LTL Model Checking of Time-Inhomogeneous Markov Chains*

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

[1] Design and Analysis of Communication Systems, University of Twente, The Netherlands
[2] Software Modelling and Verification, RWTH Aachen University, Germany
[3] Formal Methods and Tools, University of Twente, The Netherlands

Abstract. We investigate the problem of verifying linear-time properties against inhomogeneous continuous-time Markov chains (ICTMCs). A fundamental question we address is how to compute reachability probabilities. We consider two variants: time-bounded and unbounded reachability. It turns out that both can be characterized as the least solution of a system of integral equations. We show that for the time-bounded case, the obtained integral equations can be transformed into a system of ordinary differential equations; for the time-unbounded case, we identify two sufficient conditions, namely the *eventually periodic assumption* and the *eventually uniform assumption*, under which the problem can be reduced to solving a time-bounded reachability problem for the ICTMCs and a reachability problem for a DTMC. These results provide the basis for a model checking algorithm for LTL. Under the *eventually stable assumption*, we show how to compute the probability of a set of ICTMC paths which satisfy a given LTL formula. By an automata-based approach, we reduce this problem to the previous established results for reachability problems.

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. These Markov chains are typically *homogeneous*, i.e., the rates that determine the speed of changing state as well as the probabilistic nature of mode transitions are constant. However, in some situations constant rates do not adequately model real behaviors. This applies, e.g., to failure rates of hardware components [10] (that usually depend on the component's age), battery depletion [7] (where the power extraction rate non-linearly depends on the remaining amount of energy), and random phenomena that are subject to environmental influences. In these circumstances, Markov models

* Financially supported by the DFG research training group 1295 AlgoSyn, the Dutch Bsik project BRICKS, the NWO project QUPES, the EU project QUASIMODO, and the SRO DSN project of CTIT, University of Twente.

with *inhomogeneous* rates, i.e., rates that are time-varying functions, are much more appropriate [17].

Temporal logics and accompanying model-checking algorithms have been developed for discrete-time Markov chains (DTMCs for short), against linear-time properties [8, 9] and branching-time properties [11]; for CTMCs against branching-time properties [2, 3] and linear real-time properties [6]. And some of them have resulted in a number of successful model checkers such as PRISM [12] and MRMC [14]. However, the verification of *time-inhomogeneous* CTMCs (ICTMCs) has – to the best of our knowledge – not yet been investigated in depth, with the notable exception [15], which considered model checking a simple stochastic variant of Hennessy-Milner Logic (without fixed points) for *piecewise-constant* ICTMCs. The main aim of the current paper is to fill this gap by considering model checking ICTMCs w.r.t. *linear-time* properties.

One of the most fundamental linear-time properties are reachability problems. Here we address two variants: time-bounded and unbounded reachability. The former asks, given a set of goal states and a time bound, what is the probability of paths of a given ICTMC that reach the goal states within the time bound. Time-unbounded reachability is similar except that the time bound is infinity. To solve both of them, we first provide a characterization in terms of the least solution of a system of *integral equations*. This can be regarded as a generalization of similar results for CTMCs [2, 3] to ICTMCs. Furthermore, we show that for the time-bounded case, the obtained integral equations can be transformed into a system of (homogeneous) ordinary differential equations, which often enjoys an efficient numerical solution; for the time-unbounded case, generally this is not possible and one has to solve the system of integral equations directly, which is not so efficient and numerically unstable. To remedy this deficiency, we identify two sufficient conditions, i.e., the *eventually periodicity* and *eventually uniformity*, under which the problem can be reduced to the time-bounded reachability problem for ICTMCs and a (time-unbounded) reachability problem for DTMCs and thus can be solved efficiently. These classes subsume some interesting and important subclasses of ICTMCs, such as, the piecewise-constant case studied in [15] and ICTMCs with rates function representing *Weibull failure rates*. The latter distributions are important to model hazards and failures, and are popular in, e.g., reliability engineering. We then turn to model checking ICTMCs against LTL. Strictly speaking, we focus on computing the probability of the set of paths of a given ICTMC which satisfy the LTL formula. One of the main difficulties here compared to CTMCs is that in ICTMC, rates between states are functions over time instead of constants, and thus the topological structure of ICTMCs, when considered as a digraph, is not stable. To circumvent this problem, we identify a condition, i.e., the *eventually stable assumption* which intuitively means that after a (finite) time, the topological structure of the ICTMC does *not* change any more. Under this assumption, we can adapt the standard automata-based approach. A crucial ingredient is that we can construct a corresponding *separated* Büchi automaton from an LTL formula[1], based on which, one can build the product of the given ICTMC and the separated Büchi automaton while obtaining a well-defined stochastic process.

[1] Note that one can also use deterministic automata, but that would incur an extra (unnecessary) exponential blowup.

We then reduce the LTL model checking problem to the previous established results for reachability problems.

2 Preliminaries

Given a set S, let $Distr(S)$ denote the set of probability distributions over S.

Definition 1 (ICTMC). *A (labeled)* inhomogeneous continuous-time Markov chain (ICTMC) *is a tuple* $\mathcal{C} = (S, \mathrm{AP}, L, \alpha, \mathbf{R}(t))$*, where* S *is a* finite *set of* states*;* AP *is a finite set of* atomic propositions*;* $L : S \to 2^{\mathrm{AP}}$ *is a* labeling function*;* $\alpha \in Distr(S)$ *is an* initial distribution*;* $\mathbf{R}(t) : S \times S \times \mathbb{R}_{\geqslant 0} \to \mathbb{R}_{\geqslant 0}$ *is a rate matrix.*

Let diagonal matrix $\mathbf{E}(t) = \mathrm{diag}\,[E_s(t)] \in \mathbb{R}_{\geqslant 0}^{n \times n}$, where $n = |S|$ and $E_s(t) : S \times \mathbb{R}_{\geqslant 0} \to \mathbb{R}_{\geqslant 0}$ be defined as $E_s(t) = \sum_{s' \in S} \mathbf{R}_{s,s'}(t)$ for all $s \in S$, i.e., $E_s(t)$ is the *exit rate* of state s at time t. We require that all rates and exit rates, as functions of time t, are integrable. If all rates (and thus exit rates) are constant, we obtain a CTMC. A state s is *absorbing* if $\mathbf{R}_{s,s'}(t) = 0$, for $s' \neq s$.

Semantics. An ICTMC induces a stochastic process. The probability to take a transition from s to s' at time t within Δt time units is given by:

$$Prob\{s \to s', t, \Delta t\} = \int_0^{\Delta t} \mathbf{R}_{s,s'}(t+\tau) e^{-\int_0^\tau E_s(t+v)dv} d\tau = \int_t^{t+\Delta t} \mathbf{R}_{s,s'}(\tau) e^{-\int_t^\tau E_s(v)dv} d\tau.$$

Definition 2 (Timed paths). *Let* $\mathcal{C}$ *be an* ICTMC*. An infinite path starting at time* x *is a sequence* $\rho_x = s_0 \xrightarrow{t_0} s_1 \xrightarrow{t_1} s_2 \cdots$ *such that for each* $i \in \mathbb{N}$*,* $s_i \in S$*,* $t_i \in \mathbb{R}_{>0}$ *and* $\mathbf{R}_{s_i,s_{i+1}}(t) > 0$ *where* $t = x + \sum_{j=0}^{i} t_j$*. A finite path is a prefix of an infinite path ending in a state.*

We will sometimes omit the subscript of ρ_x if the starting time x is irrelevant. Let $Paths^{\mathcal{C}}$ and $Paths^{\mathcal{C}}(s, x)$ denote the set of (finite and infinite) paths in $\mathcal{C}$ and those starting from state s at time x, respectively. The superscript $\mathcal{C}$ is omitted whenever convenient. Let $\rho[n] := s_n$ be the n-th state of ρ (if it exists) and $\rho\langle n \rangle := t_n$ the time spent in state s_n. Let $\rho_x@t$ be the state occupied in ρ at time $t \in \mathbb{R}_{\geqslant 0}$, i.e. $\rho_x@t := \rho_x[n]$ where n is the smallest index such that $x + \sum_{i=0}^{n} \rho_x\langle i \rangle > t$. We assume w.l.o.g. that the time to stay in any state is strictly greater than 0.

Let $\mathcal{I}$ denote the set of all nonempty intervals $I \subseteq \mathbb{R}_{\geqslant 0}$ and let $I \oplus t$ (resp. $I \ominus t$) denote $\{x + t \mid x \in I\}$ (resp. $\{x - t \mid x \in I \wedge x \geqslant t\}$). The definition of a *Borel space* over paths through ICTMCs follows [3]. An ICTMC $\mathcal{C}$ with initial state s_0 and initial time x yields a probability measure $\Pr^{\mathcal{C}}_{s_0,x}$ on paths as follows: Let $C_x(s_0, I_0, \ldots, I_{k-1}, s_k)$ denote the *cylinder set* consisting of all paths $\rho \in Paths(s_0, x)$ such that $\rho[i] = s_i$ ($i \leqslant k$) and $\rho\langle i \rangle \in I_i$ ($i < k$). $\mathcal{F}(Paths(s_0, x))$ is the smallest σ-algebra on $Paths(s_0, x)$ which contains all cylinder sets $C_x(s_0, I_0, \ldots, I_{k-1}, s_k)$ for all state sequences $(s_0, \ldots, s_k) \in S^{k+1}$ and $I_0, \ldots, I_{k-1} \in \mathcal{I}$. The probability measure $\Pr^{\mathcal{C}}_{s_0,x}$ on $\mathcal{F}(Paths(s_0, x))$ is the unique measure recursively defined by:

$$\begin{aligned} &\Pr^{\mathcal{C}}_{s_0,x}\big(C_x(s_0, I_0, \ldots, I_{k-1}, s_k)\big) \\ &= \int_{I_0 \oplus x} \mathbf{R}_{s_0,s_1}(\tau_0) \cdot e^{-\int_x^{\tau_0} E_{s_0}(v)dv} \cdot \Pr^{\mathcal{C}}_{s_1,\tau_0}\big(C_{\tau_0}(s_1, I_1, \ldots, I_{k-1}, s_k)\big) d\tau_0 \end{aligned}$$

Example 1. An example ICTMC is illustrated in Fig. 2(a) (page 10), where $\mathrm{AP} = \{a, b, c\}$ and the rate functions are $r_i(t)$ $(1 \leqslant i \leqslant 6)$. In particular, the exit rate function of s_1 is $r_2(t) + r_3(t)$. The initial distribution is $\alpha(s_0) = 1$ and $\alpha(s) = 0$ for $s \neq s_0$. A possible rate function can be the ones depicted in Fig. 1 (page 7).

Linear temporal logic. The set of linear temporal logic (LTL) formulae over a set of atomic propositions AP is defined as follows:

Definition 3 (LTL syntax). *Given a set of atomic propositions* AP *which is ranged over by* a,b,*. . ., the syntax of* LTL *formulae is defined by:*

$$\varphi ::= \mathrm{tt} \mid a \mid \neg\varphi \mid \varphi \wedge \varphi \mid \mathsf{X}\,\varphi \mid \varphi \,\mathsf{U}\, \varphi.$$

The semantics of LTL for ICTMC $\mathcal{C}$ is defined in a standard way by a satisfaction relation, denoted $\models$, which is the least relation $\models \subseteq Paths^{\mathcal{C}} \times \mathbb{R}_{\geqslant 0} \times \mathbf{LTL}$ (here we use **LTL** to denote the set of LTL formulae) satisfying:

$(\rho, t) \models \mathrm{tt}$ $\qquad$ $(\rho, t) \models \varphi_1 \wedge \varphi_2$ iff $(\rho, t) \models \varphi_1$ and $(\rho, t) \models \varphi_2$

$(\rho, t) \models a$ iff $a \in L(\rho@t)$ $\qquad$ $(\rho, t) \models \neg\varphi$ iff $(\rho, t) \not\models \varphi$

$(\rho, t) \models \mathsf{X}\,\varphi$ iff $\exists \Delta t \geqslant 0.\ (\rho, t{+}\Delta t) \models \varphi$ and $\rho[1] = \rho@(t{+}\Delta t)$

$(\rho, t) \models \varphi_1 \,\mathsf{U}\, \varphi_2$ iff $\exists \Delta t \geqslant 0.\ (\rho, t{+}\Delta t) \models \varphi_2$ and $\forall t' < t{+}\Delta t.\ (\rho, t') \models \varphi_1$

We use $at_s{\in}\mathrm{AP}$ as an atomic proposition which holds solely at state s. For $F \subseteq S$, we write at_F for $\bigvee_{s \in F} at_s$. Let $Paths(s, x, \varphi) = \{\rho \in Paths(s, x) \mid (\rho, x) \models \varphi\}$. Note that a timed path $\rho = s_0 \xrightarrow{t_0} s_1 \xrightarrow{t_1} \cdots$ satisfies a formula φ iff the "discrete part" of ρ, namely, $s_0 s_1 s_2 \cdots (= \rho[0]\rho[1]\rho[2]\cdots)$ satisfies φ. It thus can be easily shown that the set $Paths(s, x, \varphi)$ is *measurable*. We denote the probability measure of $Paths(s, x, \varphi)$ as $Prob(s, x, \varphi) = \Pr_{s,x}(Paths(s, x, \varphi))$ and let $Prob^{\mathcal{C}}(\varphi) = \sum_{\alpha(s_0)>0} \alpha(s_0) \cdot Prob(s_0, 0, \varphi)$ be the probability that ICTMC $\mathcal{C}$ satisfies φ.

3 Reachability Analysis

In this section, we tackle reachability problems for ICTMCs. We distinguish two variants: *time-bounded* reachability and *time-unbounded* reachability. To solve both of them, we first give a characterization of $Prob(s, x, \Diamond^I at_F)$, namely, the probability of the set of paths which reach a set of goal states $F \subseteq S$ within time interval I starting from state s at time point x. This is done by resorting to a system of integral equations, which is a generalization of a similar characterization for CTMCs [3].

Proposition 1. *Let* $\mathcal{C} = (S, \mathrm{AP}, L, \alpha, \mathbf{R}(t))$ *be an* ICTMC *with* $s \in S$, $x \in \mathbb{R}_{\geqslant 0}$, $F \subseteq S$ *and interval* $I \subseteq \mathbb{R}_{\geqslant 0}$ *with* $T_1 = \inf I$ *and* $T_2 = \sup I$. *The function* $S \times \mathbb{R}_{\geqslant 0} \times \mathcal{I} \to [0, 1]$, $(s, x, I) \mapsto Prob(s, x, \Diamond^I at_F)$ *is the least fixed point of the operator*

$$\Omega : (S \times \mathbb{R}_{\geqslant 0} \times \mathcal{I} \to [0, 1]) \to (S \times \mathbb{R}_{\geqslant 0} \times \mathcal{I} \to [0, 1]) \quad ,$$

where $\Omega(f)(s, x, I) =$

$$\begin{cases} \displaystyle\int_0^{T_2} \sum_{s' \in S} \mathbf{R}_{s,s'}(x+\tau) e^{-\int_0^\tau E_s(x+v)dv} \cdot f(s', x+\tau, I \ominus \tau) d\tau, & \text{if } s \notin F \ (1) \\ \displaystyle e^{-\int_0^{T_1} E_s(x+v)dv} + \int_0^{T_1} \sum_{s' \in S} \mathbf{R}_{s,s'}(x{+}\tau) e^{-\int_0^\tau E_s(x+v)dv} \cdot f(s', x{+}\tau, I \ominus \tau) d\tau, & \text{if } s \in F \ (2) \end{cases}$$

3.1 Time-Bounded Reachability

We now solve the *time-bounded reachability* problem, i.e., given ICTMC $\mathcal{C}$, a set of goal states $F \subseteq S$ and a time bound $T \in \mathbb{R}_{\geqslant 0}$, to compute $Prob(s, x, \Diamond^{\leqslant T} at_F)$, the probability of $Paths(s, x, \Diamond^{\leqslant T} at_F)$ which is the set of paths that reach F within T time units given the initial time x. To accomplish this, we first compute $Prob(s, x, \Diamond^{=T} at_F)$, where the slightly different property $\Diamond^{=T} at_F$, in contrast to $\Diamond^{\leqslant T} at_F$, requires that states in F are reached at *exactly* time T. Note that $\Diamond^{\leqslant T} at_F$ and $\Diamond^{=T} at_F$ can also be written as $\Diamond^{[0,T]} at_F$ and $\Diamond^{[T,T]} at_F$, respectively, where $I = [0, T]$ or $I = [T, T]$ is a time interval. By instantiating (1), (2) in Prop. 1, we obtain that $Prob(s, x, \Diamond^{=T} at_F) =$

$$\begin{cases} \displaystyle\int_0^T \sum_{s' \in S} \mathbf{R}_{s,s'}(x+\tau) e^{-\int_0^\tau E_s(x+v)dv} \cdot Prob(s', x+\tau, \Diamond^{=T-\tau} at_F) d\tau, & \text{if } s \notin F \quad (3) \\ e^{-\int_0^T E_s(x+v)dv} + \displaystyle\int_0^T \sum_{s' \in S} \mathbf{R}_{s,s'}(x+\tau) e^{-\int_0^\tau E_s(x+v)dv} \cdot Prob(s', x+\tau, \Diamond^{=T-\tau} at_F) d\tau, & \text{if } s \in F \quad (4) \end{cases}$$

Intuitively, (3) and (4) are justified as follows: If $s \notin F$, the probability of reaching an F-state from s after exactly T time units given the starting time x equals the probability of reaching some direct successor s' of s in τ time units, multiplied by the probability of reaching an F-state from s' in the remaining $T - \tau$ time units. If $s \in F$ at time x, then it can either stay in s (i.e., delay) for T time units (the first summand in (4)), or regard s as a non-F state and take a transition (the second summand in (4)).

We now address the problem of solving (3) and (4), read as a *system of integral equations*. We define $\mathbf{\Pi}(x, T)$ as the matrix with entries $\mathbf{\Pi}_{i,j}(x, T)$ denoting the probability of the set of paths starting from state i at time x and reaching state j at time $x + T$. For any ICTMC, the following equation holds:

$$\mathbf{\Pi}(x, T) = \underbrace{\int_0^T \mathbf{M}(x, \tau)\mathbf{\Pi}(x + \tau, T - \tau) d\tau}_{\text{Markovian jump}} + \underbrace{\mathbf{D}(x, T)}_{\text{delay}} \quad (5)$$

$\mathbf{M}(x, T)$ is the probability density matrix where $\mathbf{M}_{i,j}(x, T) = \mathbf{R}_{i,j}(x+T) \cdot e^{-\int_0^T E_i(x+v)dv}$ is the density to move from state i to j at exactly time T and $\mathbf{D}(x, T)$ is the diagonal delay probability matrix with $\mathbf{D}_{i,i}(x, T) = e^{-\int_0^T E_i(x+v)dv}$.

We note that $\mathbf{\Pi}(x, T)$ is actually the (equivalent) matrix form of (3) and (4). For (4), it follows directly that each of its summands has a counterpart in (5). For (3), note that $\mathbf{D}(x, T)$ is a diagonal matrix where all the off-diagonal elements are 0 and that (3) does not allow a delay transition from a non-F state. This correspondence builds a half-bridge between $Prob(s, x, \Diamond^{=T} at_F)$ and $\mathbf{\Pi}(x, T)$, whereas the following proposition completes the other half bridge between $\mathbf{\Pi}(x, T)$ and the *transient probability vector* $\vec{\pi}(t)$ of ICTMCs:

Proposition 2. *Given* ICTMC $\mathcal{C}$ *with initial distribution* α *and rate matrix* $\mathbf{R}(t)$. *We have that* $\mathbf{\Pi}(0, t)$ *and* $\vec{\pi}(t)$ *satisfy the following two equations:*

$$\vec{\pi}(t) = \alpha \cdot \mathbf{\Pi}(0, t) \ , \quad (6)$$

$$\frac{d\vec{\pi}(t)}{dt} = \vec{\pi}(t) \cdot \mathbf{Q}(t), \ \vec{\pi}(0) = \alpha \ , \quad (7)$$

where $\mathbf{Q}(t) = \mathbf{R}(t) - \mathbf{E}(t)$ *is the* infinitesimal generator *of* $\mathcal{C}$.

Intuitively, this proposition implies that solving the system of integral equations $\mathbf{\Pi}(x,t)$ boils down to computing the transient probability vector $\vec{\pi}(t)$ with each element $\vec{\pi}_s(t)$ indicating the probability to be in state s at time t given the initial probability distribution $\alpha = \vec{\pi}(0)$. The transient probability is specified by a system of ODEs (7), the celebrated Chapman-Kolmogorov equations.

Given ICTMC $\mathcal{C}$, let $\mathcal{C}[F]$ be the ICTMC obtained by making the states in F absorbing in $\mathcal{C}$. We have the following theorem:

Theorem 1. *For any* ICTMC $\mathcal{C}$, $Prob^{\mathcal{C}}(s, x, \Diamond^{\leqslant T} at_F) = Prob^{\mathcal{C}[F]}(s, x, \Diamond^{=T} at_F)$.

To sum up, Proposition 1, 2 together with Theorem 1 suggest that computing time-bounded reachability probabilities in an ICTMC can be done, by first making the F states absorbing (and thus obtaining $\mathcal{C}[F]$) followed by solving a system of homogeneous ODEs (7) for $\mathcal{C}[F]$. By using standard numerical approaches, e.g., *Euler method* or *Runge-Kutta method* and their variants [16], this system of ODEs (i.e. the transient probability vector) can be solved.

3.2 Time-Unbounded Reachability

We then turn to the *time-unbounded reachability* problem, i.e., there are no constraints on the time to reach the F-states. Let $Prob(s, x, \Diamond\, at_F)$ denote the reachability probability from state s at time x to reach F within time interval $[0, \infty)$. Using Proposition 1, we can characterize $Prob(s, x, \Diamond\, at_F)$ as follows:

$$\begin{cases} \displaystyle\int_0^{\infty} \sum_{s' \in S} \mathbf{R}_{s,s'}(x+\tau) e^{-\int_0^{\tau} E_s(x+v)dv} \cdot Prob(s', x+\tau, \Diamond at_F) d\tau, & \text{if } s \notin F \quad (8) \\ 1, & \text{if } s \in F \quad (9) \end{cases}$$

The case $s \in F$ is derived from (2), where the probability to delay in an F-state for zero units of time is 1 and the probability to leave (i.e. taking a Markovian jump) an F-state in zero units of time is 0. When $s \notin F$, Eq. (8) is similar to (3) except that there is no bound on the time to leave a state $s \notin F$. Note that in contrast to the *time-bounded* case, in general it is not possible to reduce this system of integral equations to a system of ODEs. Since solving a system of integral equations is generally time consuming and numerically instable, we propose to investigate some special cases (subsets of ICTMCs), for which the reduction to ODEs is possible. Here we consider two such classes, i.e. *eventually periodic* ICTMCs and *eventually uniform* ICTMCs. Their common feature is that rate functions of the given ICTMC exhibit regular behaviors after some time T. This allows for computing time-unbounded reachability probabilities efficiently (e.g., via DTMCs). Hence the problem turns out to be reducible to computing the time-bounded reachability probabilities with time bound T, which has been tackled in the previous section, and reachability probabilities for DTMCs. Both of them, fortunately, enjoy efficient computational methods.

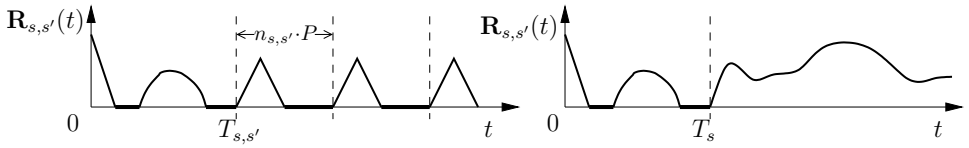

Fig. 1. Eventually periodic assumption (left) and eventually stable assumption (right)

Eventually periodic assumption. We consider *eventually periodic* ICTMCs.

Definition 4 (Eventually periodic assumption (EPA)). *An* ICTMC $\mathcal{C}$ *is* eventually periodic *if there exists some time* $P \in \mathbb{R}_{>0}$ *such that for any two states* $s, s' \in S$*, there exists some* $T_{s,s'} \in \mathbb{R}_{\geqslant 0}$ *and* $n_{s,s'} \in \mathbb{N}$ *such that*

$$\mathbf{R}_{s,s'}(t) = \begin{cases} \mathbf{R}^{(1)}_{s,s'}(t) & \text{if } t \leqslant T_{s,s'} \\ \mathbf{R}^{(2)}_{s,s'}(t) & \text{if } t > T_{s,s'} \end{cases}$$

where $\mathbf{R}^{(2)}_{s,s'}(t) = \mathbf{R}^{(2)}_{s,s'}(t + n_{s,s'}\cdot P)$.

An example rate function under the EPA is illustrated in Fig. 1 (left). After time point $T_{s,s'}$, the function $\mathbf{R}_{s,s'}(t)$ becomes periodic with the period $n_{s,s'}\cdot P$, where P is the "common factor" of all the periods of rate functions $\mathbf{R}_{s,s'}(t)$, for all $s, s' \in S$. For any ICTMC $\mathcal{C}$ satisfying EPA, let $T_{\mathrm{EP}} = \max_{s,s'\in S} T_{s,s'}$ and $P_{\mathrm{EP}} = (\gcd_{s,s'\in S} n_{s,s'})\cdot P$. Intuitively, T_{EP} is the time since when all rate functions are periodic and P_{EP} is the period of all the periodic rate functions. For instance, suppose there are two rate functions with $\mathbf{R}^{(2)}_{s_1,s_2}(t) = 2 + \cos(\frac{1}{2}t)$ and $\mathbf{R}^{(2)}_{s_2,s_3}(t) = 3 - \sin(\frac{1}{3}t)$, and let $T_{s_1,s_2} = 10$, $T_{s_2,s_3} = 15$. Then $T_{\mathrm{EP}}{=}\max\{10, 15\}{=}15$, $P{=}\pi$, $n_{s_1,s_2}{=}4$ and $n_{s_2,s_3}{=}6$, and $P_{\mathrm{EP}}{=}\gcd\{4, 6\}\cdot\pi{=}12\pi$.

Time-unbounded reachability probabilities for an ICTMC under the EPA can be computed according to Alg. 1 and justified by Theorem 2. Let us explain it in more detail. Due to (9), once F states are reached, it is irrelevant how the paths continue. This justifies the model transformation from $\mathcal{C}$ to $\mathcal{C}[F]$. The reachability problem can be divided into two subproblems: (I) first to compute the probability to reach state $s' \in S$ at exactly time T_{EC} (the second $Prob$ in (10), see below); and (II) then to compute the time-unbounded reachability from $s' \in S$ to F (the third $Prob$ in (10)). In the following we will focus on (II): Recall that we denote $Prob(s, T_{\mathrm{EP}}, \Diamond^{=P_{\mathrm{EP}}} at_{s'})$ to be the probability to reach from s to s' after time P_{EP} starting from time point T_{EP}. Since after time T_{EP} all rate functions are periodic with period P_{EP}, it holds that $Prob(s, T_{\mathrm{EP}}, \Diamond^{=P_{\mathrm{EP}}} at_{s'}) = Prob(s, T_{\mathrm{EP}}{+}n{\cdot}P_{\mathrm{EP}}, \Diamond^{=P_{\mathrm{EP}}} at_{s'})$, for all $n \in \mathbb{N}$. It then suffices to compute $Prob(s, T_{\mathrm{EP}}, \Diamond^{=P_{\mathrm{EP}}} at_{s'})$ for any $s, s' \in S$. Given the ICTMC $\mathcal{C}$ with state space S, we build a DTMC $\mathcal{D}_{\mathcal{C}} = (S, \mathbf{P})$ with $\mathbf{P}_{s,s'}{=}Prob^{\mathcal{C}}(s, T_{\mathrm{EP}}, \Diamond^{=P_{\mathrm{EP}}} at_{s'})$. Intuitively, $\mathbf{P}_{s,s'}$ is the one-step probability (one-step here means one period) to move from s to s', and the problem (II) is now reduced to computing the reachability probability from s to F-states in arbitrarily many steps (since the time-unbounded case is considered), i.e., $Prob^{\mathcal{D}_{\mathcal{C}[F]}}(s, \Diamond\, at_F)$. This can be done by standard methods, e.g., *value iteration* or *solving a system of linear equations*, see, among others, [4] (Ch. 10).

Theorem 2. *Let* $\mathcal{C} = (S, \mathrm{AP}, L, \alpha, \mathbf{R}(t))$ *be an* ICTMC *satisfying* EPA *with time* T_{EP} *and* P_{EP}*,* $s \in S$ *and* $F \subseteq S$*. Then:*

$$Prob^{\mathcal{C}}_{\mathrm{EP}}(s, 0, \Diamond\, at_F) = \sum_{s' \in S} Prob^{\mathcal{C}[F]}(s, 0, \Diamond^{=T_{\mathrm{EP}}} at_{s'}) \cdot Prob^{\mathcal{D}_{\mathcal{C}[F]}}(s', \Diamond\, at_F) \quad (10)$$

Remark 1. Sometimes we need to compute $Prob^{\mathcal{C}}_{\mathrm{EP}}(s, x, \Diamond\, at_F)$ for ICTMC $\mathcal{C} = (S, \mathrm{AP}, L, \alpha, \mathbf{R}(t))$, namely, the starting time is x instead of 0. To accomplish this, we define an ICTMC $\mathcal{C}' = (S, \mathrm{AP}, L, \alpha, \mathbf{R}'(t))$ such that $\mathbf{R}'(t) = \mathbf{R}(t + x)$ and it follows that $\mathcal{C}'$ still satisfies EPA (with $T'_{\mathrm{EP}} = T_{\mathrm{EP}} - x$ if $x \leqslant T_{\mathrm{EP}}$ and 0 otherwise; $P'_{\mathrm{EP}} = P_{\mathrm{EP}}$) and $Prob^{\mathcal{C}}_{\mathrm{EP}}(s, x, \Diamond\, at_F) = Prob^{\mathcal{C}'}_{\mathrm{EP}}(s, 0, \Diamond\, at_F)$.

Algorithm 1 Time-unbounded reachability for ICTMCs satisfying EPA

Require: ICTMC $\mathcal{C} = (S, \mathrm{AP}, L, \alpha, \mathbf{R}(t))$, EPA time T_{EP}, period P_{EP}
Ensure: $Prob^{\mathcal{C}}_{\mathrm{EP}}(s, 0, \Diamond\, at_F)$

1: For any two states $s, s' \in S$ in $\mathcal{C}[F]$, compute the *time-bounded reachability probability* with time bound T_{EP}, starting from time point x, i.e. $Prob^{\mathcal{C}[F]}(s, 0, \Diamond^{=T_{\mathrm{EP}}} at_{s'})$;
2: For any two states $s, s' \in S$ in $\mathcal{C}[F]$, compute the *time-bounded reachability probability* with time bound P_{EP}, starting from time point T_{EP}, i.e. $Prob^{\mathcal{C}[F]}(s, T_{\mathrm{EP}}, \Diamond^{=P_{\mathrm{EP}}} at_{s'})$;
3: Construct a *discrete-time Markov chain* (DTMC for short) $\mathcal{D}_{\mathcal{C}[F]} = (S, \mathbf{P})$ with $\mathbf{P}_{s,s'} = Prob^{\mathcal{C}[F]}(s, T_{\mathrm{EP}}, \Diamond^{=P_{\mathrm{EP}}} at_{s'})$. We denote the *reachability probability* from s to F in $\mathcal{D}_{\mathcal{C}}$ by $Prob^{\mathcal{D}_{\mathcal{C}[F]}}(s, \Diamond\, at_F)$;
4: Return $\sum_{s' \in S} Prob^{\mathcal{C}[F]}(s, 0, \Diamond^{=T_{\mathrm{EP}}} at_{s'}) \cdot Prob^{\mathcal{D}_{\mathcal{C}[F]}}(s', \Diamond\, at_F)$.

Eventually uniform assumption. The previous section has discussed rate functions enjoying a periodic behavior. A different class of rate functions are those which increase or decrease uniformly, e.g., an ICTMC in which all rates are a multiplicative of the *Weibull failure rate* which is characterized by the function $f(t) = \frac{\gamma}{\alpha}\left(\frac{t}{\alpha}\right)^{\gamma-1}$, where γ and α are the shape and scale parameters of the Weibull distribution, respectively. These distributions can e.g., characterize normal distributions, and are frequently used in reliability analysis. This suggests to investigate *eventually uniform* ICTMCs.

Definition 5 (Eventually uniform assumption (EUA)). *An* ICTMC $\mathcal{C}$ *is* eventually uniform *if there exists some time* $T_{\mathrm{EU}} \in \mathbb{R}_{>0}$ *and an integrable function* $f(t) : \mathbb{R}_{>0} \to \mathbb{R}_{>0}$ *such that* $\lim_{t\to\infty} \int_{T_{\mathrm{EU}}}^{t} f(\tau)d\tau \to \infty$ *and for any two states* $s, s' \in S$ *and* $t \geqslant T_{\mathrm{EU}}$*,* $\mathbf{R}_{s,s'}(t) = f(t) \cdot \mathbf{R}^{c}_{s,s'}$*, where* $\mathbf{R}^{c}_{s,s'}$ *is a constant.*

In terms of the infinitesimal generator $\mathbf{Q}(t)$ of the ICTMC $\mathcal{C}$, EUA intuitively entails that there exists some function $f(t)$ and *constant* infinitesimal generator $\mathbf{Q}^c = \mathbf{R}^c - \mathbf{E}^c$ ($\mathbf{R}^c$ and $\mathbf{E}^c$ are the constant rate matrix and exit rate matrix, respectively) such that $\mathbf{Q}(t) = f(t) \cdot \mathbf{Q}^c$ for all $t \geqslant T_{\mathrm{EU}}$. We also define the *constant transition probability matrix* $\mathbf{P}^c$ such that $\mathbf{P}^c_{s,s'} = \frac{\mathbf{R}^c_{s,s'}}{\mathbf{E}^c_s}$.

By restricting to the EUA, one can reduce the time-unbounded reachability problem for an ICTMC $\mathcal{C}$ to computing the time-bounded reachability probability with time bound T_{EU} and the reachability probability in a DTMC $\mathcal{D}^{\mathcal{C}}_{\mathrm{EU}}[F]$ with transition probability matrix $\mathbf{P}^c[F]$, where $\mathbf{P}^c[F]_{s,s'} = \mathbf{P}^c_{s,s'}$ for $s \notin F$; $\mathbf{P}^c[F]_{s,s} = 1$ and $\mathbf{P}^c[F]_{s,s'} = 0$, for $s \in F$ and $s' \neq s$. This is shown by the following theorem.

Theorem 3. *Let $\mathcal{C} = (S, \mathrm{AP}, L, \alpha, \mathbf{R}(t))$ be an* ICTMC *with $s \in S$. Given a set F of goal states and the* eventually uniform assumption *with the associated time T_{EU} and* DTMC $\mathcal{D}^{\mathcal{C}}_{\mathrm{EU}}$*, it holds that*

$$Prob^{\mathcal{C}}(s, 0, \Diamond\, at_F) = \sum_{s' \in S} Prob^{\mathcal{C}[F]}(s, 0, \Diamond^{=T_{\mathrm{EU}}} at_{s'}) \cdot Prob^{\mathcal{D}^{\mathcal{C}}_{\mathrm{EU}}[F]}(s', \Diamond\, at_F). \quad (11)$$

Remark 2. We note that the two assumptions, EUA and EPA are *incomparable*. There are rate functions (e.g. polynomials) which can *not* be represented as periodic functions but satisfy EUA; on the other hand, in case of EPA one can, for instance, assign the same sort of rate functions (e.g. $\sin$) with different periods, and thus obtain an ICTMC which invalidates EUA.

4 LTL Model Checking

In this section, we tackle the problem of model checking properties specified by LTL formulae for ICTMCs. Model checking CTMCs against LTL is not very difficult, since one can easily extract the *embedded* DTMC of the given CTMC, and thus reduce the problem to the corresponding model checking problem of DTMCs, which is well-studied, see, e.g. [8]. However, this approach does *not* work for ICTMCs, since the rates of the ICTMC vary with time. Below we shall employ an automata-based approach. For this purpose, some basic definitions are in order.

Definition 6 (Generalized Büchi automata). *A* generalized Büchi automaton (GBA) *is a tuple $\mathcal{A} = (\Sigma, Q, \Delta, Q_0, \mathcal{F})$, where Σ is a finite alphabet; Q is a finite set of states; $\Delta \subseteq Q \times \Sigma \times Q$ is a transition relation; $Q_0 \subseteq Q$ is a set of initial states, and $\mathcal{F} \subseteq 2^Q$ is a set of acceptance sets.*

We sometimes write $q \xrightarrow{\sigma} q'$ if $(q, \sigma, q') \subseteq \Delta$ for simplicity. An infinite *word* $w \in \Sigma^\omega$ is accepted by $\mathcal{A}$, if there exists an infinite *run* $\theta \in Q^\omega$ such that $\theta[0] \in Q_0$, $(\theta[i], w[i], \theta[i+1]) \subseteq \Delta$ for $i \geqslant 0$ and for each $F \in \mathcal{F}$, there exist infinitely many indices $j \in \mathbb{N}$ such that $\theta[j] \in F$. Note that $w[i]$ (resp. $\theta[i]$) denotes the i-th letter (resp. state) on w (resp. θ). The accepted language of $\mathcal{A}$, denoted $\mathcal{L}(\mathcal{A})$, is the set of all words accepted by $\mathcal{A}$. Given a GBA $\mathcal{A}$ and state q, we denote by $\mathcal{A}[q]$ the automaton $\mathcal{A}$ with q as the unique initial state. Note that $\mathcal{L}(\mathcal{A}) = \bigcup_{q \in Q_0} \mathcal{L}(\mathcal{A}[q])$. A GBA $\mathcal{A}$ is *separated*, if for any two states q, q', $\mathcal{L}(\mathcal{A}[q']) \cap \mathcal{L}(\mathcal{A}[q'']) = \varnothing$.

It follows from [9] that the correspondence between LTL formulae and separated GBA can be established:

Theorem 4. *For any* LTL *formula φ over* AP*, there exists a* separated GBA *$A_\varphi = (\Sigma, Q, \Delta, Q_0, \mathcal{F})$, where $\Sigma = 2^{\mathrm{AP}}$ and $|Q| \leqslant 2^{\mathcal{O}(|\varphi|)}$, such that $\mathcal{L}(A_\varphi)$ is the set of computations satisfying the formula φ.*

We note that the notion of *separated* is crucial for the remainder of this paper. A closely related notion, referred to as *unambiguous*, has been widely studied in automata and language theory, dating back to [1]. See also, among others, [5][13] for relevant literature. To the best of our knowledge, the notion of "separated" was firstly exploited in [9] for model checking DTMCs against LTL.

Definition 7 (Product). *Given an* ICTMC $\mathcal{C} = (S, \mathrm{AP}, L, \alpha, \mathbf{R}(t))$ *and a separated* GBA $\mathcal{A} = (\Sigma, Q, \Delta, Q_0, \mathcal{F})$*, the product* $\mathcal{C} \otimes \mathcal{A}$ *is defined as*

$$\mathcal{C} \otimes \mathcal{A} = (Loc, \mathrm{AP}, \tilde{L}, \tilde{\alpha}, \tilde{\mathbf{R}}(t)),$$

where $Loc = S \times Q$*;* $\tilde{L}(\langle s, q\rangle) = L(s)$*;* $\tilde{\alpha}(\langle s_0, q_0\rangle) = \alpha(s_0)$ *if* $\alpha(s_0) > 0$ *and* $q_0 \in Q_0$*, and undefined elsewhere; and* $\tilde{\mathbf{R}}_{\langle s,q\rangle,\langle s',q'\rangle}(t) = \mathbf{R}_{s,s'}(t)$ *if* $q \xrightarrow{L(s)} q'$.

For the sake of clarity, we call the states of a product as *locations*.

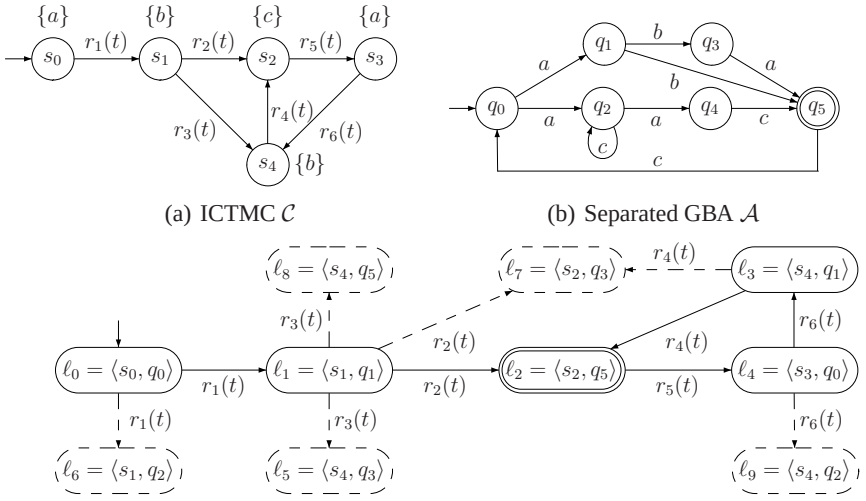

(a) ICTMC $\mathcal{C}$ (b) Separated GBA $\mathcal{A}$

(c) Product $\mathcal{C} \otimes \mathcal{A}$

Fig. 2. Example product construction of ICTMC $\mathcal{C}$ and separated GBA $\mathcal{A}$

Example 2. Given ICTMC $\mathcal{C}$ (Fig. 2(a)) and separated GBA $\mathcal{A}$ (Fig. 2(b)), the product $\mathcal{C} \otimes \mathcal{A}$ is shown in Fig. 2(c).

Remark 3. Note that in general the product itself is *not* an ICTMC. The reason is two-fold: (1) If $|Q_0| > 1$, then $\tilde{\alpha}$ is *not* a distribution; (2) The sum of the rates of outgoing transitions from a location might exceed the exit rate of the location. For instance, in Example 2, the exit rate of ℓ_0, as defined, is $\tilde{E}_{\ell_0}(t) = E_{s_0}(t) = r_1(t)$; while the sum of the rates of its outgoing transitions is $2r_1(t)$. However, due to the fact that $\mathcal{A}$ is separated, as we will see later, it would not be a problem, cf. Proposition 3.

The generalized Büchi acceptance condition, roughly speaking, requires to visit some states infinitely often. As in the tradition of model checking Markovian models, we need to identify *bottom strongly connected components* (BSCCs) of the product (when read as a graph). A *strongly connected component* (SCC for short) of the product denotes a strongly connected set of locations such that no proper superset is strongly connected. A BSCC is an SCC from which no location outside is reachable. Unfortunately, generally in ICTMCs, there is no way to define a BSCC over the product since

the rate of each transition is a *function* of time instead of a constant and thus a BSCC at time t might not be a BSCC at time t'. In other words, the topological structure (edge relation) of the product might change at any moment of time, which is one of the main difficulties of model checking ICTMCs.

To circumvent this problem, we make an (arguably mild) assumption, that is, we assume that ICTMCs are *eventually stable*, in the following sense.

Definition 8 (Eventually stable assumption (ESA)). *An* ICTMC $\mathcal{C}$ *is* eventually stable *if for each state $s \in S$, there exists some time T_s such that for any $t \geqslant T_s$ and $s' \in S$, either $\mathbf{R}_{s,s'}(t) > 0$ or $\mathbf{R}_{s,s'}(t) = 0$.*

W.l.o.g., we assume T_s is the smallest time point that the above assumption holds for state s. Let $T_{\mathrm{ES}} = \max_{s \in S} T_s$ be the smallest time point that an ICTMC is *stable*. Intuitively, an ICTMC is stable if its topological structure does *not* change any more. More specifically, transitions can alter their rates, but not from positive to zero or vice versa, i.e., no transitions will "disappear" or "newly created". An example rate function is illustrated in Fig. 1 (right), where after T_{ES} the rates keep strictly positive (note the particular value is irrelevant here). It turns out that ESA is essential for identifying stable BSCCs (also model checking LTL). A *stable* product as well as *stable* BSCC are defined in the same way, relative to the time point T_{ES}. *In the sequel, when we refer to* BSCC*s, we implicitly refer to the stable* BSCC*s in the stable product.* In accordance with this, we will sometimes write $s \mapsto s'$ for ICTMC $\mathcal{C}$ if $\mathbf{R}_{s,s'}(t) > 0$ with $t \geqslant T_{\mathrm{ES}}$; and similarly for $\ell \mapsto \ell'$ in the product.

Definition 9 (aBSCC). *Given the product $\mathcal{C} \otimes \mathcal{A}$ of an* ICTMC $\mathcal{C} = (S, \mathrm{AP}, L, \alpha, \mathbf{R}(t))$ *satisfying* ESA *and a* GBA $\mathcal{A} = (\Sigma, Q, \Delta, Q_0, \mathcal{F})$*, we define*

I. *a* SCC *is a set of locations $B \subseteq S \times Q$ such that (i) B is strongly connected meaning that for any two locations $\ell, \ell' \in B$, $\ell \mapsto^* \ell'$ where $\mapsto^*$ denotes the reflexive and transitive closure of $\mapsto$, and (ii) no proper superset of B is strongly connected;*

II. *a* SCC B *is* accepting *if $\forall F \in \mathcal{F}$, there exists some $\langle s, q \rangle \in B$ such that $q \in F$;*

III. *a* SCC B *is an* accepting bottom SCC *($B \in a$BSCC for short) if (i) B is accepting; (ii) for each location $\ell \in B$, there does* not *exist any location ℓ' such that $\ell \mapsto \ell'$ and ℓ' is in any other* accepting SCC*; (iii) for each location $\ell = \langle s, q \rangle \in B$, for any s' with $s \mapsto s'$, $\langle s', q' \rangle \in B$ for some q'.*

As an example, we note, suppose that $r_4(t), r_5(t), r_6(t) > 0$ when $t \geqslant T_{\mathrm{ES}}$, that an accepting BSCC in the stable product in Fig. 2(c) is formed by ℓ_2, ℓ_3, ℓ_4. Note that $\{\ell_7\}$ is *not* an (accepting) SCC, so III(*ii*) is not violated. $\{\ell_8\}$ is not an SCC either, since we would require that $\ell_8 \mapsto^* \ell_8$ which fails to be.

Recall that given an LTL formula, one can obtain a corresponding *separated* automaton, which renders us very nice properties for the product defined in Definition 7. A couple of lemmas, dedicated to illustrate these properties are in order. The following two essentially exploit the fact that for each accepted word of a separated GBA, there is a unique accepting path.

Lemma 1. *Given the product $\mathcal{C} \otimes \mathcal{A}$ where $\mathcal{A}$ is separated. For any a*BSCC *B of the stable product $\mathcal{C} \otimes \mathcal{A}$, it* cannot *be the case that $\langle s, q\rangle \mapsto \langle s', q'\rangle$ and $\langle s, q\rangle \mapsto \langle s', q''\rangle$ for any $\langle s, q\rangle, \langle s', q'\rangle, \langle s', q''\rangle$ in B with $q' \neq q''$.*

We say that two locations $\langle s, q\rangle$ and $\langle s', q'\rangle$ in the product $\mathcal{C} \otimes \mathcal{A}$ are *connected*, if $q \xrightarrow{L(s)} q'$[2]. We say that from location $\langle s, q\rangle$ there is a path leading to a BSCC B, if there is a sequence $\langle s_0, q_0\rangle, \langle s_1, q_1\rangle, \ldots, \langle s_n, q_n\rangle$ such that $\langle s, q\rangle = \langle s_0, q_0\rangle$, $\langle s_i, q_i\rangle$ and $\langle s_{i+1}, q_{i+1}\rangle$ are connected for $0 \leqslant i < n$ and $\langle s_n, q_n\rangle \in B$.

Lemma 2. *Given the product $\mathcal{C} \otimes \mathcal{A}$ of an* ICTMC *$\mathcal{C}$ and a separated* GBA *$\mathcal{A}$, it* cannot *be the case that there are two locations $\langle s, q\rangle$ and $\langle s, q'\rangle$ with $q \neq q'$ such that both of them have a path reaching an a*BSCC*.*

As said, given ICTMC $\mathcal{C}$ and separated GBA $\mathcal{A}$, the product $\mathcal{C} \otimes \mathcal{A}$ itself is not an ICTMC (see Example 2). However, thanks to the fact that $\mathcal{A}$ is separated, we can transform $\mathcal{C} \otimes \mathcal{A}$ into an ICTMC. Lemma 1 and 2 entail that in the product $\mathcal{C} \otimes \mathcal{A}$, we can safely remove the locations which do not lead to an accepting BSCC, and thus obtain an ICTMC model, denoted $\mathcal{C} \underline{\otimes} \mathcal{A}$. Let us illustrate this by continuing Example 2. First note that the dashed locations are the trap locations from which the accepted location ℓ_2 cannot be reached. Those locations can safely be removed since the paths passing them will never be accepted. It is not a coincidence that at most one of the outgoing transitions from those "nondeterministic" locations (i.e., $\ell_0, \ell_1, \ell_3, \ell_4$) can reach the accepted locations. This is guaranteed by the separated property of the automaton (Lemma 2). By deleting all the dashed locations, we obtain $\mathcal{C} \underline{\otimes} \mathcal{A}$.

The following proposition claims that $\mathcal{C} \underline{\otimes} \mathcal{A}$ can be viewed as an ICTMC in the sense that it defines a stochastic process exactly as an ICTMC. The crucial point is that in $\mathcal{C} \underline{\otimes} \mathcal{A}$, for each location ℓ and time t, the sum of the rates of the emanating transitions from ℓ does *not* exceed the exit rate of ℓ. (Note that the sum could be strictly less than the exit rate as for ℓ_1 in Fig. 2(c), thus it is "substochastic".) With a little abusing of terms, we call this model an ICTMC.

Proposition 3. *$\mathcal{C} \underline{\otimes} \mathcal{A}$ is an* ICTMC*. Moreover, for each accepting cylinder set, $\mathcal{C}$ and $\mathcal{C} \underline{\otimes} \mathcal{A}$ give rise to the same probability.*

Let $\mathcal{C} \underline{\otimes} \mathcal{A}^\star$ be obtained from $\mathcal{C} \underline{\otimes} \mathcal{A}$ by making each location in the aBSCCs absorbing, and define $F^\star$ as the set of locations in any aBSCC. Given an ICTMC $\mathcal{C}$ with eventually stable assumption (with T_{ES}) and an LTL formula φ, the probability of the set of paths of $\mathcal{C}$ satisfying φ, denoted $Prob^{\mathcal{C}}_{\mathrm{ES}}(\varphi)$, can be computed by Alg. 2.

Theorem 5. *For an* ICTMC *$\mathcal{C}$ with T_{ES} and an* LTL *formula φ,*

$$Prob^{\mathcal{C}}_{\mathrm{ES}}(\varphi) = \sum_{\tilde{\alpha}(\ell_0)>0} \sum_{\ell \in Loc} \tilde{\alpha}(\ell_0) \cdot Prob^{\mathcal{C} \underline{\otimes} \mathcal{A}}(\ell_0, 0, \Diamond^{=T_{\mathrm{ES}}} at_\ell) \cdot Prob^{\mathcal{C} \underline{\otimes} \mathcal{A}^\star}(\ell, T_{\mathrm{ES}}, \Diamond\, at_{F^\star}).$$

[2] Note that we do *not* require that $s \mapsto s'$. So "connected" is purely a graph-theoretic notion where the time is irrelevant.

Algorithm 2 Model checking ICTMC against LTL

Require: ICTMC $\mathcal{C}$, LTL formula φ, ESA time T_{ES};
Ensure: $Prob^{\mathcal{C}}_{\mathrm{ES}}(\varphi)$

1: Transform φ to a separated generalized Büchi automaton $\mathcal{A}$;
2: Build the product $\mathcal{C} \otimes \mathcal{A} = (Loc, \mathrm{AP}, \tilde{L}, \tilde{\alpha}, \tilde{\mathbf{R}}(t))$;
3: Find all *accepting* BSCCs in the (stable) product $\mathcal{C} \otimes \mathcal{A}$;
4: Remove all the trap locations yielding $\mathcal{C} \underline{\otimes} \mathcal{A}$;
5: Compute the *time-bounded reachability* in $\mathcal{C} \underline{\otimes} \mathcal{A}$ from initial location ℓ_0 to each $\ell \in Loc$, $Prob^{\mathcal{C} \underline{\otimes} \mathcal{A}}(\ell_0, 0, \Diamond^{=T_{\mathrm{ES}}} at_\ell)$;
6: Make each location in the aBSCCs absorbing, thus obtaining $\mathcal{C} \underline{\otimes} \mathcal{A}^\star$ and $F^\star$;
7: Compute the *time-unbounded reachability probability* in $\mathcal{C} \underline{\otimes} \mathcal{A}^\star$ from each $\ell \in Loc$ to $F^\star$, i.e., $Prob^{\mathcal{C} \underline{\otimes} \mathcal{A}^\star}(\ell, T_{\mathrm{ES}}, \Diamond\, at_{F^\star})$;
8: $Prob^{\mathcal{C}}_{\mathrm{ES}}(\varphi) = \sum_{\tilde{\alpha}(\ell_0)>0} \sum_{\ell \in Loc} \tilde{\alpha}(\ell_0) \cdot Prob^{\mathcal{C} \underline{\otimes} \mathcal{A}}(\ell_0, 0, \Diamond^{=T_{\mathrm{ES}}} at_\ell) \cdot Prob^{\mathcal{C} \underline{\otimes} \mathcal{A}^\star}(\ell, T_{\mathrm{ES}}, \Diamond\, at_{F^\star})$.

Note that $Prob^{\mathcal{C} \underline{\otimes} \mathcal{A}}(\ell_0, 0, \Diamond^{=T_{\mathrm{ES}}} at_\ell)$ and $Prob^{\mathcal{C} \underline{\otimes} \mathcal{A}^\star}(\ell, 0, \Diamond\, at_{\ell_{F^\star}})$ can be computed by the approaches in Section 3.1 and 3.2, respectively. Computing the former relies on solving a system of ODEs, whereas computing the latter, as stated in Section 3.2, one has to solve a system of integral equations in general.

*Remark 4 (*EPA, EUA *and* ESA*).* EPA and ESA are *incomparable*, i.e., there are ICTMCs that are eventually periodic but not stable (see e.g., the ICTMC with one rate function in Fig. 1 (left)), and vice versa (see, e.g., that in Fig. 1 (right)). When both assumptions are applied, we obtain ICTMCs that are "eventually positive periodic", i.e., eventually periodic and all rate function values in the periods are either strictly positive or being zero. For this subset of ICTMCs, one can resort to solving a system of ODEs and linear equations, as presented in Theorem 2 as well as Alg. 1.

EUA and ESA are *incomparable* as well. The counterexamples for both directions can be easily constructed. When both assumptions are applied, as in the previous case, the subset of ICTMCs (where $f(t)$ is eventually strictly positive) can be dealt with by solving a systems of ODEs and linear equations (Theorem 3).

The comparison of EPA and EUA can be found in Remark 2. We emphasize once again that ESA is of most importance in LTL model checking, in order to find stable BSCCs. However EPA or EUA are certain subsets of ICTMCs that we can efficiently deal with (meaning by solving a system of ODEs and linear equations). We mention that there are other approaches which can handle and solve the system of integral equations, e.g., approximation by truncating the infinite range of the integral.

Example 3. We continue Example 2 to show how to compute the set of paths of ICTMC $\mathcal{C}$ accepted by $\mathcal{A}$. Let the rate functions be defined as: $r_i(t) = i$ for $t \geqslant 0$ and $3 \leqslant i \leqslant 6$; and

$$r_1(t) = \begin{cases} t & x \in [0, 9.5) \\ 0 & x \in [9.5, 10) \\ 2 + \cos(\frac{1}{2}t) & x \in [10, \infty) \end{cases} \qquad r_2(t) = \begin{cases} 4.1 & x \in [0, 15) \\ 7.6 - \sin(\frac{1}{3}t) & x \in [15, \infty) \end{cases}$$

It is not difficult to see that this ICTMC satisfies both the ESA and EPA and $T_{\mathrm{ES}} = 10$, $T_{\mathrm{EP}} = 15$ and $P_{\mathrm{EP}} = 12\pi$.

To compute $Prob^{\mathcal{C}}_{\mathrm{ES}}(\mathcal{A})$, Alg. 2 is applied. Note that we omit the step of transforming an LTL formula to a GBA and the notation $Prob^{\mathcal{C}}_{\mathrm{ES}}(\mathcal{A})$ is self-explanatory. We consider the first $Prob$ appearing in step 8, namely, to compute $Prob^{\mathcal{C}\underline{\otimes}\mathcal{A}}(\ell_0, 0, \Diamond^{=10}\, at_\ell)$ for all ℓ in $\mathcal{C}\underline{\otimes}\mathcal{A}$. This is actually to compute the transient probability vector in $\mathcal{C}\underline{\otimes}\mathcal{A}$ at time $T_{\mathrm{ES}} = 10$, which can be done by solving a system of ODEs. We then consider the second $Prob$ appearing in step 8, namely, to compute $Prob^{\mathcal{C}\underline{\otimes}\mathcal{A}^\star}(\ell, 10, \Diamond\, at_{\ell_2})$ (note that $F^* = \{\ell_2\}$). Finally, we wrap them up as follows:

$$Prob^{\mathcal{C}}_{\mathrm{ES}}(\mathcal{A}) \stackrel{\sum}{=} \begin{cases} Prob^{\mathcal{C}\underline{\otimes}\mathcal{A}}(\ell_0, 0, \Diamond^{=10}\, at_{\ell_0}) \cdot Prob^{\mathcal{C}\underline{\otimes}\mathcal{A}^\star}(\ell_0, 10, \Diamond\, at_{\ell_2}) \\ Prob^{\mathcal{C}\underline{\otimes}\mathcal{A}}(\ell_0, 0, \Diamond^{=10}\, at_{\ell_1}) \cdot Prob^{\mathcal{C}\underline{\otimes}\mathcal{A}^\star}(\ell_1, 10, \Diamond\, at_{\ell_2}) \\ Prob^{\mathcal{C}\underline{\otimes}\mathcal{A}}(\ell_0, 0, \Diamond^{=10}\, at_{\ell_2}) \cdot 1 \\ Prob^{\mathcal{C}\underline{\otimes}\mathcal{A}}(\ell_0, 0, \Diamond^{=10}\, at_{\ell_3}) \cdot 1 \\ Prob^{\mathcal{C}\underline{\otimes}\mathcal{A}}(\ell_0, 0, \Diamond^{=10}\, at_{\ell_4}) \cdot 1 \end{cases} \tag{12}$$

The left column of (12) is the transient probability vector; we then show how to compute the elements in the right column. For this purpose, generally we have to solve a system of integral equations. Here we obtain the following one:

$$\begin{cases} f_{\ell_0}(x) = \int_0^\infty r_1(x+\tau) e^{-\int_0^\tau r_1(x+v)dv} \cdot f_{\ell_1}(x+\tau) d\tau \\ f_{\ell_1}(x) = \int_0^\infty r_2(x+\tau) e^{-\int_0^\tau r_2(x+v)+r_3(x+v)dv} \cdot f_{\ell_2}(x+\tau) d\tau \\ f_{\ell_2}(x) = 1 \end{cases}$$

In this case, one obtains that $f_{\ell_2}(x) = 1$,

$$f_{\ell_1}(10) = \int_{10}^{15} 4.1 e^{-\int_{10}^{\tau}(4.1+3)dv} d\tau + \int_{15}^{\infty} (7.6 - \sin(\frac{1}{3}\tau)) \cdot e^{-\int_{10}^{15}(4.1+3)dv - \int_{15}^{\tau}(7.6-\sin(\frac{1}{3}(x+v))+3)dv} d\tau \ .$$

and $f_{\ell_0}(10)$ can be computed accordingly. It follows that $Prob^{\mathcal{C}\underline{\otimes}\mathcal{A}^\star}(\ell_0, 10, \Diamond\, at_{\ell_2}) = f_{\ell_0}(10)$ and $Prob^{\mathcal{C}\underline{\otimes}\mathcal{A}^\star}(\ell_1, 10, \Diamond\, at_{\ell_2}) = f_{\ell_1}(10)$. Hence (12) can be obtained.

Alternatively, let us note that fortunately in this case, the EPA is satisfied. So one can apply Alg. 1 to compute $Prob^{\mathcal{C}\underline{\otimes}\mathcal{A}^\star}(\ell, 10, \Diamond\, at_{\ell_2})$. Let us illustrate for the case that $\ell = \ell_0$. (The case that $\ell = \ell_1$ is similar.) For the first $Prob$ in Alg. 1, step 4, it is again to compute the transient probability matrix $Prob^{\mathcal{C}\underline{\otimes}\mathcal{A}^\star}(\ell, 10, \Diamond^{=15-10}\, at_{\ell'})$, for $\ell, \ell' \in \{\ell_i \mid 1 \leqslant i \leqslant 4\}$ (note that ℓ_2 is already made absorbing in $\mathcal{C}\underline{\otimes}\mathcal{A}^\star$); and for the second $Prob$, we need to construct the DTMC $\mathcal{D}$ with $\mathbf{P}_{\ell,\ell'} = Prob^{\mathcal{C}\underline{\otimes}\mathcal{A}^\star}(\ell, 15, \Diamond^{=12\pi}\, at_{\ell'})$. It follows that

$$Prob^{\mathcal{C}\underline{\otimes}\mathcal{A}^\star}_{\mathrm{EP}}(\ell_0, 10, \Diamond\, at_{\ell_2}) \stackrel{\sum}{=} \begin{cases} Prob^{\mathcal{C}\underline{\otimes}\mathcal{A}^\star}(\ell_0, 10, \Diamond^{=15-10}\, at_{\ell_0}) \cdot Prob^{\mathcal{D}}(\ell_0, \Diamond\, at_{\ell_2}) \\ Prob^{\mathcal{C}\underline{\otimes}\mathcal{A}^\star}(\ell_0, 10, \Diamond^{=15-10}\, at_{\ell_1}) \cdot Prob^{\mathcal{D}}(\ell_1, \Diamond\, at_{\ell_2}) \\ Prob^{\mathcal{C}\underline{\otimes}\mathcal{A}^\star}(\ell_0, 10, \Diamond^{=15-10}\, at_{\ell_2}) \cdot 1 \end{cases}$$

The $\mathbf{P}$ matrix of the DTMC $\mathcal{D}$ is as follows (let Θ denote $\mathcal{C}\underline{\otimes}\mathcal{A}^\star$):

$$\mathbf{P} = \begin{pmatrix} Prob^{\Theta}(\ell_0, 15, \Diamond^{=12\pi}\, at_{\ell_0}) & Prob^{\Theta}(\ell_0, 15, \Diamond^{=12\pi}\, at_{\ell_1}) & Prob^{\Theta}(\ell_0, 15, \Diamond^{=12\pi}\, at_{\ell_2}) \\ 0 & Prob^{\Theta}(\ell_1, 15, \Diamond^{=12\pi}\, at_{\ell_1}) & Prob^{\Theta}(\ell_1, 15, \Diamond^{=12\pi}\, at_{\ell_2}) \\ 0 & 0 & 1 \end{pmatrix}.$$

Hence $Prob^{\mathcal{C}\underline{\otimes}\mathcal{A}^\star}_{\mathrm{EP}}(\ell_0, 10, \Diamond\, at_{\ell_2})$ can be easily computed.

5 Conclusion

We have studied the problem of verifying linear-time properties against ICTMCs. Two variants of reachability problems, i.e. time-bounded and unbounded reachability, as well as LTL properties were considered. Future work consists of identifying more classes of ICTMCs for which efficient computational methods exist such that the approach studied in this paper can be applied. Other specifications like (D)TA, M(I)TL, will also be investigated.

References

1. A. Arnold. Rational omega-languages are non-ambiguous. *Theor. Comput. Sci.*, 26:221–223, 1983.
2. 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.
3. 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.
4. C. Baier and J.-P. Katoen. *Principles of Model Checking*. MIT Press, 2008.
5. O. Carton and M. Michel. Unambiguous Büchi automata. *Theor. Comput. Sci.*, 297(1-3):37–81, 2003.
6. T. Chen, T. Han, J.-P. Katoen, and A. Mereacre. Quantitative model checking of continuous-time Markov chains against timed automata specification. In *LICS*, to appear, 2009.
7. L. Cloth, M. R. Jongerden, and B. R. Haverkort. Computing battery lifetime distributions. In *DSN*, pages 780–789, 2007.
8. C. Courcoubetis and M. Yannakakis. The complexity of probabilistic verification. *J. ACM*, 42(4):857–907, 1995.
9. J.-M. Couvreur, N. Saheb, and G. Sutre. An optimal automata approach to LTL model checking of probabilistic systems. In *LPAR*, LNCS 2850, pages 361–375, 2003.
10. S. S. Gokhale, M. R. Lyu, and K. S. Trivedi. Analysis of software fault removal policies using a non-homogeneous continuous time Markov chain. *Software Quality Control*, 12(3):211–230, 2004.
11. H. Hansson and B. Jonsson. A logic for reasoning about time and reliability. *Formal Asp. Comput.*, 6(5):512–535, 1994.
12. A. Hinton, M. Z. Kwiatkowska, G. Norman, and D. Parker. Prism: A tool for automatic verification of probabilistic systems. In *TACAS*, pages 441–444, 2006.
13. D. Kähler and T. Wilke. Complementation, disambiguation, and determinization of Büchi automata unified. In *ICALP (1)*, LNCS 5125, pages 724–735, 2008.
14. J.-P. Katoen, M. Khattri, and I. S. Zapreev. A Markov reward model checker. In *QEST*, pages 243–244, 2005.
15. J.-P. Katoen and A. Mereacre. Model checking HML on piecewise-constant inhomogeneous Markov chains. In *FORMATS*, LNCS 5215, pages 203–217, 2008.
16. J. D. Lambert. *Numerical Methods for Ordinary Differential Systems*. John Wiley & Sons, 1991.
17. A. Rindos, S. P. Woolet, Y. Viniotis, and K. S. Trivedi. Exact methods for the transient analysis of non-homogeneous continuous-time Markov chains. In *Numerical Solution of Markov Chains (NSMC)*, pages 121–134, 1995.