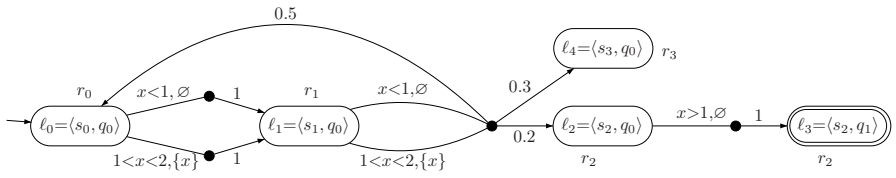

(a) DMTA $\mathcal{M} = \mathcal{C} \otimes \mathcal{A}$

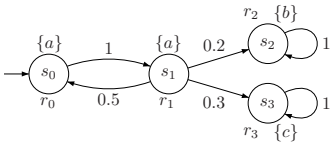

(b) CTMC $\mathcal{C}$

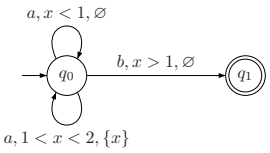

(c) DTA $\mathcal{A}$

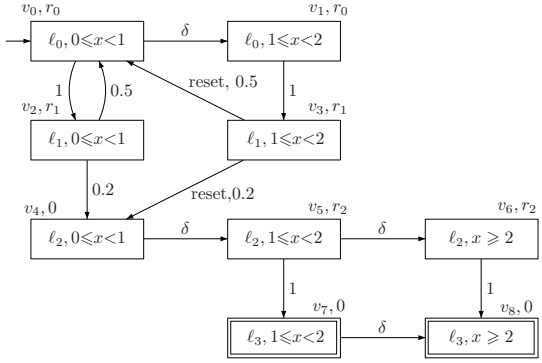

(d) Reachable region graph

Figure 2. Example product construction of CTMC $\mathcal{C}$ **and** DTA $\mathcal{A}$

We start with adopting the standard *region construction* [1] to DMTA. As we will see, this allows us to obtain a PDP from a DMTA in a natural way.

As usual, a region is a constraint. For regions $\Theta, \Theta' \in \mathcal{B}(\mathcal{X})$, 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 for all $\delta' < \delta$, $\eta + \delta' \models \Theta \vee \Theta'$. A region Θ *satisfies* a guard g (denoted $\Theta \models g$) iff $\forall \eta \models \Theta.\ \eta \models g$. A *reset operation* on region Θ is defined as $\Theta[X := 0] := \{\eta[X := 0] \mid \eta \models \Theta\}$.

Definition 7 [Region graph of DMTA] Given DMTA $\mathcal{M} = (Loc, \mathcal{X}, \ell_0, Loc_F, E, \rightsquigarrow)$, the region graph $\mathcal{G}(\mathcal{M}) = (V, v_0, V_F, \Lambda, \hookrightarrow)$ with $V := Loc \times \mathcal{B}(\mathcal{X})$ is a finite set of *vertices*, consisting of a location ℓ in $\mathcal{M}$ and a region Θ; $v_0 \in V$ is the *initial vertex* if $(\ell_0, \vec{0}) \in v_0$; $V_F := \{v \mid v\!\downarrow_1 \in Loc_F\}$ is the set of *accepting vertices*; $\hookrightarrow \subseteq V \times \big(([0,1] \times 2^{\mathcal{X}}) \cup \{\delta\}\big) \times V$ is the *transition (edge) relation*, such that:

- $v \overset{\delta}{\hookrightarrow} v'$ is a delay transition if $v\!\downarrow_1 = v'\!\downarrow_1$ and $v'\!\downarrow_2$ is a successor region of $v\!\downarrow_2$;
- $v \overset{p,X}{\hookrightarrow} v'$ is a Markovian transition if there exists some transition $v\!\downarrow_1 \overset{g,X}{\underset{p}{\longmapsto}} v'\!\downarrow_1$ in $\mathcal{M}$ such that $v\!\downarrow_2 \models g$ and $v\!\downarrow_2[X := 0] \models v'\!\downarrow_2$; and

$\Lambda : V \to \mathbb{R}_{\geqslant 0}$ is the *exit rate function* where $\Lambda(v) := E(v\!\downarrow_1)$ if there exists a Markovian transition from v, 0 otherwise.

Note that in the obtained region graph, Markovian transitions emanating from any boundary region do *not* contribute to the reachability probability as the time to hit the boundary is always zero (cf. (5)). Therefore, we can remove all the Markovian transitions emanating from boundary regions and then collapse each of them with its unique *non-boundary* (direct) successor. In the sequel we still denote this *collapsed* region graph $\mathcal{G}(\mathcal{M})$ by slightly abusing the notation.

We now define the underlying PDP of a DMTA by using the region graph $\mathcal{G}(\mathcal{M})$:

Definition 8 [PDP for DMTA] For DMTA $\mathcal{M} = (Loc, \mathcal{X}, \ell_0, Loc_F, E, \rightsquigarrow)$ and region graph $\mathcal{G}(\mathcal{M}) = (V, v_0, V_F, \Lambda, \hookrightarrow)$, let PDP $\mathcal{Z}(\mathcal{M}) = (V, \mathcal{X}, Inv, \phi, \Lambda, \mu)$ where for any $v \in V$,

- $Inv(v) := v\!\downarrow_2$ and the state space $\mathbb{S} := \{(v, \eta) \mid v \in V, \eta \in Inv(v)\}$;
- $\phi(v, \eta, t) := \eta + t$ for $\eta \models Inv(v)$;
- $\Lambda(v, \eta) := \Lambda(v)$ is the exit rate of state (v, η);
- [*boundary jump*] for each delay transition $v \overset{\delta}{\hookrightarrow} v'$ in $\mathcal{G}(\mathcal{M})$ we have $\mu(\xi, \{\xi'\}) := 1$, where $\xi = (v, \eta)$, $\xi' = (v', \eta)$ and $\eta \models \partial Inv(v)$;

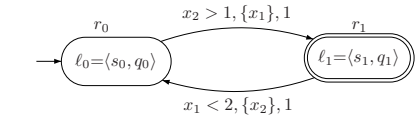

(a) DMTA $\mathcal{M} = \mathcal{C} \otimes \mathcal{A}$

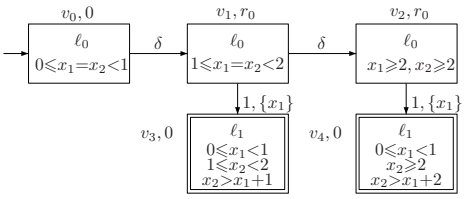

(b) Reachable region graph $\mathcal{G}(\mathcal{C} \otimes \mathcal{A})$

Figure 3. Example of a region graph

- [*Markovian jump*] for each Markovian transition $v \overset{p,X}{\hookrightarrow} v'$ in $\mathcal{G}(\mathcal{M})$ we have $\mu(\xi, \{\xi'\}) := p$, where $\xi = (v, \eta)$, $\eta \models \mathit{Inv}(v)$ and $\xi' = (v', \eta[X := 0])$.

From now on we write $\Lambda(v)$ instead of $\Lambda(v, \eta)$ as they coincide.

Example 5 For the DMTA $\mathcal{C} \otimes \mathcal{A}$ in Fig. 3(a), the reachable part (forward reachable from the initial vertex and backward reachable from the accepting vertices) of the collapsed region graph $\mathcal{G}(\mathcal{C} \otimes \mathcal{A})$ is in Fig. 3(b). The accepting vertices are sinks.

3.4 Characterizing reachability probabilities

Computing $\Pr_{\vec{0}}^{\mathcal{C}\otimes\mathcal{A}}\Big(\mathit{Paths}^{\mathcal{C}\otimes\mathcal{A}}(\Diamond\, \mathit{Loc}_F)\Big)$ is now reduced to computing the (time-unbounded) reachability probability in the PDP $\mathcal{Z}(\mathcal{C} \otimes \mathcal{A})$, given the initial state $(v_0, \vec{0})$ and the set of goal states $\{(v, \eta) \mid v \in V_F, \eta \in \mathit{Inv}(v)\}$ $\big((V_F, \cdot)$ for short$\big)$. Reachability probabilities of untimed events in a PDP $\mathcal{Z}$ can be computed in the embedded DTMP $\mathit{emb}(\mathcal{Z})$. Note that the set of locations of $\mathcal{Z}$ and $\mathit{emb}(\mathcal{Z})$ are equal. In the sequel, let $\mathcal{D}$ denote $\mathit{emb}(\mathcal{Z})$.

For each vertex $v \in V$, we define recursively $\mathit{Prob}^{\mathcal{D}}\big((v, \eta), (V_F, \cdot)\big)$ $\big($or shortly $\mathit{Prob}^{\mathcal{D}}_v(\eta)\big)$ as the probability to reach the goal states $(V_F, \cdot)$ in $\mathcal{D}$ from state (v, η).

– for the delay transition $v \overset{\delta}{\hookrightarrow} v'$,

$$\mathit{Prob}^{\mathcal{D}}_{v,\delta}(\eta) = e^{-\Lambda(v)\flat(v,\eta)} \cdot \mathit{Prob}^{\mathcal{D}}_{v'}\big(\eta + \flat(v, \eta)\big). \quad (4)$$

Recall that $\flat(v, \eta)$ is the minimal time for (v, η) to hit the boundary $\partial \mathit{Inv}(v)$.

– for the Markovian transition $v \overset{p,X}{\hookrightarrow} v'$,

$$\mathit{Prob}^{\mathcal{D}}_{v,v'}(\eta) = \int_0^{\flat(v,\eta)} p{\cdot}\Lambda(v){\cdot}e^{-\Lambda(v)\tau}{\cdot}\mathit{Prob}^{\mathcal{D}}_{v'}\big((\eta{+}\tau)[X := 0]\big)\, d\tau. \quad (5)$$

Overall, for each vertex $v \in V$, we obtain:

$$\mathit{Prob}^{\mathcal{D}}_v(\eta) = \begin{cases} \mathit{Prob}^{\mathcal{D}}_{v,\delta}(\eta) + \sum_{v \overset{p,X}{\hookrightarrow} v'} \mathit{Prob}^{\mathcal{D}}_{v,v'}(\eta), & \text{if } v \notin V_F \\ 1, & \text{otherwise} \end{cases}. \quad (6)$$

Note that here the notation η is slightly abused. It represents a vector of clock variables (see Example 6). Eq. (4) and (5) are derived based on (3) and (2), respectively. In particular the multi-step reachability probability is computed using a sequence of one-step transition probabilities.

Hence we obtain a system of *integral equations* (6). One can read (6) either in the form $f(\xi) = \int_{\mathit{Dom}(\xi)} K(\xi, \xi') f(d\xi')$, where K is the kernel and $\mathit{Dom}(\xi)$ 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, (6) does *not* necessarily have a unique solution. It turns out that the reachability probability $\mathit{Prob}^{\mathcal{D}}_{v_0}(\vec{0})$ coincides with the least fixpoint of the operator $\mathcal{J}$ (denoted by $\mathrm{lfp}\mathcal{J}$) i.e., $\mathit{Prob}^{\mathcal{D}}_{v_0}(\vec{0}) = (\mathrm{lfp}\mathcal{J})(v_0, \vec{0})$. Formally, we have:

Theorem 3 For any CTMC $\mathcal{C}$ and DTA $\mathcal{A}$, $\Pr_{\vec{0}}^{\mathcal{C}\otimes\mathcal{A}}\big(\mathit{Paths}^{\mathcal{C}\otimes\mathcal{A}}(\Diamond\, \mathit{Loc}_F)\big)$ is the least solution of $\mathit{Prob}^{\mathcal{D}}_{v_0}(\cdot)$, where $\mathcal{D}$ is the embedded DTMP of $\mathcal{C} \otimes \mathcal{A}$.

Remark 2 *Clock valuations η and η' in region Θ may induce different reachability probabilities. The reason is 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 the traditional timed automata theory as well as probabilistic timed automata [14], where η and η' are not distinguished.*

Example 6 For the region graph in Fig. 3(b), the system of integral equations for v_1 in location ℓ_0 is as follows for $1 \leqslant x_1 = x_2 < 2$:

$$\mathit{Prob}^{\mathcal{D}}_{v_1}(x_1, x_2) = \mathit{Prob}^{\mathcal{D}}_{v_1,\delta}(x_1, x_2) + \mathit{Prob}^{\mathcal{D}}_{v_1,v_3}(x_1, x_2),$$

where $\mathit{Prob}^{\mathcal{D}}_{v_1,\delta}(x_1, x_2) = e^{-(2-x_1)r_0}{\cdot}\mathit{Prob}^{\mathcal{D}}_{v_2}(2, 2)$ and $\mathit{Prob}^{\mathcal{D}}_{v_1,v_3}(x_1, x_2) = \int_0^{2-x_1} r_0{\cdot}e^{-r_0\tau}{\cdot}\mathit{Prob}^{\mathcal{D}}_{v_3}(0, x_2{+}\tau)\, d\tau$ where $\mathit{Prob}^{\mathcal{D}}_{v_3}(0, x_2 + \tau) = 1$. The integral equations for v_2 can be derived in a similar way.

3.5 Approximating reachability probabilities

Finally, we discuss how to obtain a solution of (6). The integral equations (6) are *Volterra equations of the*

second type [2]. For a general reference on solutions to Volterra equations, cf., e.g. [11]. As an alternative option to solve (6), we proceed to give a general formulation of $\Pr^{\mathcal{C}}\left(Paths^{\mathcal{C}}(\mathcal{A})\right)$ using a system of *partial differential equations* (PDEs). Let the *augmented* DTA $\mathcal{A}[t_f]$ be obtained from $\mathcal{A}$ by adding a new clock variable y which is never reset and a clock constraint $y < t_f$ on all edges entering the accepting locations in Loc_F, where t_f is a finite (and usually very large) integer. The purpose of this augmentation is to ensure that the value of all clocks reaching Loc_F is bounded. It is clear that $Paths^{\mathcal{C}}(\mathcal{A}[t_f]) \subseteq Paths^{\mathcal{C}}(\mathcal{A})$. More precisely, $Paths^{\mathcal{C}}(\mathcal{A}[t_f])$ coincides with those paths which can reach the accepting states of $\mathcal{A}$ within the time bound t_f. Note that $\lim_{t_f\to\infty}\Pr^{\mathcal{C}}(Paths^{\mathcal{C}}(\mathcal{A}[t_f])) = \Pr^{\mathcal{C}}(Paths^{\mathcal{C}}(\mathcal{A}))$. We can approximate $\Pr^{\mathcal{C}}(Paths^{\mathcal{C}}(\mathcal{A}))$ by solving the PDEs with a large t_f as follows:

Proposition 1 Given a CTMC $\mathcal{C}$, an augmented DTA $\mathcal{A}[t_f]$ and the underlying PDP $\mathcal{Z}(\mathcal{C}\otimes\mathcal{A}[t_f]) = (V,\mathcal{X},Inv,\phi,\Lambda,\mu)$, $\Pr^{\mathcal{C}}\left(Paths^{\mathcal{C}}(\mathcal{A}[t_f])\right) = \hbar_{v_0}(0,\vec{0})$ (which is the probability to reach the final states in $\mathcal{Z}$ starting from initial state $(v_0,\vec{0}_{\mathcal{X}\cup\{y\}}{}^{4})$) is the unique solution of the following system of PDEs:

$$\frac{\partial\hbar_v(y,\eta)}{\partial y}+\sum_{i=1}^{|\mathcal{X}|}\frac{\partial\hbar_v(y,\eta)}{\partial\eta^{(i)}}+$$
$$\Lambda(v)\cdot\sum_{v\overset{p,X}{\hookrightarrow}v'}p\cdot(\hbar_{v'}(y,\eta[X:=0])-\hbar_v(y,\eta))=0,$$

where $v\in V\setminus V_F$, $\eta\models Inv(v)$, $\eta^{(i)}$ is the i'th clock variable and $y\in[0,t_f[$. For every $\eta\models\partial Inv(v)$ and transition $v\overset{\delta}{\hookrightarrow}v'$, the boundary conditions take the form: $\hbar_v(y,\eta)=\hbar_{v'}(y,\eta)$. For every vertex $v\in V_F$, $\eta\models Inv(v)$ and $y\in[0,t_f[$, we have the following PDE:

$$\frac{\partial\hbar_v(y,\eta)}{\partial y}+\sum_{i=1}^{|\mathcal{X}|}\frac{\partial\hbar_v(y,\eta)}{\partial\eta^{(i)}}+1=0.$$

The final boundary conditions are that for every vertex $v\in V$ and $\eta\models Inv(v)\cup\partial Inv(v)$, $\hbar_v(t_f,\eta)=0$.

4 Single-clock DTA specifications

For single-clock DTA specifications, we can simplify the system of integral equations obtained in the previous section to a system of *linear* equations where the coefficients are a solution of a system of ODEs that can be calculated efficiently.

Given a DMTA $\mathcal{M}$, we denote the set of constants appearing in the clock constraints of $\mathcal{M}$ as $\{c_0,\ldots,c_m\}$ with $c_0=0$. We assume the following order: $0=c_0<c_1<\cdots<c_m$. Let $\Delta c_i=c_{i+1}-c_i$ for $0\leqslant i<m$. Note that for one clock DMTA, the regions in the region graph $\mathcal{G}(\mathcal{M})$ can be represented by the following intervals: $[c_0,c_1),\ldots,[c_m,\infty)$. We partition the region graph $\mathcal{G}(\mathcal{M})=(V,v_0,V_F,\Lambda,\hookrightarrow)$, or $\mathcal{G}$ for short, into a set of subgraphs $\mathcal{G}_i=(V_i,V_{Fi},\Lambda_i,\{M_i,F_i,B_i\})$, where $0\leqslant i\leqslant m$ and $\Lambda_i(v)=\Lambda(v)$, if $v\in V_i$, 0 otherwise. These subgraphs are obtained by partitioning V, V_F and $\hookrightarrow$ as follows:

- $V=\bigcup_{0\leqslant i\leqslant m}\{V_i\}$, where $V_i=\{(\ell,\Theta)\in V\mid\Theta\subseteq[c_i,c_{i+1})\}$;
- $V_F=\bigcup_{0\leqslant i\leqslant m}\{V_{Fi}\}$, where $v\in V_{Fi}$ iff $v\in V_i\cap V_F$;
- $\hookrightarrow=\bigcup_{0\leqslant i\leqslant m}\{M_i\cup F_i\cup B_i\}$, where M_i is the set of *Markovian transitions (without reset)* between vertices inside $\mathcal{G}_i$; F_i is the set of *delay transitions* from the vertices in $\mathcal{G}_i$ to that in $\mathcal{G}_{i+1}$ (Forward) and B_i is the set of *Markovian transitions (with reset)* from $\mathcal{G}_i$ to $\mathcal{G}_0$ (Backward). It is easy to see that M_i, F_i, and B_i are pairwise disjoint.

Since the initial vertex of $\mathcal{G}_0$ is v_0 and the initial vertices of $\mathcal{G}_i$ for $0<i\leqslant m$ are implicitly given by $\mathbf{F}_{i-1}$, we omit them in the definition. As an example, the vertices in Fig. 4 are partitioned by the ovals and the M_i edges are unlabeled while the F_i and B_i edges are labeled with δ and "reset", respectively. The V_F vertices (double circles) may appear in any $\mathcal{G}_i$. Actually, if $v=(\ell,[c_i,c_{i+1}))\in V_F$, then $v'=(\ell,[c_j,c_{j+1}))\in V_F$ for $i<j\leqslant m$. This is true because $V_F=\{(\ell,\mathsf{true})\mid\ell\in Loc_F\}$. It implies that for each final vertex not in the last region, there is a delay transition from it to the next region, see the final vertex in $\mathcal{G}_{i+1}$ in Fig. 4. The exit rate functions and the probabilities on Markovian edges are omitted in the graph.

Given a subgraph $\mathcal{G}_i$ ($0\leqslant i\leqslant m$) of $\mathcal{G}$ with k_i states, let the probability vector $\vec{U}_i(x)=[u_i^1(x),\ldots,u_i^{k_i}(x)]^\top\in\mathbb{R}^{k_i\times1}$ where $u_i^j(x)$ is the probability to go from vertex $v_i^j\in V_i$ to vertices in V_F (in $\mathcal{G}$) at time x. Starting from (4-6), we provide a set of integral equations for $\vec{U}_i(x)$ which we later on reduce to a system of linear equations. Distinguish two cases:
Case $0\leqslant i<m$**:** $\vec{U}_i(x)$ is given by:

$$\vec{U}_i(x) = \int_0^{\Delta c_i-x}\mathbf{M}_i(\tau)\vec{U}_i(x+\tau)d\tau \quad (7)$$
$$+ \int_0^{\Delta c_i-x}\mathbf{B}_i(\tau)d\tau\cdot\vec{U}_0(0) \quad (8)$$
$$+ \mathbf{D}_i(\Delta c_i-x)\cdot\mathbf{F}_i\vec{U}_{i+1}(0), \quad (9)$$

where $x\in[0,\Delta c_i]$ and

[4] denoting the valuation η with $\eta(x)=0$ for $x\in\mathcal{X}\cup\{y\}$.

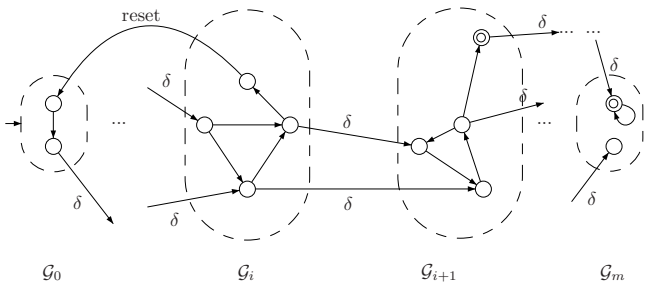

Figure 4. Partitioning the region graph

- $\mathbf{D}_i(x) \in \mathbb{R}^{k_i \times k_i}$ is the delay probability matrix, where for any $0 \leqslant j \leqslant k_i$, $\mathbf{D}_i(x)[j,j] = e^{-E(v_i^j)x}$ (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 the Markovian transitions inside $\mathcal{G}_i$, where $\mathbf{P}_i$ and $\mathbf{E}_i$ are the transition probability matrix and exit rate matrix for vertices inside $\mathcal{G}_i$, respectively;
- $\mathbf{B}_i(x) \in \mathbb{R}^{k_i \times k_0}$ is the probability density matrix for the reset edges B_i, where $\mathbf{B}_i(x)[j,j']$ indicates the probability density function to take the Markovian jump with reset from the j-th vertex in $\mathcal{G}_i$ to the j'-th vertex in $\mathcal{G}_0$; and
- $\mathbf{F}_i \in \mathbb{R}^{k_i \times k_{i+1}}$ is the incidence matrix for delay edges F_i. More specifically, $\mathbf{F}_i[j,j'] = 1$ indicates that there is a delay transition from the j-th vertex in $\mathcal{G}_i$ to the j'-th vertex in $\mathcal{G}_{i+1}$; 0 otherwise.

Let us explain these equations. Eq. (9) is obtained from (4) where $\mathbf{D}_i(\Delta c_i - x)$ indicates 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). In a similar way, (7) and (8) are obtained from (5); the former reflects the case where clock x is not reset, while the latter considers the reset of x (and returning to $\mathcal{G}_0$).

Case $i = m$: $\vec{U}_m(x)$ is simplified as follows:

$$\vec{U}_m(x) = \int_0^\infty \hat{\mathbf{M}}_m(\tau)\vec{U}_m(x+\tau)d\tau + \vec{1}_F + \int_0^\infty \mathbf{B}_m(\tau)d\tau \cdot \vec{U}_0(0) \tag{10}$$

where $\hat{\mathbf{M}}_m(\tau)[v,\cdot] = \mathbf{M}_m(\tau)[v,\cdot]$ for $v \notin V_F$, 0 otherwise. $\vec{1}_F$ is a vector such that $\vec{1}_F[v] = 1$ if $v \in V_F$, 0 otherwise. We note that $\vec{1}_F$ stems from the second clause of (6), and $\hat{\mathbf{M}}_m$ is obtained by setting the corresponding elements of $\mathbf{M}_m$ to 0. Also note that as the last subgraph $\mathcal{G}_m$ involves infinite regions, it has no delay transitions.

Before solving the system of integral equations (7-10), we first make the following observations:

(i) Due to the fact that inside $\mathcal{G}_i$ there are only Markovian jumps with neither resets nor delay transitions, $\mathcal{G}_i$ with (V_i, Λ_i, M_i) forms a CTMC $\mathcal{C}_i$, say. For each $\mathcal{G}_i$ we define an *augmented* CTMC $\mathcal{C}_i^a$ with state space $V_i \cup V_0$, such that all V_0-vertices are made absorbing in $\mathcal{C}_i^a$. The edges connecting V_i to V_0 are kept and all the edges inside $\mathcal{C}_0$ are removed. The augmented CTMC is used to calculate the probability to start from a vertex in $\mathcal{G}_i$ and take a reset edge in a certain time.

(ii) Given any CTMC $\mathcal{C}$ with k states and rate matrix $\mathbf{P} \cdot \mathbf{E}$, the matrix $\mathbf{\Pi}(x)$ is given by:

$$\mathbf{\Pi}(x) = \int_0^x \mathbf{M}(\tau)\mathbf{\Pi}(x-\tau)d\tau + \mathbf{D}(x). \tag{11}$$

Intuitively, $\mathbf{\Pi}(t)[j,j']$ indicates the probability to start from vertex j and reach j' at time t.

The following proposition states the close relationship between $\mathbf{\Pi}(x)$ and the transient probability vector:

Proposition 2 Given a CTMC $\mathcal{C}$ with initial distribution α, rate matrix $\mathbf{P} \cdot \mathbf{E}$ and $\mathbf{\Pi}(t)$, $\vec{\pi}(t)$ satisfies the following two equations:

$$\vec{\pi}(t) = \alpha \cdot \mathbf{\Pi}(t), \tag{12}$$

$$\frac{d\vec{\pi}(t)}{dt} = \vec{\pi}(t) \cdot \mathbf{Q}, \tag{13}$$

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

$\vec{\pi}(t)$ is the *transient probability vector* with $\vec{\pi}(t)[s]$ indicating the probability to be in state s at time t given the initial probability distribution α. Eq. (13) is the celebrated forward Chapman-Kolmogorov equations. According to this proposition, solving the integral equation $\mathbf{\Pi}(t)$ boils down to selecting the appropriate initial distribution vector α and solving the system of ODEs (13), which can be done very efficiently using the *uniformization technique*.

Prior to exposing how to solve the system of integral equations by solving a system of *linear* equations, we define $\bar{\mathbf{\Pi}}_i^a \in \mathbb{R}^{k_i \times k_0}$ for an augmented CTMC $\mathcal{C}_i^a$ to be part of $\mathbf{\Pi}_i^a$, where $\bar{\mathbf{\Pi}}_i^a$ only keeps the probabilities starting from V_i and ending in V_0. Actually,

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

where $\mathbf{0} \in \mathbb{R}^{k_0 \times k_i}$ is the matrix with all elements zero and $\mathbf{I} \in \mathbb{R}^{k_0 \times k_0}$ is the identity matrix.

Theorem 4 For subgraph $\mathcal{G}_i$ of $\mathcal{G}$ with k_i states, it holds for $0 \leqslant i < m$ that:

$$\vec{U}_i(0) = \mathbf{\Pi}_i(\Delta c_i) \cdot \mathbf{F}_i \vec{U}_{i+1}(0) + \bar{\mathbf{\Pi}}_i^a(\Delta c_i) \cdot \vec{U}_0(0), \tag{14}$$

where $\mathbf{\Pi}_i(\Delta c_i)$ and $\bar{\mathbf{\Pi}}^a_i(\Delta c_i)$ are for CTMC $\mathcal{C}_i$ and the augmented CTMC $\mathcal{C}^a_i$, respectively. For case $i = m$,

$$\vec{U}_m(0) = \hat{\mathbf{P}}_i \cdot \vec{U}_m(0) + \vec{1}_F + \hat{\mathbf{B}}_m \cdot \vec{U}_0(0), \qquad (15)$$

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

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.

Now we explain how (14) is derived from (7)-(9). The term $\mathbf{\Pi}_i(\Delta c_i) \cdot \mathbf{F}_i \vec{U}_{i+1}(0)$ is for the delay transitions, where $\mathbf{F}_i$ specifies how the delay transitions are connected between $\mathcal{G}_i$ and $\mathcal{G}_{i+1}$. The term $\bar{\mathbf{\Pi}}^a_i(\Delta c_i) \cdot \vec{U}_0(0)$ is for Markovian transitions with reset. $\bar{\mathbf{\Pi}}^a_i(\Delta c_i)$ in the augmented CTMC $\mathcal{C}^a_i$ specifies the probabilities to first take transitions inside $\mathcal{G}_i$ and then a one-step Markovian transition back to $\mathcal{G}_0$. Eq. (15) is derived from (10). Since it is the last region and time goes to infinity, the time to enter the region is irrelevant (thus set to 0). Thus $\int_0^\infty \hat{\mathbf{M}}_i(\tau)d\tau$ boils down to $\hat{\mathbf{P}}_i$. In fact, the Markovian jump probability inside $\mathcal{G}_m$ can be taken from the embedded DTMC of $\mathcal{C}_m$, which is $\hat{\mathbf{P}}_i$.

Remark 3 *We note that for two-clock* DTA *which yield two-clock* DMTA*, the approach given in this section fails in general. In the single-clock case, the reset guarantees to jump to $\mathcal{G}_0(0)$ and delay to $\mathcal{G}_{i+1}(0)$ when it is in $\mathcal{G}_i$. However, in the two-clock case, after delay or reset generally only one clock has a fixed value while the value of the other one is not determined.*

The time-complexity of computing the reachability probability in the single-clock DTA case is $\mathcal{O}(m \cdot |S|^2 \cdot |Loc|^2 \cdot \lambda \cdot \Delta c + m^3 \cdot |S|^3 \cdot |Loc|^3)$, where m is the number of constants appearing in DTA, $|S|$ is the number of states in the CTMC, $|Loc|$ is the number of locations in the DTA, λ is the maximal exit rate in the CTMC and $\Delta c = \max_{0 \le i < m}\{c_{i+1} - c_i\}$. The first term $m \cdot |S|^2 \cdot |Loc|^2 \cdot \lambda \cdot \Delta c$ is due to the uniformization technique for computing transient distributions; and the second term $m^3 \cdot |S|^3 \cdot |Loc|^3$ is the time complexity for solving a system of linear equations with $\mathcal{O}(m \cdot |S| \cdot |Loc|)$ variables.

5 Conclusion

We addressed the quantitative verification of a CTMC $\mathcal{C}$ against a DTA $\mathcal{A}$. As a key result, we obtained that the probability of $\mathcal{C} \models \mathcal{A}$ can be reduced to computing reachability probabilities in PDPs. For single-clock DTA, this reduces to solving a system of linear equations yielding an equivalent, though simpler, characterization as in [13]. Moreover, it essentially provides a proof of the procedure proposed in [13]. Future work is planned for considering non-deterministic TA, and M(I)TL model checking.

Acknowledgement We thank Jeremy Sproston and the anonymous reviewers for insightful discussions and/or comments. This research is funded by the DFG research training group 1295 AlgoSyn, the Dutch Bsik project BRICKS, the NWO project QUPES and the EU FP7 project QUASIMODO.

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 continous-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*, 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] 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.

[9] 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.

[10] T. Chen, T. Han, J.-P. Katoen, and A. Mereacre. Quantitative model checking of continuous-time Markov chains against timed automata specification. Technical report, AIB-2009-02, RWTH Aachen University, Germany, 2009.

[11] C. Corduneanu. *Integral Equations and Applications*. Cambridge University Press, 1991.

[12] M. H. A. Davis. *Markov Models and Optimization*. Chapman and Hall, 1993.

[13] 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.

[14] 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.

[15] M. Y. Vardi. Automatic verification of probabilistic concurrent finite-state programs. In *FOCS*, pages 327–338, 1985.