Probabilistic Timed Automata with Clock-Dependent Probabilities

Jeremy Sproston(✉)

Dipartimento di Informatica, University of Turin, Turin, Italy
sproston@di.unito.it

Abstract. Probabilistic timed automata are classical timed automata extended with discrete probability distributions over edges. We introduce clock-dependent probabilistic timed automata, a variant of probabilistic timed automata in which transition probabilities can depend linearly on clock values. Clock-dependent probabilistic timed automata allow the modelling of a continuous relationship between time passage and the likelihood of system events. We show that the problem of deciding whether the maximum probability of reaching a certain location is above a threshold is undecidable for clock-dependent probabilistic timed automata. On the other hand, we show that the maximum and minimum probability of reaching a certain location in clock-dependent probabilistic timed automata can be approximated using a region-graph-based approach.

1 Introduction

Reactive systems are increasingly required to satisfy a combination of qualitative criteria (such as safety and liveness) and quantitative criteria (such as timeliness, reliability and performance). This trend has led to the development of techniques and tools for the formal verification of both qualitative and quantitative properties. In this paper, we consider a formalism for real-time systems that exhibit randomised behaviour, namely probabilistic timed automata (PTA) [10,17]. PTAs extend classical Alur-Dill timed automata [4] with discrete probabilistic branching over automata edges; alternatively a PTA can be viewed as a Markov decision process [20] or a Segala probabilistic automaton [21] extended with timed-automata-like clock variables and constraints over those clocks. PTAs have been used previously to model case studies including randomised protocols and scheduling problems with uncertainty [16,19], some of which have become standard benchmarks in the field of probabilistic model checking.

We recall briefly the behaviour of a PTA: as time passes, the model stays within a particular discrete state, and the values of its clocks increase at the same rate; at a certain point in time, the model can leave the discrete state if the current values of the clocks satisfy a constraint (called a guard) labelling one of the probability distributions over edges leaving the state; then a probabilistic choice as to which discrete state to then visit is made according to the chosen edge distribution. In the standard presentation of PTAs, any dependencies between time and probabilities over edges must be defined by utilising multiple

M. Hague and I. Potapov (Eds.): RP 2017, LNCS 10506, pp. 144–159, 2017.
DOI: 10.1007/978-3-319-67089-8_11

distributions enabled with different sets of clock values. For example, to model the fact that a packet loss is more likely as time passes, we can use clock x to measure time, and two distributions μ_1 and μ_2 assigning probability λ_1 and λ_2 (for $\lambda_1 < \lambda_2$), respectively, to taking edges leading to a discrete state corresponding to packet loss, where the guard of μ_1 is $x \leq c$ and the guard of μ_2 is $x > c$, for some constant $c \in \mathbb{N}$. Hence, when the value of clock x is not more than c, a packet loss occurs with probability λ_1, otherwise it occurs with probability λ_2. A more direct way of expressing the relationship between time and probability would be letting the probability of making a transition to a discrete state representing packet loss be dependent on the value of the clock, i.e., let the value of this probability be equal to $f(x)$, where f is an increasing function from the values of x to probabilities. We note that such a kind of dependence of discrete branching probabilities on values of continuous variables is standard in the field of stochastic hybrid systems, for example in [1].

In this paper we consider such a formalism based on PTAs, in which all probabilities used by edge distributions can be expressed as functions of values of the clocks used by the model: the resulting formalism is called *clock-dependent probabilistic timed automata* (cdPTA). We focus on a simple class of functions from clock values to probabilities, namely those that can be expressed as sums of continuous piecewise linear functions, and consider a basic problem in the context of probabilistic model checking, namely probabilistic reachability: determine whether the maximum (respectively, minimum) probability of reaching a certain set of locations from the initial state is above (respectively, below) a threshold. After introducing cdPTAs (in Sect. 2), our first result (in Sect. 3) is that the probabilistic reachability problem is undecidable for cdPTA with a least three clocks. This result is inspired from recent related work on stochastic timed Markov decision processes [2]. Furthermore, we give an example of cdPTA with one clock for which the maximal probability of reaching a certain location involves a particular edge being taken when the clock has an irrational value. This suggests that classical techniques for partitioning the state space into a finite number of equivalence classes on the basis of a fixed, rational-numbered time granularity, such as the region graph [4] or the corner-point abstraction [8], cannot be applied directly to the case of cdPTA to obtain optimal reachability probabilities, because they rely on the fact that optimal choices can be made either at or arbitrarily closely to clock values that are multiples of the chosen rational-numbered time granularity. In Sect. 4, we present a conservative approximation method for cdPTA, i.e., maximum (respectively, minimum) probabilities are bounded from above (respectively, from below) in the approximation. This method is based on the region graph but uses concepts from the corner-point abstraction to define transition distributions. We show that successive refinement of the approximation, obtained by increasing the time granularity by a constant factor, does not lead to a more conservative approximation: in practice, in many cases such a refinement can lead to a substantial improvement in the computed probabilities, which we show using a small example.

2 Clock-Dependent Probabilistic Timed Automata

Preliminaries. We use $\mathbb{R}_{\geq 0}$ to denote the set of non-negative real numbers, $\mathbb{Q}$ to denote the set of rational numbers and $\mathbb{N}$ to denote the set of natural numbers. A (discrete) probability *distribution* over a countable set Q is a function $\mu : Q \rightarrow [0, 1]$ such that $\sum_{q \in Q} \mu(q) = 1$. For a function $\mu : Q \rightarrow \mathbb{R}_{\geq 0}$ we define $\mathsf{support}(\mu) = \{q \in Q : \mu(q) > 0\}$. Then for an uncountable set Q we define $\mathsf{Dist}(Q)$ to be the set of functions $\mu : Q \rightarrow [0, 1]$, such that $\mathsf{support}(\mu)$ is a countable set and μ restricted to $\mathsf{support}(\mu)$ is a (discrete) probability distribution. Given $q \in Q$, we use $\{q \mapsto 1\}$ to denote the distribution that assigns probability 1 to the single element q.

A *probabilistic transition system* (PTS) $\mathcal{T} = (S, \overline{s}, Act, \Delta)$ comprises the following components: a set S of *states* with an *initial state* $\overline{s} \in S$, a set Act of *actions*, and a *probabilistic transition relation* $\Delta \subseteq S \times Act \times \mathsf{Dist}(S)$. The sets of states, actions and the probabilistic transition relation can be uncountable. Transitions from state to state of a PTS are performed in two steps: if the current state is s, the first step concerns a nondeterministic selection of a probabilistic transition $(s, a, \mu) \in \Delta$; the second step comprises a probabilistic choice, made according to the distribution μ, as to which state to make the transition (that is, a transition to a state $s' \in S$ is made with probability $\mu(s')$). We denote such a completed transition by $s \xrightarrow{a,\mu} s'$. We assume that for each state $s \in S$ there exists some $(s, a, \mu) \in \Delta$.

An *infinite run* of the PTS $\mathcal{T}$ is an infinite sequence of consecutive transitions $r = s_0 \xrightarrow{a_0,\mu_0} s_1 \xrightarrow{a_1,\mu_1} \cdots$ (i.e., the target state of one transition is the source state of the next). Similarly, a *finite run* of $\mathcal{T}$ is a finite sequence of consecutive transitions $r = s_0 \xrightarrow{a_0,\mu_0} s_1 \xrightarrow{a_1,\mu_1} \cdots \xrightarrow{a_{n-1},\mu_{n-1}} s_n$. We use $\mathit{InfRuns}^{\mathcal{T}}$ to denote the set of infinite runs of $\mathcal{T}$, and $\mathit{FinRuns}^{\mathcal{T}}$ the set of finite runs of $\mathcal{T}$. If r is a finite run, we denote by $\mathit{last}(r)$ the last state of r. For any infinite run r and $i \in \mathbb{N}$, let $r(i) = s_i$ be the $(i+1)$th state along r. Let $\mathit{InfRuns}^{\mathcal{T}}(s)$ refer to the set of infinite runs of $\mathcal{T}$ commencing in state $s \in S$.

A *strategy* of a PTS $\mathcal{T}$ is a function σ mapping every finite run $r \in \mathit{FinRuns}^{\mathcal{T}}$ to a distribution in $\mathsf{Dist}(\Delta)$ such that $(s, a, \mu) \in \mathsf{support}(\sigma(r))$ implies that $s = \mathit{last}(r)$. From [11, Lemma 4.10], without loss of generality we can assume henceforth that strategies map to distributions assigning positive probability to finite sets of elements, i.e., strategies σ for which $|\mathsf{support}(\sigma(r))|$ is finite for all $r \in \mathit{FinRuns}^{\mathcal{T}}$. For any strategy σ, let $\mathit{InfRuns}^{\sigma}$ denote the set of infinite runs resulting from the choices of σ. For a state $s \in S$, let $\mathit{InfRuns}^{\sigma}(s) = \mathit{InfRuns}^{\sigma} \cap \mathit{InfRuns}^{\mathcal{T}}(s)$. Given a strategy σ and a state $s \in S$, we define the probability measure Pr^{σ}_{s} over $\mathit{InfRuns}^{\sigma}(s)$ in the standard way [14].

Given a set $S_F \subseteq S$, define $\Diamond S_F = \{r \in \mathit{InfRuns}^{\mathcal{T}} : \exists i \in \mathbb{N} \text{ s.t. } r(i) \in S_F\}$ to be the set of infinite runs of $\mathcal{T}$ such that some state of S_F is visited along the run. Given a set $\Sigma' \subseteq \Sigma$ of strategies, we define the *maximum value over* Σ' *with respect to* S_F as $\mathbb{P}^{\max}_{\mathcal{T},\Sigma'}(S_F) = \sup_{\sigma \in \Sigma'} \mathrm{Pr}^{\sigma}_{\overline{s}}(\Diamond S_F)$. Similarly, the *minimum value over* Σ' *with respect to* S_F is defined as $\mathbb{P}^{\min}_{\mathcal{T},\Sigma'}(S_F) = \inf_{\sigma \in \Sigma'} \mathrm{Pr}^{\sigma}_{\overline{s}}(\Diamond S_F)$. The *maximal reachability problem* for $\mathcal{T}$, $S_F \subseteq S$, $\Sigma' \subseteq \Sigma$, $\trianglerighteq \in \{\geq, >\}$ and

$\lambda \in [0,1]$ is to decide whether $\mathbb{P}^{\max}_{\mathcal{T},\Sigma'}(S_F) \trianglerighteq \lambda$. Similarly, the *minimal reachability problem* for $\mathcal{T}$, $S_F \subseteq S$, $\Sigma' \subseteq \Sigma$, $\trianglelefteq \in \{\leq, <\}$ and $\lambda \in [0,1]$ is to decide whether $\mathbb{P}^{\min}_{\mathcal{T},\Sigma'}(S_F) \trianglelefteq \lambda$.

Clock-Dependent Probabilistic Timed Automata. Let $\mathcal{X}$ be a finite set of real-valued variables called *clocks*, the values of which increase at the same rate as real-time and which can be reset to 0. A function $v : \mathcal{X} \to \mathbb{R}_{\geq 0}$ is referred to as a *clock valuation* and the set of all clock valuations is denoted by $\mathbb{R}^{\mathcal{X}}_{\geq 0}$. For $v \in \mathbb{R}^{\mathcal{X}}_{\geq 0}$, $t \in \mathbb{R}_{\geq 0}$ and $X \subseteq \mathcal{X}$, we use $v{+}t$ to denote the clock valuation that increments all clock values in v by t, and $v[X{:=}0]$ to denote the clock valuation in which clocks in X are reset to 0.

For a set Q, a *distribution template* $\mathfrak{d} : \mathbb{R}^{\mathcal{X}}_{\geq 0} \to \mathsf{Dist}(Q)$ gives a distribution over Q for each clock valuation. In the following, we use notation $\mathfrak{d}[v]$, rather than $\mathfrak{d}(v)$, to denote the distribution corresponding to distribution template $\mathfrak{d}$ and clock valuation v. Let $\mathfrak{Dist}(Q)$ be the set of distribution templates over Q.

The set $CC(\mathcal{X})$ of *clock constraints* over $\mathcal{X}$ is defined as the set of conjunctions over atomic formulae of the form $x \sim c$, where $x \in \mathcal{X}$, $\sim \in \{<, \leq, \geq, >\}$, and $c \in \mathbb{N}$. A clock valuation v satisfies a clock constraint ψ, denoted by $v \models \psi$, if ψ resolves to `true` when substituting each occurrence of clock x with $v(x)$.

A *clock-dependent probabilistic timed automaton* (cdPTA) $\mathcal{P} = (L, \bar{l}, \mathcal{X}, inv, prob)$ comprises the following components: a finite set L of *locations* with an *initial location* $\bar{l} \in L$; a finite set $\mathcal{X}$ of clocks; a function $inv : L \to CC(\mathcal{X})$ associating an *invariant condition* with each location; a set $prob \subseteq L \times CC(\mathcal{X}) \times \mathfrak{Dist}(2^{\mathcal{X}} \times L)$ of *probabilistic edges*. A probabilistic edge $(l, g, \mathfrak{p}) \in prob$ comprises: (1) a source location l; (2) a clock constraint g, called a *guard*; and (3) a distribution template $\mathfrak{p}$ with respect to pairs of the form $(X, l') \in 2^{\mathcal{X}} \times L$ (i.e., pairs consisting of a set X of clocks to be reset and a target location l').

The behaviour of a cdPTA takes a similar form to that of a standard probabilistic timed automaton [10,17]: in any location time can advance as long as the invariant holds, and the choice as to how much time elapses is made nondeterministically; a probabilistic edge can be taken if its guard is satisfied by the current values of the clocks and, again, the choice as to which probabilistic edge to take is made nondeterministically; for a taken probabilistic edge, the choice of which clocks to reset and which target location to make the transition to is *probabilistic*. The key difference with cdPTAs is that the distribution used to make this probabilistic choice depends on the probabilistic edge taken *and* on the current clock valuation.

Example 1. In Fig. 1 we give an example of a cdPTA modelling a simple robot that must reach a certain geographical area and then carry out a particular task. The usual conventions for the graphical representation of timed automata are used in the figure. Black squares denote the distributions of probabilistic edges, and expressions on probabilities used by distribution templates are written with a grey background on their outgoing arcs. The robot can be in one of four

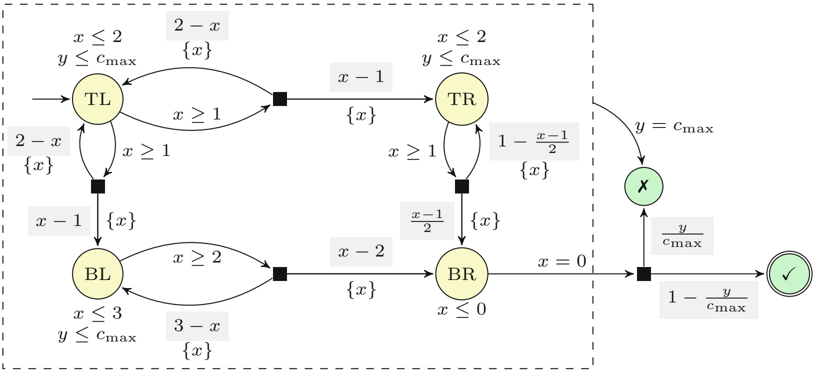

Fig. 1. A cdPTA modelling a simple robot example.

geographical areas, which can be thought of as cells in a 2 × 2 grid, each of which corresponds to a cdPTA location. The robot begins in the top-left cell (corresponding to location TL), and its objective is to reach the bottom-right cell (location BR). The robot can move either to the top-right cell (location TR), or to the bottom-left cell (location BL), then to the bottom-right cell. In each cell, the robot must wait a certain amount of time (1 time units in the top cells and 2 time units in the bottom-left cell) before attempting to leave the cell (for example, to recharge solar batteries), after which it can spend at most 1 time unit attempting to leave the cell. With a certain probability, the attempt to leave the cell will fail, and the robot must wait before trying to leave the cell again; the more time is dedicated to leaving the cell, the more likely the robot will succeed. Although passing through the top-right cell is not slower than passing through the bottom-left cell, the probability of leaving the cell successfully increases at a slower rate than in other cells (representing, for example, terrain in which the robot finds it difficult to navigate). On arrival in the bottom-right cell, the robot successfully carries out its task with a probability that is inversely proportional to the total time elapsed (for example, the robot could be transporting medical supplies, the efficacy of which may be inversely proportional to the time elapsed). The clock x is used to represent the amount of time used by the robot in its attempt to move from cell to cell, whereas the clock y represents the total amount of time since the start of the robot's mission. If the clock y reaches its maximum amount $c_{\max}$, then the mission fails (as denoted by the edge to the location denoted by ✗, which is available in locations TL, TR, BL and BR, as indicated by the dashed box). The objective of the robot's controller is to maximise the probability of reaching the location denoted by ✓. Note that there is a trade-off between dedicating more time to movement between the cells, which increases the probability of successful navigation and therefore progress towards the target point, and spending less time on the overall mission, which increases the probability of carrying out the required task at the target point. □

A *state* of a cdPTA is a pair comprising a location and a clock valuation satisfying the location's invariant condition, i.e., $(l, v) \in L \times \mathbb{R}_{\geq 0}^{\mathcal{X}}$ such that $v \models inv(l)$. In any state (l, v), either a certain amount of time $\delta \in \mathbb{R}_{\geq 0}$ elapses, or a probabilistic edge is traversed. If time elapses, then the choice of δ requires that the invariant $inv(l)$ remains continuously satisfied while time passes. The resulting state after this transition is $(l, v+\delta)$. A probabilistic edge $(l', g, \mathfrak{p}) \in prob$ can be chosen from (l, v) if $l = l'$ and it is *enabled*, i.e., the clock constraint g is satisfied by v. Once a probabilistic edge $(l, g, \mathfrak{p})$ is chosen, a set of clocks to reset and a successor location are selected at random, according to the distribution $\mathfrak{p}[v]$.

We make a number of assumptions concerning the cdPTA models considered. Firstly, we restrict our attention to cdPTAs for which it is always possible to take a probabilistic edge, either immediately or after letting time elapse. This condition holds generally for PTA models in practice [16]. A sufficient syntactic condition for this property has been presented formally in [12]. Secondly, we consider cdPTAs that feature invariant conditions that prevent clock values from exceeding some bound: formally, for each location $l \in L$, we have that $inv(l)$ contains a constraint of the form $x \leq c$ or $x < c$ for each clock $x \in \mathcal{X}$. Thirdly, we assume that all possible target states of probabilistic edges satisfy their invariants: for all probabilistic edges $(l, g, \mathfrak{p}) \in prob$, for all clock valuations $v \in \mathbb{R}_{\geq 0}^{\mathcal{X}}$ such that $v \models g$, and for all $(X, l') \in 2^{\mathcal{X}} \times L$, we have that $\mathfrak{p}[v](X, l') > 0$ implies $v[X := 0] \models inv(l')$. Finally, we assume that any clock valuation that satisfies the guard of a probabilistic edge also satisfies the invariant of the source location: this can be achieved, without changing the underlying semantic PTS, by replacing each probabilistic edge $(l, g, \mathfrak{p}) \in prob$ by $(l, g \wedge inv(l), \mathfrak{p})$.

Let $\mathbf{0} \in \mathbb{R}_{\geq 0}^{\mathcal{X}}$ be the clock valuation which assigns 0 to all clocks in $\mathcal{X}$. The semantics of the cdPTA $\mathcal{P} = (L, \bar{l}, \mathcal{X}, inv, prob)$ is the PTS $[\![\mathcal{P}]\!] = (S, \bar{s}, Act, \Delta)$ where:

- $S = \{(l, v) : l \in L \text{ and } v \in \mathbb{R}_{\geq 0}^{\mathcal{X}} \text{ s.t. } v \models inv(l)\}$ and $\bar{s} = \{(\bar{l}, \mathbf{0})\}$;
- $Act = \mathbb{R}_{\geq 0} \cup prob$;
- $\Delta = \overrightarrow{\Delta} \cup \widehat{\Delta}$, where $\overrightarrow{\Delta} \subseteq S \times \mathbb{R}_{\geq 0} \times \mathsf{Dist}(S)$ and $\widehat{\Delta} \subseteq S \times prob \times \mathsf{Dist}(S)$ such that:
  - $\overrightarrow{\Delta}$ is the smallest set such that $((l, v), \delta, \{(l, v+\delta) \mapsto 1\}) \in \overrightarrow{\Delta}$ if there exists $\delta \in \mathbb{R}_{\geq 0}$ such that $v + \delta' \models inv(l)$ for all $0 \leq \delta' \leq \delta$;
  - $\widehat{\Delta}$ is the smallest set such that $((l, v), (l, g, \mathfrak{p}), \mu) \in \widehat{\Delta}$ if
    1. $v \models g$;
    2. for any $(l', v') \in S$, we have $\mu(l', v') = \sum_{X \in \mathsf{Reset}(v, v')} \mathfrak{p}[v](X, l')$, where $\mathsf{Reset}(v, v') = \{X \subseteq \mathcal{X} \mid v[X := 0] = v'\}$.

When considering maximum and minimum values for cdPTAs, we henceforth consider strategies that alternate between transitions from $\overrightarrow{\Delta}$ (time elapse transitions) and transitions from $\widehat{\Delta}$ (probabilistic edge transitions). Formally, a *cdPTA strategy* σ is a strategy such that, for a finite run $r \in FinRuns^{[\![\mathcal{P}]\!]}$ that has $s \xrightarrow{a,\mu} s'$ as its final transition, either $(s, a, \mu) \in \overrightarrow{\Delta}$ and $\mathsf{support}(\sigma(r)) \in \widehat{\Delta}$, or $(s, a, \mu) \in \widehat{\Delta}$ and $\mathsf{support}(\sigma(r)) \in \overrightarrow{\Delta}$. We write $\mathbf{\Sigma}$ for the set of cdPTA strategies

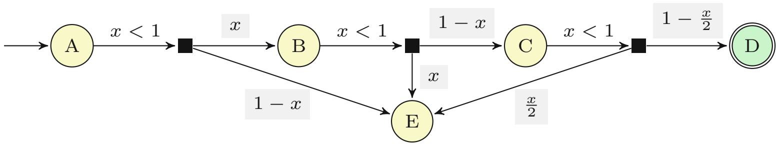

Fig. 2. A one-clock cdPTA for which the maximum probability is attained by a time delay corresponding to an irrational number.

of $[\![\mathcal{P}]\!]$. Given a set $F \subseteq L$ of locations, subsequently called *target locations*, we let $S_F = \{(l, v) \in S : l \in F\}$. Let $\trianglerighteq \in \{\geq, >\}$, $\trianglelefteq \in \{\leq, <\}$ and $\lambda \in [0, 1]$: then the maximal (respectively, minimal) reachability problem for cdPTA is to decide whether $\mathbb{P}^{\max}_{[\![\mathcal{P}]\!],\mathbf{\Sigma}}(S_F) \trianglerighteq \lambda$ (respectively, $\mathbb{P}^{\min}_{[\![\mathcal{P}]\!],\mathbf{\Sigma}}(S_F) \trianglelefteq \lambda$).

Piecewise Linear Clock Dependencies. In this paper, we concentrate on a particular subclass of distribution templates based on continuous piecewise linear functions. Let $x \in \mathcal{X}$ be a clock and $p = (l, g, \mathfrak{p}) \in \mathit{prob}$ be a probabilistic edge. Let I^p_x be the interval containing the values of x of clock valuations that satisfy g: formally $I^p_x = \{v(x) \in \mathbb{R}_{\geq 0} : v \in \mathbb{R}^{\mathcal{X}}_{\geq 0} \text{ s.t. } v \models g\}$. For example, for $g = (x \geq 3) \wedge (x < 5) \wedge (y \leq 8)$, we have $I^p_x = [3, 5)$ and $I^p_y = [0, 8]$. We equip each probabilistic edge $p = (l, g, \mathfrak{p}) \in \mathit{prob}$ and $e = (X, l') \in 2^{\mathcal{X}} \times L$ with a continuous piecewise linear function $f^{p,e}_x$ with domain I^p_x for each clock $x \in \mathcal{X}$. Formally, we consider a partition $\mathcal{I}^{p,e}_x$ of I^p_x (i.e., $\bigcup_{I \in \mathcal{I}^{p,e}_x} I = I^p_x$ and $I \cap I' = \emptyset$ for each $I, I' \in \mathcal{I}^{p,e}_x$ such that $I \neq I'$), and sets $\{c^{p,e}_{x,I}\}_{I \in \mathcal{I}^{p,e}_x}$ and $\{d^{p,e}_{x,I}\}_{I \in \mathcal{I}^{p,e}_x}$ of constants in $\mathbb{Q}$ such that: (a) for every $I \in \mathcal{I}^{p,e}_x$ and $\gamma \in I$, we have $f^{p,e}_x(\gamma) = c^{p,e}_{x,I} + d^{p,e}_{x,I} \cdot \gamma$; (b) $f^{p,e}_x$ is continuous (i.e., for each $\gamma \in I^p_x$, we have $\lim_{\zeta \to \gamma} f^{p,e}_x(\zeta) = f^{p,e}_x(\gamma)$). We make the following assumptions for each probabilistic edge $p \in \mathit{prob}$: (1) all endpoints of intervals in $\mathcal{I}^{p,e}_x$ are natural numbers, for all clocks $x \in \mathcal{X}$ and $e \in 2^{\mathcal{X}} \times L$; (2) $\sum_{x \in \mathcal{X}} f^{p,e}_x(v(x)) \in [0, 1]$ for each $e \in 2^{\mathcal{X}} \times L$ and $v \in \mathbb{R}^{\mathcal{X}}_{\geq 0}$ such that $v \models g$; (3) $\sum_{e \in 2^{\mathcal{X}} \times L} \sum_{x \in \mathcal{X}} f^{p,e}_x(v(x)) = 1$ for each $v \in \mathbb{R}^{\mathcal{X}}_{\geq 0}$ such that $v \models g$. Then the probabilistic edge p is *piecewise linear* if, for each $e \in 2^{\mathcal{X}} \times L$ and each $v \in \mathbb{R}^{\mathcal{X}}_{\geq 0}$ such that $v \models g$, we have $\mathfrak{p}[v](e) = \sum_{x \in \mathcal{X}} f^{p,e}_x(v(x))$. We assume henceforth that all probabilistic edges of cdPTAs are piecewise linear.

Example 2. Standard methods for the analysis of timed automata typically consist of a finite-state system that represents faithfully the original model. In particular, the region graph [4] and the corner-point abstraction [8] both involve the division of the state space according to a fixed, rational-numbered granularity. The example of a one-clock cdPTA $\mathcal{P}$ of Fig. 2 shows that such an approach cannot be used for the exact computation of optimal reachability probabilities in cdPTAs, because optimality may be attained when the clock has an irrational value. For an example of the formal description of a piecewise linear probabilistic edge, consider the probabilistic edge from location C, which we denote by p_{C}: then we have $\mathcal{I}^{p_{\mathrm{C}},(\emptyset,\mathrm{D})}_x = \mathcal{I}^{p_{\mathrm{C}},(\emptyset,\mathrm{E})}_x = \{[0, 1)\}$, with $c^{p_{\mathrm{C}},(\emptyset,\mathrm{D})}_{x,[0,1)} = 1$, $d^{p_{\mathrm{C}},(\emptyset,\mathrm{D})}_{x,[0,1)} = -\frac{1}{2}$, $c^{p_{\mathrm{C}},(\emptyset,\mathrm{E})}_{x,[0,1)} = 0$, and $d^{p_{\mathrm{C}},(\emptyset,\mathrm{E})}_{x,[0,1)} = \frac{1}{2}$. Now consider the maximum probability of reaching location D (that is, $\mathbb{P}^{\max}_{[\![\mathcal{P}]\!],\mathbf{\Sigma}}(S_{\{\mathrm{D}\}})$). Intuitively, the longer the cdPTA remains in location A, the lower the probability of making a transition to location E from A, but the higher the probability of making a transition to E from B and C. Note that, after A is left, the choice resulting in the maximum probability of reaching D is to take the outgoing transitions from B and C as soon as possible (delaying in B and C will increase the value of x, therefore increasing the probability of making a transition to E). Denoting by δ the amount of time elapsed in A, the maximum probability of reaching D is equal to $\delta(1-\delta)(1-\frac{\delta}{2})$, which (within the interval $[0,1)$) reaches its maximum at $1-\frac{\sqrt{3}}{3}$. Hence, this example indicates that abstractions based on the optimality of choices made at (or arbitrarily close to) rational-numbered clock values (such as the region graph or corner-point abstraction) do not yield exact analysis methods for cdPTAs. □

3 Undecidability of Maximal Reachability of cdPTAs

Theorem 1. *The maximal reachability problem is undecidable for cdPTAs with at least 3 clocks.*

Proof (sketch). We proceed by reducing the non-halting problem for two-counter machines to the maximal reachability problem for cdPTAs. The reduction has close similarities to a reduction presented in [2].

A two-counter machine $\mathcal{M} = (\mathcal{L}, \mathcal{C})$ comprises a set $\mathcal{L} = \{\ell_1, ..., \ell_n\}$ of instructions and a set $\mathcal{C} = \{\mathsf{c}_1, \mathsf{c}_2\}$ of counters. The instructions are of the following form (for $1 \le i, j, k \le n$ and $l \in \{1, 2\}$):

1. $\ell_i : \mathsf{c}_l := \mathsf{c}_l + 1;$ goto ℓ_j (increment c_l);
2. $\ell_i : \mathsf{c}_l := \mathsf{c}_l - 1;$ goto ℓ_j (decrement c_l);
3. $\ell_i :$ if $(\mathsf{c}_l > 0)$ them goto ℓ_j else goto ℓ_k (zero check c_l);
4. ℓ_n : HALT (halting instruction).

A configuration $(\ell, \mathsf{v}_1, \mathsf{v}_2)$ of a two-counter machine comprises an instruction ℓ and values v_1 and v_2 of counters c_1 and c_2, respectively. A run of a two-counter machine consists of a finite or infinite sequence of configurations, starting from configuration $(\ell_1, 0, 0)$, and where subsequent configurations are successively generated by following the rule specified in the associated configuration. A run is finite if and only if the final instruction visited along the run is ℓ_n (the halting instruction). The halting problem for two-counter machines concerns determining whether the unique run of the two-counter machine is finite, and is undecidable [18]; hence the non-halting problem (determining whether the unique run of the two-counter machine is infinite) is also undecidable.

Consider a two-counter machine $\mathcal{M}$. We reduce the non-halting problem for $\mathcal{M}$ to the maximal reachability problem in the following way. We construct a cdPTA $\mathcal{P}_{\mathcal{M}}$ with three clocks $\{x_1, x_2, x_3\}$ by considering modules for each form that the instructions of a two-counter machine can take. On entry to each module, we have that $x_1 = \frac{1}{2^{\mathsf{c}_1}}$, $x_2 = \frac{1}{2^{\mathsf{c}_2}}$ and $x_3 = 0$. The module for simulating an increment instruction is shown in Fig. 3. In location ℓ_i, there is a delay of

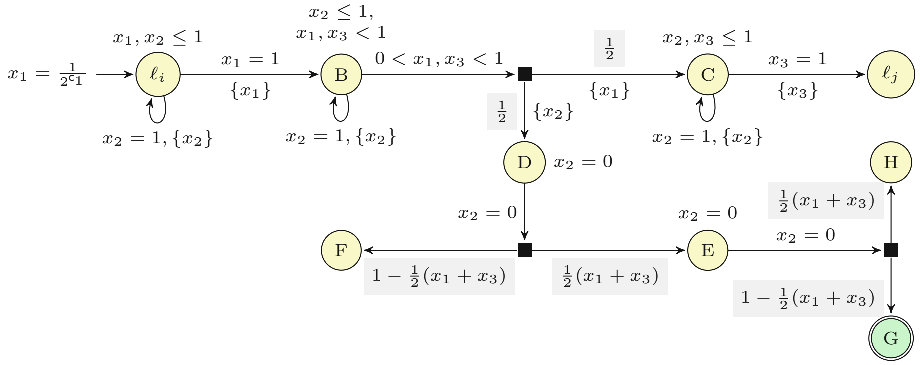

Fig. 3. The cdPTA module for simulating an increment instruction for counter c_1.

$1 - \frac{1}{2^{\mathsf{c}_1}}$, and hence the values of the clocks on entry to location B are $x_1 = 0$, $x_2 = \frac{1}{2^{\mathsf{c}_2}} + 1 - \frac{1}{2^{\mathsf{c}_1}} \mod 1$ and $x_3 = 1 - \frac{1}{2^{\mathsf{c}_1}}$. A nondeterministic choice is then made concerning the amount of time that elapses in location B: note that this amount must be in the interval $(0, \frac{1}{2^{\mathsf{c}_1}})$. In order to correctly simulate the increment of counter c_1, the choice of delay in location B should be equal to $\frac{1}{2^{\mathsf{c}_1+1}}$. On leaving location B, a probabilistic choice is made: the rightward outcome corresponds to continuing the simulation of the two-counter machine, whereas the downward outcome corresponds to checking that the delay in location B was correctly $\frac{1}{2^{\mathsf{c}_1+1}}$. We write the delay in location B as $\frac{1}{2^{\mathsf{c}_1+1}} + \epsilon$, where $-\frac{1}{2^{\mathsf{c}_1+1}} < \epsilon < \frac{1}{2^{\mathsf{c}_1+1}}$: hence, for a correct simulation of the increment of c_1, we require that $\epsilon = 0$.

Consider the case in which the downward outcome (from the outgoing probabilistic edge of location B) is taken: then the cdPTA fragment from location D has the role of checking whether $\epsilon = 0$. Note that, after entering location D, no time elapses in locations D and E (as enforced by the reset of x_2 to zero and the invariant condition $x_2 = 0$), and hence both clocks x_1 and x_3 retain the same values that they had when location B was left. We show that the probability of reaching the target location G from location D is $\frac{1}{4} - \epsilon^2$, and hence equal to $\frac{1}{4}$ if and only if $\epsilon = 0$. To see that the probability of reaching G from D is $\frac{1}{4} - \epsilon^2$, observe that the probability is equal to $\frac{1}{2}(x_1 + x_3) = \frac{1}{2}(\frac{1}{2^{\mathsf{c}_1+1}} + \epsilon + (1 - \frac{1}{2^{\mathsf{c}_1+1}}) + \epsilon) = \frac{1}{2} + \epsilon$ multiplied by $1 - \frac{1}{2}(x_1 + x_3) = \frac{1}{2} - \epsilon$, i.e., equal to $\frac{1}{4} - \epsilon^2$. Hence the probability of reaching location G from location D is equal to $\frac{1}{4}$ if and only if $\epsilon = 0$ (otherwise, the probability is less than $\frac{1}{4}$).

The module for simulating a decrement instruction is shown in Fig. 4. In a similar manner to the cdPTA fragment in Fig. 3 for the simulation of an increment instruction, the only nondeterministic choice made is with regard to the amount of time spent in location ℓ_i, which is denoted by δ. For the correct simulation of the decrement instruction, δ should equal $1 - \frac{1}{2^{\mathsf{c}_1-1}}$. The rightward outcome is taken from the probabilistic edge leaving location ℓ_i corresponds to the continuation of the simulation of the two-counter machine: hence, on entry

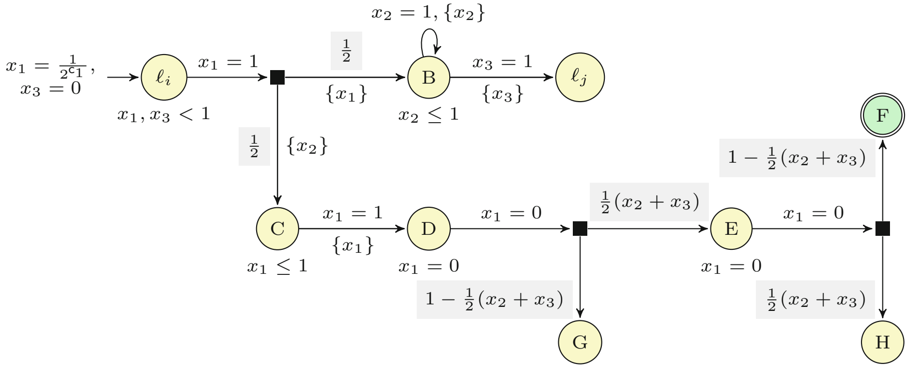

Fig. 4. The cdPTA module for simulating a decrement instruction for counter c_1.

to location B, we have $x_1 = 0$, $x_2 = \frac{1}{2^{\mathsf{c}_2}} + \delta$ and $x_3 = \delta$; then, on entry to location ℓ_j, we have $x_1 = \delta$, $x_2 = \frac{1}{2^{\mathsf{c}_2}}$ and $x_3 = 0$.

Let $\delta = 1 - \frac{1}{2^{\mathsf{c}_1 - 1}} + \epsilon$. For the correct simulation of the decrement instruction, we require that $\epsilon = 0$. The downward outcome from the probabilistic edge leaving location ℓ_i corresponds to checking that $\epsilon = 0$, and takes a similar form to the analogous downward edge of the cdPTA fragment for the increment instruction, as shown in Fig. 3. Note that, on entry to location C, we have that $x_1 = 1 - \frac{1}{2^{\mathsf{c}_1}} + \epsilon$, $x_2 = 0$ and $x_3 = 1 - \frac{1}{2^{\mathsf{c}_1 - 1}} + \epsilon$. Then, on entry to location D, we have that $x_1 = 0$, $x_2 = \frac{1}{2^{\mathsf{c}_1}} - \epsilon$ and $x_3 = 1 - \frac{1}{2^{\mathsf{c}_1}}$. As no time elapses in locations D and E, we have that target location F is then reached with probability $\frac{1}{2}(x_2 + x_3) = \frac{1}{2}(\frac{1}{2^{\mathsf{c}_1}} - \epsilon + 1 - \frac{1}{2^{\mathsf{c}_1}}) = \frac{1}{2} + \frac{\epsilon}{2}$ multiplied by the probability $1 - \frac{1}{2}(x_2 + x_3) = \frac{1}{2} - \frac{\epsilon}{2}$, which equals $\frac{1}{4} - \frac{\epsilon^2}{4}$. Hence we conclude that the probability of reaching location F from location C is equal to $\frac{1}{4}$ if and only if $\epsilon = 0$.

Finally, the module for a zero test instruction $\ell_i :$ if $(\mathsf{c}_1 > 0)$ then goto ℓ_j else goto ℓ_k is shown in Fig. 5. The module is almost identical to that of [3], and we present it here only for completeness. After entry to location ℓ_i, two probabilistic edges are enabled: the rightward one is taken if $\mathsf{c}_1 = 0$ (i.e., if $x_1 = \frac{1}{2^0} = 1$), whereas the leftward one is taken otherwise. Both probabilistic edges involve an outcome leading to a target location with probability $\frac{1}{4}$: if this outcome is not taken, the cdPTA fragment then proceeds to location ℓ_j or ℓ_j, depending on which probabilistic edge was taken.

Given the construction of a cdPTA simulating the two-counter machine using the modules described above, we can now proceed to show Theorem 1. The reasoning is the same as that of Lemma 5 of [2]. If the two-counter machine halts in k steps, and the strategy of the cdPTA correctly simulates the two-counter machine the probability of reaching a target location will be $\frac{1}{2} \cdot \frac{1}{4} + (\frac{1}{2})^2 \cdot \frac{1}{4} + ... + (\frac{1}{2})^k \cdot \frac{1}{4} < \frac{1}{4}$. If the two-counter machine halts in k steps, and the strategy of the cdPTA does not correctly simulate the two-counter machine, then this means that the probability of reaching a target location is strictly less than that corresponding to correct simulation, given that deviation from simulation of a certain step

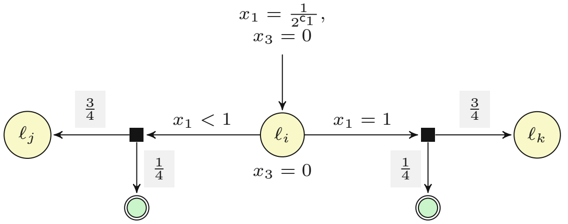

Fig. 5. The cdPTA module for simulating a zero-test instruction for counter c_1.

corresponds to reaching the target locations with probability strictly less than $\frac{1}{4}$ in that step. Now consider the case in which the two-counter machine does not halt: in this case, faithful simulation in the cdPTA corresponds to reaching target locations with probability $\sum_{i=1}^{\infty}(\frac{1}{2})^i \cdot \frac{1}{4} = \frac{1}{4}$, whereas unfaithful simulation in the cdPTA corresponds to reaching the target locations with probability $\sum_{i=1}^{\infty}(\frac{1}{2})^i \cdot \gamma_i$ where $\gamma_i \leq \frac{1}{4}$ for all $i \in \mathbb{N}$ and $\gamma_j < \frac{1}{4}$ for at least one $j \in \mathbb{N}$, and hence $\sum_{i=1}^{\infty}(\frac{1}{2})^i \cdot \gamma_i < \frac{1}{4}$. Therefore the two-counter machine does not halt if and only if there exists a strategy in the constructed cdPTA that reaches the target locations with probability at least $\frac{1}{4}$, concluding the proof of Theorem 1. □

4 Approximation of Reachability Probabilities

We now consider the approximation of maximal and minimal reachability probabilities of cdPTAs. Our approach is to utilise concepts from the corner-point abstraction [8]. However, while the standard corner-point abstraction is a finite-state system that extends the classical region graph by encoding corner points within states, the states of our finite-state system correspond to regions, and we use corners of regions only to define available distributions. Furthermore, in contrast to the widespread use of the corner-point abstraction in the context of weighted (or priced) timed automata (see [7] for a survey), and in line with the undecidability results presented in Sect. 3, our variant of the corner-point abstraction does not result in a finite-state system that can be used to obtain a quantitative measure that is arbitrarily close to the actual one: in the context of cdPTAs, we will present a method that approximates maximal and minimal reachability properties, and show that successive refinement of regions leads to a more accurate approximation.

First we define regions and corner points. Let $\mathcal{P} = (L, \bar{l}, \mathcal{X}, \mathit{inv}, \mathit{prob})$ be a cdPTA, which we assume to be fixed throughout this section, and let $M \in \mathbb{N}$ denote the upper bound on clocks in $\mathcal{P}$. We choose $k \in \mathbb{N}$, which we will refer to as the *(time) granularity*, and let $[k] = \{\frac{c}{k} : c \in \mathbb{N}\}$ be the set of multiples of $\frac{1}{k}$. A *k-region* $(h, [X_0, ..., X_n])$ over $\mathcal{X}$ comprises:

1. a function $h : \mathcal{X} \to ([k] \cap [0, M])$ assigning a multiple of $\frac{1}{k}$ no greater than M to each clock and

2. a partition $[X_0, ..., X_n]$ of $\mathcal{X}$, where $X_i \neq \emptyset$ for all $1 \leq i \leq n$ and $h(x) = M$ implies $x \in X_0$ for all $x \in \mathcal{X}$.

Given clock valuation $v \in \mathbb{R}_{\geq 0}^{\mathcal{X}}$ and granularity k, the *k-region* $R = (h, [X_0, ..., X_n])$ *containing* v (written $v \in R$) satisfies the following conditions:

1. $\lfloor k \cdot v(x) \rfloor = k \cdot h(x)$ for all clocks $x \in \mathcal{X}$;
2. $v(x) = h(x)$ for all clocks $x \in X_0$;
3. $k \cdot v(x) - \lfloor k \cdot v(x) \rfloor \leq k \cdot v(y) - \lfloor k \cdot v(y) \rfloor$ if and only if $x \in X_i$ and $y \in X_j$ with $i \leq j$, for all clocks $x, y \in \mathcal{X}$.

Note that, rather than considering regions delimited by valuations corresponding to natural numbers, in our definition regions are delimited by valuations corresponding to multiples of $\frac{1}{k}$. We use Regs_k to denote the set of k-regions. For $R, R' \in \mathsf{Regs}_k$ and clock constraint $\psi \in CC(\mathcal{X})$, we say that R' is a *ψ-satisfying time successor* of R if there exist $v \in R$ and $\delta \in \mathbb{R}_{\geq 0}$ such that $(v+\delta) \in R'$ and $(v+\delta') \models \psi$ for all $0 \leq \delta' \leq \delta$. For a given k-region $R \in \mathsf{Regs}_k$, we let $R[X := 0]$ be the k-region that corresponds to resetting clocks in X to 0 from clock valuations in R (that is, $R[X := 0]$ contains valuations $v[X := 0]$ for $v \in R$). We use $R_{\mathbf{0}}$ to denote the k-region that contains the valuation $\mathbf{0}$.

A *corner point* $\alpha = \langle a_i \rangle_{0 \leq i \leq n} \in ([k] \cap [0, M])^n$ of k-region $(h, [X_0, ..., X_n])$ is defined by:

$$a_i(x) = \begin{cases} h(x) & \text{if } x \in X_j \text{ with } j \leq i \\ h(x) + \frac{1}{k} & \text{if } x \in X_j \text{ with } j > i\ . \end{cases}$$

Note that a k-region $(h, [X_0, ..., X_n])$ is associated with $n+1$ corner points. Let $\mathsf{CP}(R)$ be the set of corner points of k-region R. Given granularity k, we let $\mathsf{CornerPoints}_k$ be the set of all corner points.

Next we define the *clock-dependent region graph with granularity k* as the finite-state PTS $\mathcal{A}_k = (\mathsf{S}_k, \bar{\mathsf{s}}, \mathsf{Act}_k, \mathit{\Gamma}_k)$, where $\mathsf{S}_k = L \times \mathsf{Regs}_k$, $\bar{\mathsf{s}} = (\bar{l}, R_{\mathbf{0}})$, $\mathsf{Act}_k = \{\tau\} \cup (\mathsf{CornerPoints}_k \times \mathit{prob})$, and $\mathit{\Gamma}_k = \overrightarrow{\mathit{\Gamma}_k} \cup \widehat{\mathit{\Gamma}_k}$ where $\overrightarrow{\mathit{\Gamma}_k} \subseteq \mathsf{S}_k \times \{\tau\} \times \mathsf{Dist}(\mathsf{S}_k)$ and $\widehat{\mathit{\Gamma}_k} \subseteq \mathsf{S}_k \times \mathsf{CornerPoints}_k \times \mathit{prob} \times \mathsf{Dist}(\mathsf{S}_k)$ such that:

- $\overrightarrow{\mathit{\Gamma}_k}$ is the smallest set of transitions such that $((l, R), \tau, \{(l, R') \mapsto 1\}) \in \overrightarrow{\mathit{\Gamma}_k}$ if (l, R') is an $\mathit{inv}(l)$-satisfying time successor of (l, R);
- $\widehat{\mathit{\Gamma}_k}$ is the smallest set such that $((l, R), (\alpha, (l, g, \mathfrak{p})), \nu) \in \widehat{\mathit{\Gamma}_k}$ if:
  1. $R \models g$;
  2. $\alpha \in \mathsf{CP}(R)$;
  3. for any $(l', R') \in \mathsf{S}_k$, we have that $\nu(l', R') = \sum_{X \in \mathsf{Reset}(R, R')} \mathfrak{p}[\alpha](X, l')$, where $\mathsf{Reset}(R, R') = \{X \subseteq \mathcal{X} \mid R[X := 0] = R'\}$.

Hence the clock-dependent region graph of a cdPTA encodes corner points within (probabilistic-edge-based) transitions, in contrast to the corner-point abstraction, which encodes corner points within states. In fact, a literal application of the standard corner-point abstraction, as presented in [7], does not result in a conservative approximation, which we now explain with reference to Example 2.

Example 2 (continued). Recall that the states of the corner-point abstraction comprise a location, a region and a corner point of the region, and transitions maintain consistency between corner points of the source and target states. For example, for the cdPTA of Fig. 2, consider the state $(\mathrm{A}, 0 < x < 1, x = 1)$, where $0 < x < 1$ is used to refer to the state's region component and $x = 1$ is used to refer to the state's corner point. Then the probabilistic edge leaving location A is enabled (because the state represents the situation in which clock x is in the interval $(0, 1)$ and arbitrarily close to 1). Standard intuition on the corner-point abstraction (adapted from weights in [7] to probabilities in distribution templates in this paper) specifies that, when considering probabilities of outgoing probabilistic edges, the state $(\mathrm{A}, 0 < x < 1, x = 1)$ should be associated with probabilities for which $x = 1$. Hence the probability of making a transition to location B is 1, and the target corner-point-abstraction state is $(\mathrm{B}, 0 < x < 1, x = 1)$. However, now consider the probabilistic edge leaving location B: in this case, given that the corner point under consideration is $x = 1$, the probability of making a transition to location C is 0, and hence the target location D is reachable with probability 0. Furthermore, consider the state $(\mathrm{A}, 0 < x < 1, x = 0)$: in this case, if the probabilistic edge leaving location A is taken, then location B is reached with probability 0, and hence location D is again reachable with probability 0. We can conclude that such a direct application of the corner-point abstraction to cdPTA is not a conservative approximation of the cdPTA, because the maximum reachability probability in the corner-point abstraction is 0, i.e., less than the maximum reachability probability of the cdPTA (which we recall is $1 - \frac{\sqrt{3}}{3}$). Instead, in our definition of the clock-dependent region graph, we allow "inconsistent" corner points to be used in successive transitions: for example, from location A, the outgoing probabilistic edge can be taken using the value of x corresponding to the corner point $x = 1$; then, from locations B and C, the outgoing probabilistic edge can be taken using corner point $x = 0$. Hence maximum probability of reaching the target location D, with $k = 1$, is 1. □

Analogously to the case of cdPTA strategies, we consider strategies of clock-dependent region graphs that alternate between transitions from $\overrightarrow{\Gamma_k}$ (time elapse transitions) and transitions from $\widehat{\Gamma_k}$ (probabilistic edge transitions). Formally, a *region graph strategy* σ is a strategy of $\mathcal{A}_k$ such that, for a finite run $r \in \mathit{FinRuns}^{\mathcal{A}_k}$ that has $(l, R) \xrightarrow{a,\nu} (l', R')$ as its final transition, either $((l, R), a, \nu) \in \overrightarrow{\Gamma_k}$ and $\mathsf{support}(\sigma(r)) \in \widehat{\Gamma_k}$, or $((l, R), a, \nu) \in \widehat{\Gamma_k}$ and $\mathsf{support}(\sigma(r)) \in \overrightarrow{\Gamma_k}$. We write $\mathbf{\Pi}_k$ for the set of region graph strategies of $\mathcal{A}_k$.

Let $F \subseteq L$ be the set of target locations, which we assume to be fixed in the following. Recall that $S_F = \{(l, v) \in L \times \mathbb{R}_{\geq 0}^{\mathcal{X}} : l \in F\}$ and let $\mathsf{Regs}_k^F = \{(l, R) \in \mathsf{S}_k : l \in F\}$. The following result specifies that the maximum (minimum) probability for reaching target locations from the initial state of a cdPTA is bounded from above (from below, respectively) by the corresponding maximum (minimum, respectively) probability in the clock-dependent region graph with granularity k. Similarly, the maximum (minimum) probability computed in the region graph of granularity k is an upper (lower, respectively) bound on the maximum (minimum, respectively) probability computed in the

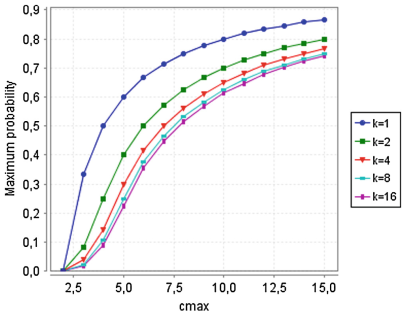

Fig. 6. Maximum probability of reaching location ✓ in the cdPTA of Fig. 1.

region graph of granularity $2k$ (we note that this result can be adapted to hold for granularity ck rather than $2k$, for any $c \in \mathbb{N} \setminus \{0, 1\}$). The proof of the proposition can be found in [22].

Proposition 1

1. $\mathbb{P}^{\max}_{[\![\mathcal{P}]\!],\mathbf{\Sigma}}(S_F) \leq \mathbb{P}^{\max}_{\mathcal{A}_k,\mathbf{\Pi}_k}(\mathsf{Regs}^F_k)$, $\mathbb{P}^{\min}_{[\![\mathcal{P}]\!],\mathbf{\Sigma}}(S_F) \geq \mathbb{P}^{\min}_{\mathcal{A}_k,\mathbf{\Pi}_k}(\mathsf{Regs}^F_k)$.

2. $\mathbb{P}^{\max}_{\mathcal{A}_{2k},\mathbf{\Pi}_{2k}}(\mathsf{Regs}^F_{2k}) \leq \mathbb{P}^{\max}_{\mathcal{A}_k,\mathbf{\Pi}_k}(\mathsf{Regs}^F_k)$, $\mathbb{P}^{\min}_{\mathcal{A}_{2k},\mathbf{\Pi}_{2k}}(\mathsf{Regs}^F_{2k}) \geq \mathbb{P}^{\min}_{\mathcal{A}_k,\mathbf{\Pi}_k}(\mathsf{Regs}^F_k)$.

Example 2 (continued). We give the intuition underlying Proposition 1 using Example 2 (Fig. 2), considering the maximum probability of reaching the target location D. When $k = 1$, as described above, the maximum probability of reaching D is 1. Instead, for $k = 2$, the maximum probability of reaching location D corresponds to taking the probabilistic edge from location A for the corner point $x = \frac{1}{2}$ corresponding to the 2-region $0 < x < \frac{1}{2}$ and the probabilistic edges from locations B and C for corner point $x = 0$, again for the 2-region $0 < x < \frac{1}{2}$ i.e., the probability is $\frac{1}{2}$. With granularity $k = 4$, the maximum probability of reaching location D is 0.328125, obtained by taking the probabilistic edge from A for the corner point $x = \frac{1}{2}$, and the probabilistic edges from B and C for corner point $x = \frac{1}{4}$, where the 4-region used in all cases is $\frac{1}{4} < x < \frac{1}{2}$. □

Example 1 (continued). In Fig. 6 we plot the values of the maximum probability of reaching location ✓ in the example of Fig. 1 for various values of $c_{\max}$ and k, obtained by encoding the clock-dependent region graph as a finite-state PTS and using Prism [15]. For this example, the difference between the probabilities obtained from low values of k is substantial. We note that the number of states of the largest instance that we considered here (for $k = 16$ and $c_{\max} = 15$) was 140174. □

5 Conclusion

In this paper we presented cdPTAs, an extension of PTAs in which probabilities can depend on the values of clocks. We have shown that a basic probabilistic model checking problem, maximal reachability, is undecidable for cdPTAs with at least three clocks. One direction of future research could be attempting to improve these results by considering cdPTAs with one or two clocks, or identifying other kinds of subclass of cdPTAs for which for which probabilistic reachability is decidable: for example, we conjecture decidability can be obtained for cdPTAs in which all clock variables are reset after utilising a probabilistic edge that depends non-trivially on clock values. Furthermore, we conjecture that qualitative reachability problems (whether there exists a strategy such that the target locations are reached with probability strictly greater than 0, or equal to 1) are decidable (and in exponential time) for cdPTAs for which the piecewise linear functions are bounded away from 0 by a region graph construction. The case of piecewise linear functions that can approach arbitrarily closely to 0 requires more care (because non-forgetful cycles, in the terminology of [5], can lead to convergence of a probability used along a cdPTA path to 0). We also presented a conservative overapproximation method for cdPTAs. At present this method gives no guarantees on the distance of the obtained bounds to the actual optimal probability: future work could address this issue, by extending the region graph construction from a PTS to a stochastic game (to provide upper and lower bounds on the maximum/minimum probability in the manner of [13]), or by considering approximate relations (by generalising the results of [6,9] from Markov chains to PTSs).

Acknowledgments. The inspiration for cdPTA arose from a discussion with Patricia Bouyer on the corner-point abstraction. Thanks also to Holger Hermanns, who expressed interest in a cdPTA-like formalism in a talk at Dagstuhl Seminar 14441.

References

1. Abate, A., Katoen, J., Lygeros, J., Prandini, M.: Approximate model checking of stochastic hybrid systems. Eur. J. Control **16**(6), 624–641 (2010)
2. Akshay, S., Bouyer, P., Krishna, S.N., Manasa, L., Trivedi, A.: Stochastic timed games revisited. In: Proceedings of the 41st International Symposium on Mathematical Foundations of Computer Science (MFCS 2016). LIPIcs, vol. 58, pp. 8:1–8:14. Leibniz-Zentrum für Informatik (2016)
3. Akshay, S., Bouyer, P., Krishna, S.N., Manasa, L., Trivedi, A.: Stochastic timed games revisited. CoRR, abs/1607.05671 (2016)
4. Alur, R., Dill, D.L.: A theory of timed automata. Theoret. Comput. Sci. **126**(2), 183–235 (1994)
5. Basset, N., Asarin, E.: Thin and thick timed regular languages. In: Fahrenberg, U., Tripakis, S. (eds.) FORMATS 2011. LNCS, vol. 6919, pp. 113–128. Springer, Heidelberg (2011). doi:10.1007/978-3-642-24310-3_9
6. Bian, G., Abate, A.: On the relationship between bisimulation and trace equivalence in an approximate probabilistic context. In: Esparza, J., Murawski, A.S. (eds.) FoSSaCS 2017. LNCS, vol. 10203, pp. 321–337. Springer, Heidelberg (2017). doi:10.1007/978-3-662-54458-7_19
7. Bouyer, P.: On the optimal reachability problem in weighted timed automata and games. In: Proceedings of the 7th Workshop on Non-Classical Models of Automata and Applications (NCMA 2015). books@ocg.at, vol. 318, pp. 11–36. Austrian Computer Society (2015)

8. Bouyer, P., Brinksma, E., Larsen, K.G.: Optimal infinite scheduling for multi-priced timed automata. Formal Methods Syst. Des. **32**(1), 2–23 (2008)
9. D'Innocenzo, A., Abate, A., Katoen, J.: Robust PCTL model checking. In: Proceedings of the 15th ACM International Conference on Hybrid Systems: Computation and Control (HSCC 2012), pp. 275–286. ACM (2012)
10. Gregersen, H., Jensen, H.E.: Formal design of reliable real time systems. Master's thesis, Department of Mathematics and Computer Science, Aalborg University (1995)
11. Hahn, E.M.: Model checking stochastic hybrid systems. Ph.D. thesis, Universität des Saarlandes (2013)
12. Jurdziński, M., Laroussinie, F., Sproston, J.: Model checking probabilistic timed automata with one or two clocks. Log. Methods Comput. Sci. **4**(3), 1–28 (2008)
13. Kattenbelt, M., Kwiatkowska, M., Norman, G., Parker, D.: A game-based abstraction-refinement framework for Markov decision processes. Formal Methods Syst. Des. **36**(3), 246–280 (2010)
14. Kemeny, J.G., Snell, J.L., Knapp, A.W.: Denumerable Markov Chains. Graduate Texts in Mathematics, 2nd edn. Springer, New York (1976)
15. Kwiatkowska, M., Norman, G., Parker, D.: PRISM 4.0: verification of probabilistic real-time systems. In: Gopalakrishnan, G., Qadeer, S. (eds.) CAV 2011. LNCS, vol. 6806, pp. 585–591. Springer, Heidelberg (2011). doi:10.1007/978-3-642-22110-1_47
16. Kwiatkowska, M., Norman, G., Parker, D., Sproston, J.: Performance analysis of probabilistic timed automata using digital clocks. Formal Methods Syst. Des. **29**, 33–78 (2006)
17. Kwiatkowska, M., Norman, G., Segala, R., Sproston, J.: Automatic verification of real-time systems with discrete probability distributions. Theoret. Comput. Sci. **286**, 101–150 (2002)
18. Minsky, M.: Computation: Finite and Infinite Machines. Prentice Hall International, Upper Saddle River (1967)
19. Norman, G., Parker, D., Sproston, J.: Model checking for probabilistic timed automata. Formal Methods Syst. Des. **43**(2), 164–190 (2013)
20. Puterman, M.L.: Markov Decision Processes. Wiley, Hoboken (1994)
21. Segala, R.: Modeling and verification of randomized distributed real-time systems. Ph.D. thesis, Massachusetts Institute of Technology (1995)
22. Sproston, J.: Probabilistic timed automata with clock-dependent probabilities. CoRR (2017)