



HAL
open science

Probabilistic Timed Automata with One Clock and Initialised Clock-Dependent Probabilities

Jeremy Sproston

► **To cite this version:**

Jeremy Sproston. Probabilistic Timed Automata with One Clock and Initialised Clock-Dependent Probabilities. 40th International Conference on Formal Techniques for Distributed Objects, Components, and Systems (FORTE), Jun 2020, Valletta, Malta. pp.150-168, 10.1007/978-3-030-50086-3_9. hal-03283239

HAL Id: hal-03283239

<https://inria.hal.science/hal-03283239>

Submitted on 9 Jul 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Probabilistic Timed Automata with One Clock and Initialised Clock-Dependent Probabilities

Jeremy Sproston

Dipartimento di Informatica, University of Turin, Italy

Abstract. Clock-dependent probabilistic timed automata extend classical timed automata with discrete probabilistic choice, where the probabilities are allowed to depend on the exact values of the clocks. Previous work has shown that the quantitative reachability problem for clock-dependent probabilistic timed automata with at least three clocks is undecidable. In this paper, we consider the subclass of clock-dependent probabilistic timed automata that have *one clock*, that have clock dependencies described by affine functions, and that satisfy an *initialisation* condition requiring that, at some point between taking edges with non-trivial clock dependencies, the clock must have an integer value. We present an approach for solving in polynomial time quantitative and qualitative reachability problems of such one-clock initialised clock-dependent probabilistic timed automata. Our results are obtained by a transformation to interval Markov decision processes.

1 Introduction

The diffusion of complex systems with timing requirements that operate in unpredictable environments has led to interest in formal modelling and verification techniques for timed and probabilistic systems. One such formal verification technique is *model checking* [12, 3], in which a system model is verified automatically against formally-specified properties. A well-established modelling formalism for timed systems is *timed automata* [2]. A timed automaton consists of a finite graph equipped with a set of real-valued variables called *clocks*, which increase at the same rate as real-time and which can be used to constrain the relative time of events. To model formally probabilistic systems, frameworks such as Markov chains or Markov decision processes are used typically. Model-checking algorithms for these formalisms have been presented in the literature: for overviews of these techniques see, for example, [7] for timed automata, and [3, 14] for Markov chains and Markov decision processes. Furthermore, timed automata and Markov decision processes have been combined to obtain the formalism of *probabilistic timed automata* [16, 25, 28], which can be viewed as timed automata with probabilities associated with their edges (or, equivalently, as Markov decision processes equipped with clocks and their associated constraints).

For the modelling of certain systems, it may be advantageous to model the fact that the probability of some events, in particular those concerning the environment in which the system is operating, vary as time passes. For example,




Fig. 1. An example of one-clock clock-dependent probabilistic automaton \mathcal{P} .

in automotive and aeronautic contexts, the probability of certain reactions of human operators may depend on factors such as fatigue, which can increase over time (see, for example, [13]); an increase in the amount of time that an unmanned aerial vehicle spends performing a search and rescue operation in a hazardous zone may increase the probability that the vehicle incurs damage from the environment; an increase in the time elapsed before a metro train arrives at a station can result in an increase in the number of passengers on the station's platform, which can in turn increase the probability of the doors failing to shut at the station, due to overcrowding of the train (see [4]). A natural way of representing such a dependency of probability of events on time is using a continuous function: for example, for the case in which a task can be completed between 1 and 3 time units in the future, we could represent the successful completion of the task by probability $\frac{x+1}{4}$, where the clock variable x (measuring the amount of time elapsed) ranges over the interval $[1, 3]$. The standard probabilistic timed automaton formalism cannot express such a continuous relationship between probabilities and time, being limited to step functions (where the intervals along which the function is constant have rational-numbered endpoints). This limitation led to the development of an extension of probabilistic timed automata called clock-dependent probabilistic timed automata [32], in which the probabilities of crossing edges can depend on clock values according to piecewise constant functions. Figure 1 gives an example of such a clock-dependent probabilistic timed automaton, using the standard conventions for the graphical representation of (probabilistic) timed automata (the model has one clock denoted by x , and black boxes denote probabilistic choices over outgoing edges). In location W , the system is *working* on a task, which is completed after between 1 and 3 units of time. When the task is completed, it is either *successful* (edge to location S), *fails* (edge to location F) or leads to system *termination* (edge to location T). For the case in which the task completion fails, between 4 and 5 time units after work on the task started the system may either try again to work on the task (edge to location W , resetting x to 0) or to terminate (edge to location T). The edges corresponding to probabilistic choices are labelled with

expressions over the clock x , which describe how the probability of those edges changes in accordance with changes in the value of x . For example, the longer the time spent in location W, the higher the value of x when location W is left, and the higher the probability of making a transition to location S, which corresponds to the successful completion of the task.

Previous work on clock-dependent probabilistic timed automata showed that a basic *quantitative* reachability problem (regarding whether there is a scheduler of nondeterministic choice such that the probability of reaching a set of target locations exceeds some probability threshold) is undecidable, but that an approach based on the region graph can be employed to approximate optimal reachability probabilities [32]. The undecidability result relied on the presence of at least three clocks: in this paper, following similar precedents in the context of (non-probabilistic and probabilistic) variants of timed automata (for example, [26, 8, 23, 5, 6, 1]), we restrict our attention to clock-dependent probabilistic timed automata with a *single* clock variable. As in [32], we consider the case in which the dependencies of transition probabilities on the value of the clock are described by affine functions. Furthermore, we assume that, between any two edges with a non-constant dependence on the clock, the clock must have a natural-numbered value, either through being reset to 0 or by increasing as time passes. We call this condition *initialisation*, following the precedents of [1] and [21], in which similar conditions are used to obtain decidability results for stochastic timed systems with one clock, and hybrid automata, respectively; intuitively, the value of the clock is “reinitialised” (either explicitly, through a reset to 0, or implicitly, through the passage of time) to a known, natural value between non-constant dependencies of probability on the value of the clock. Note that the clock-dependent probabilistic timed automaton of Figure 1 satisfies this assumption (although clock x is not reset on the edge to location F, it must take values 3 and 4 before location F can be left). We show that, for such clock-dependent probabilistic timed automata, quantitative reachability problems can be solved in polynomial time. Similarly, we can also solve in polynomial time *qualitative* reachability problems, which ask whether there exists a scheduler of nondeterminism such that a set of target locations can be reached with probability 1 (or 0), or whether all schedulers of nondeterminism result in the target locations being reached with probability 1 (or 0).

These results rely on the construction of an *interval Markov decision process* from the one-clock clock-dependent probabilistic timed automaton. Interval Markov decision processes have been well-studied in the verification context (for example, in [30, 20, 17]), and also in other contexts, such as planning [15] and control [27, 36]. They comprise a finite state space where transitions between states are achieved in the following manner: for each state, there is a nondeterministic choice between a set of actions, where each action is associated with a decoration of the set of edges from the state with *intervals* in $[0, 1]$; then a nondeterministic choice as to the exact probabilities associated with each outgoing edge is chosen from the intervals associated with the action chosen in the first step; finally, a probabilistic choice is made over the edges according to the probabilities chosen

in the second step, thus determining the next state. In contrast to the standard formulation of interval Markov decision processes, we allow edges corresponding to probabilistic choices to be labelled not only with closed intervals, but also with open and half-open intervals. While (half-)open intervals have been considered previously in the context of interval Markov chains in [9, 33], we are unaware of any work considering them in the context of interval Markov decision processes. The presence of open intervals is vital to obtain a precise representation of the one-clock clock-dependent probabilistic timed automaton.

We proceed by giving some preliminary concepts in Section 2: this includes a reduction from interval Markov decision processes to interval Markov chains [22, 24, 31] with the standard Markov decision process-based semantics, which may be of independent interest. The reduction takes open and half-open intervals into account; while [9] has shown that open interval Markov chains can be reduced to closed Markov chains for the purposes of quantitative properties, [33] shows that the open/closed distinction is critical for the evaluation of qualitative properties. In Section 3, we present the definition of one-clock clock-dependent probabilistic timed automata, and present the transformation to interval Markov decision processes in Section 4. Proofs of the results can be found in [34].

2 Interval Markov Decision Processes

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 $\sum_{q \in Q} \mu(q) = 1$. Let $\text{Dist}(Q)$ be the set of distributions over Q . For a (possibly uncountable) set Q and a function $\mu : Q \rightarrow [0, 1]$, we define $\text{support}(\mu) = \{q \in Q \mid \mu(q) > 0\}$. Then, for an uncountable set Q , we define $\text{Dist}(Q)$ to be the set of functions $\mu : Q \rightarrow [0, 1]$ such that $\text{support}(\mu)$ is a countable set and μ restricted to $\text{support}(\mu)$ is a distribution. Given a binary function $f : Q \times Q \rightarrow [0, 1]$ and element $q \in Q$, we denote by $f(q, \cdot) : Q \rightarrow [0, 1]$ the unary function such that $f(q, \cdot)(q') = f(q, q')$ for each $q' \in Q$.

A *Markov chain* (MC) \mathcal{C} is a pair (S, \mathbf{P}) where S is a set of *states* and $\mathbf{P} : S \times S \rightarrow [0, 1]$ is a *transition probability function*, such that $\mathbf{P}(s, \cdot) \in \text{Dist}(S)$ for each state $s \in S$. A *path* of MC \mathcal{C} is a sequence $s_0 s_1 \dots$ of states such that $\mathbf{P}(s_i, s_{i+1}) > 0$ for all $i \geq 0$. Given a path $\mathbf{r} = s_0 s_1 \dots$ and $i \geq 0$, we let $\mathbf{r}(i) = s_i$ be the $(i + 1)$ -th state along \mathbf{r} . The set of paths of \mathcal{C} starting in state $s \in S$ is denoted by $\text{Paths}^{\mathcal{C}}(s)$. In the standard manner (see, for example, [3, 14]), given a state $s \in S$, we can define a probability measure $\text{Pr}_s^{\mathcal{C}}$ over $\text{Paths}^{\mathcal{C}}(s)$.

A *Markov decision process* (MDP) $\mathcal{M} = (S, A, \Delta)$ comprises a set S of *states*, a set A of *actions*, and a *probabilistic transition function* $\Delta : S \times A \rightarrow \text{Dist}(S) \cup \{\perp\}$. The symbol \perp is used to represent the unavailability of an action in a state, i.e., $\Delta(s, a) = \perp$ signifies that action $a \in A$ is not available in state $s \in S$. For each state $s \in S$, let $A(s) = \{a \in A \mid \Delta(s, a) \neq \perp\}$, and assume that $A(s) \neq \emptyset$, i.e., there is at least one available action in each state. Transitions from state to state of an MDP are performed in two steps: if the current state is s , the first

step concerns a nondeterministic selection of an action $a \in A(s)$; the second step comprises a probabilistic choice, made according to the distribution $\Delta(s, a)$, as to which state to make the transition (that is, a transition to a state $s' \in S$ is made with probability $\Delta(s, a)(s')$). In general, the sets of states and actions can be uncountable. We say that an MDP is *finite* if S and A are finite sets.

A (n infinite) path of an MDP \mathcal{M} is a sequence $s_0 a_0 s_1 a_1 \cdots$ such that $a_i \in A(s_i)$ and $\Delta(s_i, \mu_i)(s_{i+1}) > 0$ for all $i \geq 0$. Given an infinite path $\mathbf{r} = s_0 a_0 s_1 a_1 \cdots$ and $i \geq 0$, we let $\mathbf{r}(i) = s_i$ be the $(i+1)$ -th state along \mathbf{r} . Let $Paths^{\mathcal{M}}$ be the set of infinite paths of \mathcal{M} . A finite path is a sequence $r = s_0 a_0 s_1 a_1 \cdots a_{n-1} s_n$ such that $a_i \in A(s_i)$ and $\Delta(s_i, \mu_i)(s_{i+1}) > 0$ for all $0 \leq i < n$. Let $last(r) = s_n$ denote the final state of r . We use ras to denote the finite path $s_0 a_0 s_1 a_1 \cdots a_{n-1} s_n a s$. Let $Paths_*^{\mathcal{M}}$ be the set of finite paths of the MDP \mathcal{M} . Let $Paths^{\mathcal{M}}(s)$ and $Paths_*^{\mathcal{M}}(s)$ be the sets of infinite paths and finite paths, respectively, of \mathcal{M} starting in state $s \in S$.

A *scheduler* is a function $\sigma : Paths_*^{\mathcal{M}} \rightarrow \bigcup_{s \in S} \text{Dist}(A(s))$ such that $\sigma(r) \in \text{Dist}(A(last(r)))$ for all $r \in Paths_*^{\mathcal{M}}$.¹ Let $\Sigma^{\mathcal{M}}$ be the set of schedulers of the MDP \mathcal{M} . We say that infinite path $\mathbf{r} = s_0 a_0 s_1 a_1 \cdots$ is *generated* by σ if $\sigma(s_0 a_0 s_1 a_1 \cdots a_{i-1} s_i)(a_i) > 0$ for all $i \in \mathbb{N}$. Let $Paths^{\sigma}$ be the set of paths generated by σ . The set $Paths_*^{\sigma}$ of finite paths generated by σ is defined similarly. Let $Paths^{\sigma}(s) = Paths^{\sigma} \cap Paths^{\mathcal{M}}(s)$ and $Paths_*^{\sigma}(s) = Paths_*^{\sigma} \cap Paths_*^{\mathcal{M}}(s)$. Given a scheduler $\sigma \in \Sigma^{\mathcal{M}}$, we can define a countably infinite-state MC \mathcal{C}^{σ} that corresponds to the behaviour of the scheduler σ : we let $\mathcal{C}^{\sigma} = (Paths_*^{\sigma}, \mathbf{P})$, where, for $r, r' \in Paths_*^{\sigma}$, we have $\mathbf{P}(r, r') = \sigma(r)(a) \cdot \Delta(last(r), a)(s)$ if $r' = ras$ and $a \in A(last(r))$, and $\mathbf{P}(r, r') = 0$ otherwise. For $r = s_0 a_0 s_1 a_1 \cdots a_{n-1} s_n$, we denote the $(i+1)$ -th prefix of r by r_i , i.e., $r_i = s_0 a_0 s_1 a_1 \cdots a_{i-1} s_i$, for $i \leq n$. Then, given $s \in S$ and $r \in Paths_*^{\sigma}(s)$, we let $\text{Pr}_{*,s}^{\sigma}(r) = \mathbf{P}(r_0, r_1) \cdots \mathbf{P}(r_{n-1}, r_n)$. Let $\text{Cyl}(r) \subseteq Paths^{\mathcal{M}}(s)$ be the set of infinite paths starting in s that have the finite path r as a prefix. Then we let Pr_s^{σ} be the unique probability measure over $Paths^{\sigma}(s)$ such that $\text{Pr}_s^{\sigma}(\text{Cyl}(r)) = \text{Pr}_{*,s}^{\sigma}(r)$ (for more details, see [3, 14]).

Given a set $T \subseteq S$, we define $\diamond T = \{\mathbf{r} \in Paths^{\mathcal{M}} \mid \exists i \in \mathbb{N}. \mathbf{r}(i) \in T\}$ as the set of infinite runs of \mathcal{M} such that some state of T is visited along the run. Let $s \in S$. We define the *maximum probability of reaching T from s* as $\mathbb{P}_{\mathcal{M},s}^{\max}(\diamond T) = \sup_{\sigma \in \Sigma^{\mathcal{M}}} \text{Pr}_s^{\sigma}(\diamond T)$. Similarly, the *minimum probability of reaching T from s* is defined as $\mathbb{P}_{\mathcal{M},s}^{\min}(\diamond T) = \inf_{\sigma \in \Sigma^{\mathcal{M}}} \text{Pr}_s^{\sigma}(\diamond T)$. The *maximal reachability problem* for \mathcal{M} , $T \subseteq S$, $s \in S$, $\triangleright \in \{\geq, >\}$ and $\lambda \in [0, 1]$ is to decide whether $\mathbb{P}_{\mathcal{M},s}^{\max}(\diamond T) \triangleright \lambda$. Similarly, the *minimal reachability problem* for \mathcal{M} , $T \subseteq S$, $s \in S$, $\triangleleft \in \{\leq, <\}$ and $\lambda \in [0, 1]$ is to decide whether $\mathbb{P}_{\mathcal{M},s}^{\min}(\diamond T) \triangleleft \lambda$. The maximal and minimal reachability problems are called *quantitative* problems. We also consider the following *qualitative* problems: $(\forall 0)$ decide whether $\text{Pr}_s^{\sigma}(\diamond T) = 0$ for all $\sigma \in \Sigma^{\mathcal{M}}$; $(\exists 0)$ decide whether there exists $\sigma \in \Sigma^{\mathcal{M}}$ such that $\text{Pr}_s^{\sigma}(\diamond T) = 0$;

¹ From [18, Lemma 4.10], without loss of generality we can assume henceforth that schedulers map to distributions assigning positive probability to *finite* sets of actions, i.e., schedulers σ for which $|\text{support}(\sigma(r))|$ is finite for all $r \in Paths_*^{\mathcal{M}}$.

($\exists 1$) decide whether there exists $\sigma \in \Sigma^{\mathcal{M}}$ such that $\Pr_s^\sigma(\diamond T) = 1$; ($\forall 1$) decide whether $\Pr_s^\sigma(\diamond T) = 1$ for all $\sigma \in \Sigma^{\mathcal{M}}$.

Interval Markov Chains. We let \mathcal{I} denote the set of (open, half-open or closed) intervals that are subsets of $[0, 1]$ and that have rational-numbered endpoints. Given an interval $I \in \mathcal{I}$, we let $\text{left}(I)$ (respectively, $\text{right}(I)$) be the left (respectively, right) endpoint of I . Note that $\text{left}(I), \text{right}(I) \in [0, 1] \cap \mathbb{Q}$.

An *interval distribution* over a finite set Q is a function $\mathfrak{d} : Q \rightarrow \mathcal{I}$ such that (1) $\sum_{q \in Q} \text{left}(\mathfrak{d}(q)) \leq 1 \leq \sum_{q \in Q} \text{right}(\mathfrak{d}(q))$, (2a) $\sum_{q \in Q} \text{left}(\mathfrak{d}(q)) = 1$ implies that $\mathfrak{d}(q)$ is left-closed for all $q \in Q$, and (2b) $\sum_{q \in Q} \text{right}(\mathfrak{d}(q)) = 1$ implies that $\mathfrak{d}(q)$ is right-closed for all $q \in Q$. We define $\mathfrak{Dist}(Q)$ as the set of interval distributions over Q . An *assignment for interval distribution* \mathfrak{d} is a distribution $\alpha \in \mathfrak{Dist}(Q)$ such that $\alpha(q) \in \mathfrak{d}(q)$ for each $q \in Q$. Note that conditions (1), (2a) and (2b) in the definition of interval distributions guarantee that there exists at least one assignment for each interval distribution. Let $\mathfrak{G}(\mathfrak{d})$ be the set of assignments for \mathfrak{d} .

An (open) *interval Markov chain* (IMC) \mathfrak{C} is a pair (S, \mathfrak{P}) , where S is a finite set of *states*, and $\mathfrak{P} : S \times S \rightarrow \mathcal{I}$ is a *interval-based transition function* such that $\mathfrak{P}(s, \cdot)$ is an interval distribution for each $s \in S$ (formally, $\mathfrak{P}(s, \cdot) \in \mathfrak{Dist}(S)$). An IMC makes a transition from a state $s \in S$ in two steps: first an assignment α is chosen from the set $\mathfrak{G}(\mathfrak{P}(s, \cdot))$ of assignments for $\mathfrak{P}(s, \cdot)$, then a probabilistic choice over target states is made according to α . The semantics of an IMC corresponds to an MDP that has the same state space as the IMC, and for which each state is associated with a set of distributions, the precise transition probabilities of which are chosen from the interval distribution of the state. Formally, the semantics of an IMC $\mathfrak{C} = (S, \mathfrak{P})$ is the MDP $\llbracket \mathfrak{C} \rrbracket = (S, \mathfrak{G}(\mathfrak{P}), \Delta)$, where $\mathfrak{G}(\mathfrak{P}) = \bigcup_{s \in S} \mathfrak{G}(\mathfrak{P}(s, \cdot))$, and for which $\Delta(s, \alpha) = \alpha$ for all states $s \in S$ and assignments $\alpha \in \mathfrak{G}(\mathfrak{P}(s, \cdot))$ for $\mathfrak{P}(s, \cdot)$. In previous literature (for example, [31, 10, 11]), this semantics is called the “IMDP semantics”.

Computing $\mathbb{P}_{\llbracket \mathfrak{C} \rrbracket, s}^{\max}(\diamond T)$ and $\mathbb{P}_{\llbracket \mathfrak{C} \rrbracket, s}^{\min}(\diamond T)$ can be done for an IMC \mathfrak{C} simply by transforming the IMC by closing all of its (half-)open intervals, then employing a standard maximum/minimum reachability probability computation on the new, “closed” IMC (for example, the algorithms of [31, 11]): the correctness of this approach is shown in [9]. Algorithms for qualitative problems of IMCs (with open, half-open and closed intervals) are given in [33]. All of the aforementioned algorithms run in polynomial time in the size of the IMC, which is obtained as the sum over all states $s, s' \in S$ of the binary representation of the endpoints of $\mathfrak{P}(s, s')$, where rational numbers are encoded as the quotient of integers written in binary.

Interval Markov Decision Processes. An (open) *interval Markov decision process* (IMDP) $\mathfrak{M} = (S, \mathfrak{A}, \mathfrak{D})$ comprises a finite set S of states, a finite set \mathfrak{A} of actions, and an *interval-based transition function* $\mathfrak{D} : S \times \mathfrak{A} \rightarrow \mathfrak{Dist}(S) \cup \{\perp\}$. Let $\mathfrak{A}(s) = \{\mathfrak{a} \in \mathfrak{A} \mid \mathfrak{D}(s, \mathfrak{a}) \neq \perp\}$, and assume that $\mathfrak{A}(s) \neq \emptyset$ for each state $s \in S$. In contrast to IMCs, an IMDP makes a transition from a state $s \in S$

in *three* steps: (1) an action $\mathbf{a} \in \mathfrak{A}(s)$ is chosen, then (2) an assignment α for $\mathfrak{D}(s, \mathbf{a})$ is chosen, and finally (3) a probabilistic choice over target states to make the transition to is performed according to α . Formally, the semantics of an IMDP $\mathfrak{M} = (S, \mathfrak{A}, \mathfrak{D})$ is the MDP $\llbracket \mathfrak{M} \rrbracket = (S, A, \Delta)$ where $A(s) = \{(\mathbf{a}, \alpha) \in \mathfrak{A} \times \text{Dist}(S) \mid \mathbf{a} \in \mathfrak{A}(s) \text{ and } \alpha \in \mathfrak{G}(\mathfrak{D}(s, \mathbf{a}))\}$ for each state $s \in S$, and $\Delta(s, (\mathbf{a}, \alpha)) = \alpha$ for each state $s \in S$ and action/assignment pair $(\mathbf{a}, \alpha) \in A(s)$. Note that (as in, for example, [30, 20, 17]) we adopt a *cooperative* resolution of nondeterminism for IMDPs, in which the choice of action and assignment (steps (1) and (2) above) is combined into a *single* nondeterministic choice in the semantic MDP.

Given the cooperative nondeterminism for IMDPs, we can show that, given an IMDP, an IMC can be constructed such that the maximal and minimal reachability probabilities for the IMDP and IMC coincide, and furthermore qualitative properties agree on the IMDP and the IMC. Formally, given the IMDP $\mathfrak{M} = (S, \mathfrak{A}, \mathfrak{D})$, we construct an IMC $\mathfrak{C}[\mathfrak{M}] = (\tilde{S}, \tilde{\mathfrak{P}})$ in the following way:

- the set of states is defined as $\tilde{S} = S \cup (S \otimes \mathfrak{A})$, where $S \otimes \mathfrak{A} = \bigcup_{s \in S} \{(s, \mathbf{a}) \in S \times \mathfrak{A} \mid \mathbf{a} \in \mathfrak{A}(s)\}$;
- for $s \in S$ and $\mathbf{a} \in \mathfrak{A}(s)$, let $\tilde{\mathfrak{P}}(s, (s, \mathbf{a})) = [0, 1]$, and let $\tilde{\mathfrak{P}}((s, \mathbf{a}), \cdot) = \mathfrak{D}(s, \mathbf{a})$.

The following proposition states the correctness of this construction with respect to quantitative and qualitative problems.

Proposition 1. *Let $\mathfrak{M} = (S, \mathfrak{A}, \mathfrak{D})$ be an IMDP, and let $s \in S$, $T \subseteq S$ and $\lambda \in \{0, 1\}$. Then:*

- $\mathbb{P}_{\llbracket \mathfrak{M} \rrbracket, s}^{\max}(\diamond T) = \mathbb{P}_{\mathfrak{C}[\mathfrak{M}], s}^{\max}(\diamond T)$ and $\mathbb{P}_{\llbracket \mathfrak{M} \rrbracket, s}^{\min}(\diamond T) = \mathbb{P}_{\mathfrak{C}[\mathfrak{M}], s}^{\min}(\diamond T)$;
- there exists $\sigma \in \Sigma^{\llbracket \mathfrak{M} \rrbracket}$ such that $\text{Pr}_s^\sigma(\diamond T) = \lambda$ if and only if there exists $\sigma' \in \Sigma^{\mathfrak{C}[\mathfrak{M}]}$ such that $\text{Pr}_s^{\sigma'}(\diamond T) = \lambda$;
- $\text{Pr}_s^\sigma(\diamond T) = \lambda$ for all $\sigma \in \Sigma^{\llbracket \mathfrak{M} \rrbracket}$ if and only if $\text{Pr}_s^{\sigma'}(\diamond T) = \lambda$ for all $\sigma' \in \Sigma^{\mathfrak{C}[\mathfrak{M}]}$.

3 Clock-Dependent Probabilistic Timed Automata with One Clock

In this section, we introduce the formalism of clock-dependent probabilistic timed automata. The definition of clock-dependent probabilistic timed automata of [32] features an arbitrary number of clock variables. In contrast, we consider models with only *one* clock variable, which will be denoted x for the remainder of the paper.

A *clock valuation* is a value $v \in \mathbb{R}_{\geq 0}$, interpreted as the current value of clock x . Following the usual notational conventions for modelling formalisms based on timed automata, for clock valuation $v \in \mathbb{R}_{\geq 0}$ and $X \in \{\{x\}, \emptyset\}$, we write $v[X:=0]$ to denote the clock valuation in which clocks in X are reset to 0; in the one-clock setting, we have $v[\{x\}:=0] = 0$ and $v[\emptyset:=0] = v$. In the following, we write $2^{\{x\}}$ rather than $\{\{x\}, \emptyset\}$.

The set Ψ of *clock constraints* over x is defined as the set of conjunctions over atomic formulae of the form $x \sim c$, where $\sim \in \{<, \leq, \geq, >\}$ 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 .

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

A *one-clock clock-dependent probabilistic timed automaton* (1c-cdPTA) $\mathcal{P} = (L, \text{inv}, \text{prob})$ comprises the following components:

- a finite set L of *locations*;
- a function $\text{inv} : L \rightarrow \Psi$ associating an *invariant condition* with each location;
- a set $\text{prob} \subseteq L \times \Psi \times \text{Temp}(2^{\{x\}} \times L)$ of *probabilistic edges*.

A probabilistic edge $(l, g, \wp) \in \text{prob}$ comprises: (1) a source location l ; (2) a clock constraint g , called a *guard*; and (3) a distribution template \wp with respect to pairs of the form $(X, l') \in 2^{\{x\}} \times L$ (i.e., pairs consisting of a first element indicating whether x should be reset to 0 or not, and a second element corresponding to a target location l'). We refer to pairs $(X, l') \in 2^{\{x\}} \times L$ as *outcomes*.

The behaviour of a 1c-cdPTA takes a similar form to that of a standard (one-clock) probabilistic timed automaton [16, 25, 23]: in any location time can advance as long as the invariant condition 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 value of the clock, and the choice as to which probabilistic edge to take is made nondeterministically; for a taken probabilistic edge, the choice of whether to reset the clock and which target location to make the transition to is *probabilistic*. In comparison with one-clock probabilistic timed automata, the key novelty of 1c-cdPTAs is that the distribution used to make this probabilistic choice depends on the probabilistic edge taken *and* on the current clock valuation.

A *state* of a 1c-cdPTA is a pair comprising a location and a clock valuation satisfying the location's invariant condition, i.e., $(l, v) \in L \times \mathbb{R}_{\geq 0}$ such that $v \models \text{inv}(l)$. In any state (l, v) , a certain amount of time $t \in \mathbb{R}_{\geq 0}$ elapses, then a probabilistic edge is traversed. The choice of t requires that the invariant $\text{inv}(l)$ remains continuously satisfied while time passes. The resulting state after the elapse of time is $(l, v+t)$. A probabilistic edge $(l', g, \wp) \in \text{prob}$ can then be chosen from $(l, v+t)$ if $l = l'$ and it is *enabled*, i.e., the clock constraint g is satisfied by $v+t$. Once a probabilistic edge (l, g, \wp) is chosen, a successor location, and whether to reset the clock to 0, is chosen at random, according to the distribution $\wp[v+t]$.

We make the following assumptions on 1c-cdPTAs, in order to simplify the definition of their semantics. Firstly, we consider 1c-cdPTAs featuring invariant conditions that prevent the clock from exceeding some bound, and impose no lower bound: formally, for each location $l \in L$, we have that $\text{inv}(l)$ is a constraint $x \leq c$ for some $c \in \mathbb{N}$, or a constraint $x < c$ for some $c \in \mathbb{N} \setminus \{0\}$. Secondly, we restrict our attention to 1c-cdPTAs for which it is always possible to take a

probabilistic edge, either immediately or after letting time elapse. Formally, for each location $l \in L$, if $inv(l) = (x \leq c)$ then (viewing c as a clock valuation) $c \models g$ for some $(l, g, \wp) \in prob$; instead, if $inv(l) = (x < c)$ then $c - \varepsilon \models g$ for all $\varepsilon \in (0, 1)$ and $(l, g, \wp) \in prob$. Thirdly, we assume that all possible target states of probabilistic edges satisfy their invariants. Observe that, given the first assumption, this may not be the case only when the clock is *not* reset. Formally, for all probabilistic edges $(l, g, \wp) \in prob$, for all clock valuations $v \in \mathbb{R}_{\geq 0}$ such that $v \models g$, and for all $l' \in L$, we have that $\wp[v](\emptyset, l') > 0$ implies $v[\emptyset := 0] \models inv(l')$, i.e., $v \models inv(l')$. Note that we relax some of these assumptions when depicting 1c-cdPTAs graphically (for example, the 1c-cdPTA of Figure 1 can be made to satisfy these assumptions by adding invariant conditions and self-looping probabilistic edges to locations S and T).

The semantics of the 1c-cdPTA $\mathcal{P} = (L, inv, prob)$ is the MDP $\llbracket \mathcal{P} \rrbracket = (S, A, \Delta)$ where:

- $S = \{(l, v) \in L \times \mathbb{R}_{\geq 0} \mid v \models inv(l)\}$;
- $A = \mathbb{R}_{\geq 0} \times prob$;
- for $(l, v) \in S$, $\tilde{v} \in \mathbb{R}_{\geq 0}$ and $(l, g, \wp) \in prob$ such that (1) $\tilde{v} \geq v$, (2) $\tilde{v} \models g$ and (3) $w \models inv(l)$ for all $v \leq w \leq \tilde{v}$, then we let $\Delta((l, v), (\tilde{v}, (l, g, \wp)))(l', v')$ be the distribution such that, for $(l', v') \in S$:

$$\Delta((l, v), (\tilde{v}, (l, g, \wp)))(l', v') = \begin{cases} \wp[\tilde{v}](\{x\}, l') + \wp[\tilde{v}](\emptyset, l') & \text{if } v' = \tilde{v} = 0 \\ \wp[\tilde{v}](\emptyset, l') & \text{if } v' = \tilde{v} > 0 \\ \wp[\tilde{v}](\{x\}, l') & \text{if } v' = 0 \text{ and } \tilde{v} > 0 \\ 0 & \text{otherwise.} \end{cases}$$

Let $F \subseteq L$ be a set of locations, and let $T_F = \{(l, v) \in S \mid l \in F\}$ be the set of states of $\llbracket \mathcal{P} \rrbracket$ that have their location component in F . Then the maximum value of reaching F from state $(l, v) \in S$ corresponds to $\mathbb{P}_{\llbracket \mathcal{P} \rrbracket, (l, v)}^{\max}(\diamond T_F)$. Similarly, the minimum value of reaching F from state $(l, v) \in S$ corresponds to $\mathbb{P}_{\llbracket \mathcal{P} \rrbracket, (l, v)}^{\min}(\diamond T_F)$. As in Section 2, we can define a number of quantitative and qualitative reachability problems on 1c-cdPTA, where the initial state is set as $(l, 0)$ for a particular $l \in L$. The maximal reachability problem for \mathcal{P} , $F \subseteq L$, $l \in L$, $\triangleright \in \{\geq, >\}$ and $\lambda \in [0, 1]$ is to decide whether $\mathbb{P}_{\llbracket \mathcal{P} \rrbracket, (l, 0)}^{\max}(\diamond T_F) \triangleright \lambda$; similarly, the minimal reachability problem for \mathcal{P} , $F \subseteq L$, $l \in L$, $\triangleleft \in \{\leq, <\}$ and $\lambda \in [0, 1]$ is to decide whether $\mathbb{P}_{\llbracket \mathcal{P} \rrbracket, (l, 0)}^{\min}(\diamond T_F) \triangleleft \lambda$. Furthermore, we can define analogues of the qualitative problems featured in Section 2: ($\forall 0$) decide whether $\Pr_{(l, 0)}^{\sigma}(\diamond T_F) = 0$ for all $\sigma \in \Sigma^{\llbracket \mathcal{P} \rrbracket}$; ($\exists 0$) decide whether there exists $\sigma \in \Sigma^{\llbracket \mathcal{P} \rrbracket}$ such that $\Pr_{(l, 0)}^{\sigma}(\diamond T_F) = 0$; ($\exists 1$) decide whether there exists $\sigma \in \Sigma^{\llbracket \mathcal{P} \rrbracket}$ such that $\Pr_{(l, 0)}^{\sigma}(\diamond T_F) = 1$; ($\forall 1$) decide whether $\Pr_{(l, 0)}^{\sigma}(\diamond T_F) = 1$ for all $\sigma \in \Sigma^{\llbracket \mathcal{P} \rrbracket}$.

Affine Clock Dependencies. In this paper, we consider distribution templates that are defined in terms of sets of affine functions in the following way.

Given probabilistic edge $p = (l, g, \wp) \in prob$, let I^p be the set of clock valuations in which p is enabled, i.e., $I^p = \{v \in \mathbb{R}_{\geq 0} \mid v \models g \wedge inv(l)\}$. Note that

$I^p \subseteq \mathbb{R}_{\geq 0}$ corresponds to an interval with natural-numbered endpoints. Let $\overline{I^p}$ be the closure of I^p . We say that p is *affine* if, for each $e \in 2^{\{x\}} \times L$, there exists a pair $(c_e^p, d_e^p) \in \mathbb{Q}^2$ of rational constants, such that $\wp[v](e) = c_e^p + d_e^p \cdot v$ for all $v \in \overline{I^p}$. Note that, by the definition of distribution templates, for all $v \in \overline{I^p}$, we have $c_e^p + d_e^p \cdot v \geq 0$ for each $e \in 2^{\{x\}} \times L$, and $\sum_{e \in 2^{\{x\}} \times L} (c_e^p + d_e^p \cdot v) = 1$. A 1c-cdPTA is affine if all of its probabilistic edges are affine. Henceforth we assume that the 1c-cdPTAs we consider are affine. An affine probabilistic edge p is *constant* if, for each $e \in 2^{\{x\}} \times L$, we have $d_e^p = 0$, i.e., $\wp[v](e) = c_e^p$ for some $c_e^p \in \mathbb{Q}$, for all $v \in \overline{I^p}$. Note that, for a probabilistic edge $p \in \text{prob}$, outcome $e \in 2^{\{x\}} \times L$ and open interval $I \subseteq I^p$, if $d_e^p \neq 0$, then $\wp[v](e) > 0$ for all $v \in I$ (because the existence of $v_{=0} \in I$ such that $\wp[v_{=0}](e) = 0$, together with $d_e^p \neq 0$, would mean that there exists $v' \in I$ such that $\wp[v'](e) < 0$, which contradicts the definition of distribution templates).

Initialisation. In this paper, we also introduce a specific requirement for 1c-cdPTAs that allows us to analyse faithfully 1c-cdPTA using IMDPs in Section 4. A *symbolic path fragment* is a sequence $(l_0, g_0, \wp_0)(X_0, l_1)(l_1, g_1, \wp_1)(X_1, l_2) \cdots (l_n, g_n, \wp_n) \in (\text{prob} \times (2^{\{x\}} \times L))^+ \times \text{prob}$ of probabilistic edges and outcomes such that $\wp_i[v](X_i, l_{i+1}) > 0$ for all $v \in I^{(l_i, g_i, \wp_i)}$ for all $i < n$. In this paper, we consider 1c-cdPTAs for which each symbolic path fragment that begins and ends with a non-constant probabilistic edge requires that the clock takes a natural numbered value, either from being reset or from passing through guards that have at most one (natural numbered) value in common. Formally, a 1c-cdPTA is *initialised* if, for any symbolic path fragment $(l_0, g_0, \wp_0)(X_0, l_1)(l_1, g_1, \wp_1)(X_1, l_2) \cdots (l_n, g_n, \wp_n)$ such that (l_0, g_0, \wp_0) and (l_n, g_n, \wp_n) are non-constant, either (1) $X_i = \{x\}$ or (2) $I^{(l_i, g_i, \wp_i)} \cap I^{(l_{i+1}, g_{i+1}, \wp_{i+1})}$ is empty or contains a single valuation, for some $0 < i < n$. We henceforth assume that all 1c-cdPTAs considered in this paper are initialised.

4 Translation from 1c-cdPTAs to IMDPs

In this section, we show that we can solve quantitative and qualitative problems of (affine and initialised) 1c-cdPTAs. In contrast to the approach for quantitative problems of multiple-clock cdPTAs presented in [32], which involves the construction of an *approximate* MDP, we represent the 1c-cdPTA *precisely* using an IMDP, by adapting the standard region-graph construction for one-clock (probabilistic) timed automata of [26, 23].

Let $\mathcal{P} = (L, \text{inv}, \text{prob})$ be a 1c-cdPTA. Let $\text{Cst}(\mathcal{P})$ be the set of constants that are used in the guards of probabilistic edges and invariants of \mathcal{P} , and let $\mathbb{B} = \text{Cst}(\mathcal{P}) \cup \{0\}$. We write $\mathbb{B} = \{b_0, b_1, \dots, b_k\}$, where $0 = b_0 < b_1 < \dots < b_k$. The set \mathbb{B} defines the set $\mathcal{I}_{\mathbb{B}} = \{[b_0, b_0], (b_0, b_1), [b_1, b_1], \dots, [b_k, b_k]\}$. We define a total order on $\mathcal{I}_{\mathbb{B}}$ in the following way: $[b_0, b_0] < (b_0, b_1) < [b_1, b_1] < \dots < [b_k, b_k]$. Given an open interval $B = (b, b') \in \mathcal{I}_{\mathbb{B}}$, its closure is written as \overline{B} , i.e., $\overline{B} = [b, b']$. Furthermore, let $\text{lb}(B) = b$ and $\text{rb}(B) = b'$ refer to the left- and right-endpoints of B . For a closed interval $[b, b] \in \mathcal{I}_{\mathbb{B}}$, we let $\text{lb}(B) = \text{rb}(B) = b$.

Let ψ be a guard of a probabilistic edge or an invariant of \mathcal{P} . By definition, we have that, for each $B \in \mathcal{I}_{\mathbb{B}}$, either $B \subseteq \{v \in \mathbb{R}_{\geq 0} \mid v \models \psi\}$ or $B \cap \{v \in \mathbb{R}_{\geq 0} \mid v \models \psi\} = \emptyset$. We write $B \models \psi$ in the case of $B \subseteq \{v \in \mathbb{R}_{\geq 0} \mid v \models \psi\}$ (to represent the fact that all valuations of B satisfy ψ).

Example 1. Consider the 1c-cdPTA of Figure 1. We have $\mathbb{B} = \{0, 1, 3, 4, 5\}$ and $\mathcal{I}_{\mathbb{B}} = \{[0, 0], (0, 1), [1, 1], (1, 3), [3, 3], (3, 4), [4, 4], (4, 5), [5, 5]\}$. Consider the clock constraint $x < 3$: we have $B \models (x < 3)$ for all $B \in \{[0, 0], (0, 1), [1, 1], (1, 3)\}$. Similarly, for the clock constraint $4 < x < 5$, we have $(4, 5) \models (4 < x < 5)$.

\mathbb{B} -minimal Schedulers. The following technical lemma specifies that any scheduler of the 1c-cdPTA can be made “more deterministic” in the following way: for each interval $\tilde{B} \in \mathcal{I}_{\mathbb{B}}$ and probabilistic edge $p \in \text{prob}$, if, after executing a certain finite path, a scheduler chooses (assigns positive probability to) multiple actions $(\tilde{v}_1, p), \dots, (\tilde{v}_n, p)$ that share the same probabilistic edge p and for which $\tilde{v}_i \in \tilde{B}$ for all $1 \leq i \leq n$, then we can obtain another scheduler for which the aforementioned actions are replaced by an action (\tilde{v}, p) such that $\tilde{v} \in \tilde{B}$. Formally, we say that a scheduler $\sigma \in \Sigma^{\llbracket \mathcal{P} \rrbracket}$ of $\llbracket \mathcal{P} \rrbracket$ is \mathbb{B} -minimal if, for all finite paths $r \in \text{Paths}_{*}^{\llbracket \mathcal{P} \rrbracket}$, for all probabilistic edges $p \in \text{prob}$, and for all pairs of actions $(\tilde{v}_1, p_1), (\tilde{v}_2, p_2) \in \text{support}(\sigma(r))$, either $p_1 \neq p_2$ or v_1 and v_2 belong to distinct intervals in $\mathcal{I}_{\mathbb{B}}$, i.e., the intervals $\tilde{B}_1, \tilde{B}_2 \in \mathcal{I}_{\mathbb{B}}$ for which $\tilde{v}_1 \in \tilde{B}_1$ and $\tilde{v}_2 \in \tilde{B}_2$ are such that $\tilde{B}_1 \neq \tilde{B}_2$. Let $\Sigma_{\mathbb{B}}^{\llbracket \mathcal{P} \rrbracket}$ be the set of schedulers of $\Sigma^{\llbracket \mathcal{P} \rrbracket}$ that are \mathbb{B} -minimal.

Lemma 1. *Let $(l, v) \in S_{\mathcal{P}}$ and $F \subseteq L$. Then, for each $\sigma \in \Sigma^{\llbracket \mathcal{P} \rrbracket}$, there exists $\pi \in \Sigma_{\mathbb{B}}^{\llbracket \mathcal{P} \rrbracket}$ such that $\text{Pr}_{(l,v)}^{\sigma}(\diamond T_F) = \text{Pr}_{(l,v)}^{\pi}(\diamond T_F)$.*

The underlying idea of the proof of the lemma (which can be found in [34]) is that every finite path of π corresponds to a set of finite paths of σ , where all of these paths have the same length, visit the same locations in order, choose the same probabilistic edges in order, and visit the same intervals of clock valuations in order. Consider the choice of π after a finite path: to replicate the choices made at the end of the corresponding set of finite paths of σ , the choice of π is derived from a weighted average of the choices of σ , where the weights correspond to the probabilities of the finite paths of σ under consideration. Another key point for the construction of π is that, when a non-constant probabilistic edge is chosen, the clock valuation used by π when taking the probabilistic edge reflects the clock valuation used by σ when taking some probabilistic edge from the finite paths of σ under consideration: the clock valuation chosen by π is obtained as a weighted average of the clock valuations chosen by σ . Lemma 1 allows us to consider only \mathbb{B} -minimal schedulers in the sequel, permitting us to obtain a close correspondence between the schedulers of $\llbracket \mathcal{P} \rrbracket$ and the schedulers of the constructed IMDP, the definition of which we consider in the subsequent subsection.

Example 2. Consider the 1c-cdPTA of Figure 1. In the following, we denote the outgoing probabilistic edges from W and F as p_W and p_F , respectively. Consider

a scheduler $\sigma \in \Sigma^{\llbracket \mathcal{P} \rrbracket}$, where $\sigma(W, 0)$ (i.e., the choice of σ after the finite path comprising the single state $(W, 0)$) assigns probability $\frac{1}{2}$ to the action $(\frac{5}{4}, p_W)$ and probability $\frac{1}{2}$ to the action $(\frac{7}{4}, p_W)$ (where the two actions refer to either $\frac{5}{4}$ or $\frac{7}{4}$ time units elapsing, after which the probabilistic edge p_W is taken). Then we can construct a \mathbb{B} -minimal scheduler $\pi \in \Sigma_{\mathbb{B}}^{\llbracket \mathcal{P} \rrbracket}$ such that $\pi(W, 0)$ assigns probability 1 to the action $(\frac{3}{2}, p_W)$ (i.e., where $\frac{3}{2} = \frac{1}{2} \cdot \frac{5}{4} + \frac{1}{2} \cdot \frac{7}{4}$). Now consider finite paths $r = (W, 0)(\frac{5}{4}, p_W)(F, \frac{5}{4})$ and $r' = (W, 0)(\frac{7}{4}, p_W)(F, \frac{7}{4})$. Note that $\Pr_{*,(W,0)}^{\sigma}(r) = \frac{1}{2} \cdot \frac{11 - \frac{3 \cdot 5}{4}}{16}$ and $\Pr_{*,(W,0)}^{\sigma}(r') = \frac{1}{2} \cdot \frac{11 - \frac{3 \cdot 7}{4}}{16}$. Say that $\sigma(r)$ assigns probability 1 to $\frac{17}{4}$ and probability 1 to $\frac{19}{4}$. Then $\pi((W, 0)(\frac{3}{2}, p_W)(F, \frac{3}{2}))$ assigns probability 1 to action (\tilde{v}, p_F) , where $\tilde{v} = \Pr_{*,(W,0)}^{\sigma}(r) \cdot 1 \cdot \frac{17}{4} + \Pr_{*,(W,0)}^{\sigma}(r') \cdot 1 \cdot \frac{19}{4}$, i.e., a weighted sum of the time delays chosen by σ after r and r' , where the weights correspond to the probabilities of r and r' under σ . Repeating this reasoning for all finite paths yields a \mathbb{B} -minimal scheduler π such that the probability of reaching a set of target states from $(W, 0)$ is the same for both σ and π .

IMDP Construction. We now present the idea of the IMDP construction. The states of the IMDP fall into two categories: (1) pairs comprising a location and an interval from $\mathcal{I}_{\mathbb{B}}$, with the intuition that the state $(l, B) \in L \times \mathcal{I}_{\mathbb{B}}$ of the IMDP represents all states (l, v) of $\llbracket \mathcal{P} \rrbracket$ such that $v \in B$; (2) triples comprising an interval from $\mathcal{I}_{\mathbb{B}}$, a probabilistic edge and a bit that specifies whether the state refers to the left- or right-endpoint of the interval. A single transition of the semantics of the 1c-cdPTA, which we recall represents the elapse of time (therefore increasing the value of the clock) followed by the traversal of a probabilistic edge, is represented by a sequence of *two* transitions in the IMDP: the first IMDP transition represents the choice of (i) the probabilistic edge, (ii) the interval in $\mathcal{I}_{\mathbb{B}}$ which contains the valuation of the clock after letting time elapse and immediately before the probabilistic edge is traversed, and (iii) in the case in which the aforementioned interval is open, the position of the clock valuation within the interval; the second IMDP transition represents the probabilistic choice made from the extreme (left and right) endpoints of the aforementioned interval with the chosen probabilistic edges.

Example 3. The IMDP construction, applied to the example of Figure 1, is shown in Figure 2 (note that transitions corresponding to probability 0 are shown with a dashed line). The location W , and the value of the clock being 0, is represented by the state $(W, [0, 0])$. Recall that the outgoing probabilistic edge from W is enabled when the clock is between 1 and 3: hence the single action $((1, 3), p_W)$ is available from $(W, [0, 0])$ (representing the set of actions (\tilde{v}, p_W) of $\llbracket \mathcal{P} \rrbracket$ with $\tilde{v} \in (1, 3)$). The action $((1, 3), p_W)$ is associated with two target states, $((1, 3), p_W, \text{lb})$ and $((1, 3), p_W, \text{rb})$, each corresponding to the probability interval $(0, 1)$. The choice of probability within the interval can be done in the IMDP to represent a choice of clock valuation in $(1, 3)$: for example, the valuation $\frac{3}{2}$ would be represented by the assignment that associates probability $\frac{3}{4}$ with $((1, 3), p_W, \text{lb})$ and $\frac{1}{4}$ with $((1, 3), p_W, \text{rb})$ (i.e., assigns a weight of $\frac{3}{4}$ to the lower bound of $(1, 3)$, and a weight of $\frac{1}{4}$ to the upper bound of

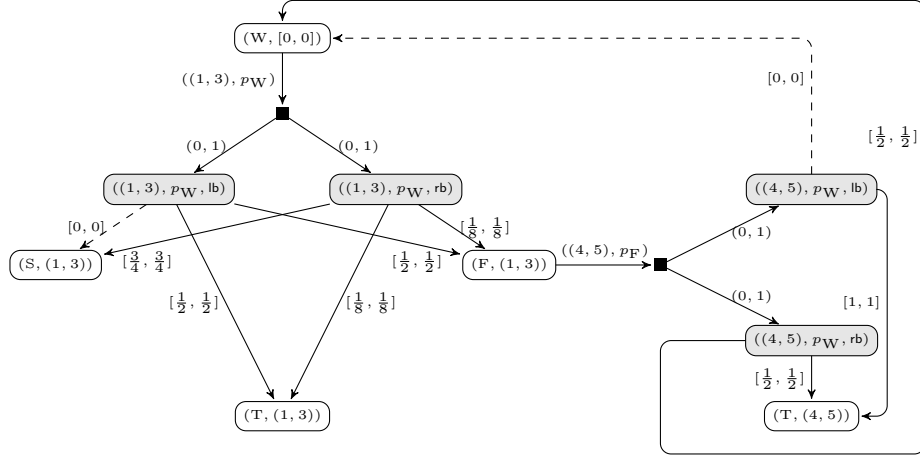


Fig. 2. Interval Markov decision process $\mathfrak{M}[\mathcal{P}]$ obtained from \mathcal{P} .

$(1, 3)$, obtaining the weighted combination $\frac{3}{4} \cdot 1 + \frac{1}{4} \cdot 3 = \frac{3}{2}$. Then, from both $((1, 3), p_W, lb)$ and $((1, 3), p_W, rb)$, there is a probabilistic choice made regarding the target IMDP state to make the subsequent transition to, i.e., the transitions from $((1, 3), p_W, lb)$ and $((1, 3), p_W, rb)$ do not involve nondeterminism, because there is only one action available, and because the resulting interval distribution assigns singleton intervals to all possible target states.² The probabilities of the transitions from $((1, 3), p_W, lb)$ and $((1, 3), p_W, rb)$ are derived from the clock dependencies associated with 1 (i.e., the left endpoint of $(1, 3)$) and 3 (i.e., the right endpoint of $(1, 3)$), respectively. Hence the multiplication of the probabilities of the two aforementioned transitions (from $(W, [0, 0])$ to either $((1, 3), p_W, lb)$ or $((1, 3), p_W, rb)$, and then to $(S, (1, 3))$, $(T, (1, 3))$ or $(F, (1, 3))$) represents exactly the probability of a single transition in the 1c-cdPTA: for example, in the 1c-cdPTA, considering again the example of the clock valuation associating $\frac{3}{2}$ with x , the probability of making a transition to location S is $\frac{3x-3}{8} = \frac{3}{16}$; in the IMDP, assigning $\frac{3}{4}$ to the transition to $((1, 3), p_W, lb)$ and $\frac{1}{4}$ to the transition to $((1, 3), p_W, rb)$, we then obtain that the probability of making a transition to $(S, (1, 3))$ from $(W, [0, 0])$ is $\frac{3}{4} \cdot 0 + \frac{1}{4} \cdot \frac{3}{4} = \frac{3}{16}$. Similar reasoning applies to the transitions available from $(F, (1, 3))$.

We now describe formally the construction of the IMDP $\mathfrak{M}[\mathcal{P}] = (S_{\mathfrak{M}[\mathcal{P}]}, \mathcal{A}_{\mathfrak{M}[\mathcal{P}]}, \mathcal{D}_{\mathfrak{M}[\mathcal{P}]})$. The set of states of $\mathfrak{M}[\mathcal{P}]$ is defined as $S_{\mathfrak{M}[\mathcal{P}]} = S_{\mathfrak{M}[\mathcal{P}]}^{\text{reg}} \cup S_{\mathfrak{M}[\mathcal{P}]}^{\text{end}}$, where $S_{\mathfrak{M}[\mathcal{P}]}^{\text{reg}} = \{(l, B) \in L \times \mathcal{I}_{\mathbb{B}} \mid B \models \text{inv}(l)\}$ and $S_{\mathfrak{M}[\mathcal{P}]}^{\text{end}} = \{(\tilde{B}, (l, g, \varphi), \text{dir}) \in \mathcal{I}_{\mathbb{B}} \times \text{prob} \times \{\text{lb}, \text{rb}\} \mid \tilde{B} \models g \wedge \text{inv}(l)\}$. In order to distinguish states of $\llbracket \mathcal{P} \rrbracket$ and states of $\mathfrak{M}[\mathcal{P}]$, we refer to elements of $S_{\mathfrak{M}[\mathcal{P}]}^{\text{reg}}$ as *regions*, and elements of $S_{\mathfrak{M}[\mathcal{P}]}^{\text{end}}$ as *endpoint indicators*. The set of actions of $\mathfrak{M}[\mathcal{P}]$ is de-

² Given that there is only one action available from states such as $((1, 3), p_W, lb)$ and $((1, 3), p_W, rb)$, we omit both the action and the usual black box from the figure.

defined as $\mathfrak{A}_{\mathfrak{M}[\mathcal{P}]} = \{(\tilde{B}, (l, g, \wp)) \in \mathcal{I}_{\mathbb{B}} \times \text{prob} \mid \tilde{B} \models g \wedge \text{inv}(l)\} \cup \{\tau\}$ (i.e., there is an action for each combination of interval from $\mathcal{I}_{\mathbb{B}}$ and probabilistic edge such that all valuations from the interval satisfy both the guard of the probabilistic edge and the invariant condition of its source location). For each region $(l, B) \in S_{\mathfrak{M}[\mathcal{P}]}^{\text{reg}}$, let $\mathfrak{A}_{\mathfrak{M}[\mathcal{P}]}(l, B) = \{(\tilde{B}, (l', g, \wp)) \in \mathfrak{A}_{\mathfrak{M}[\mathcal{P}]} \mid l = l' \text{ and } \tilde{B} \geq B\}$.³ For each $(\tilde{B}, p, \text{dir}) \in S_{\mathfrak{M}[\mathcal{P}]}^{\text{end}}$, let $\mathfrak{A}_{\mathfrak{M}[\mathcal{P}]}(\tilde{B}, p, \text{dir}) = \{\tau\}$. The transition function $\mathfrak{D}_{\mathfrak{M}[\mathcal{P}]} : S_{\mathfrak{M}[\mathcal{P}]} \times \mathfrak{A}_{\mathfrak{M}[\mathcal{P}]} \rightarrow \mathfrak{Dist}(S_{\mathfrak{M}[\mathcal{P}]}) \cup \{\perp\}$ is defined as follows⁴:

- For each $(l, B) \in S_{\mathfrak{M}[\mathcal{P}]}^{\text{reg}}$ and $(\tilde{B}, p) \in \mathfrak{A}_{\mathfrak{M}[\mathcal{P}]}(l, B)$, we let $\mathfrak{D}_{\mathfrak{M}[\mathcal{P}]}((l, B), (\tilde{B}, p))$ be the interval distribution such that $\mathfrak{D}_{\mathfrak{M}[\mathcal{P}]}((l, B), (\tilde{B}, p))(\tilde{B}, p, \text{lb}) = (0, 1)$, $\mathfrak{D}_{\mathfrak{M}[\mathcal{P}]}((l, B), (\tilde{B}, p))(\tilde{B}, p, \text{rb}) = (0, 1)$, and $\mathfrak{D}_{\mathfrak{M}[\mathcal{P}]}((l, B), (\tilde{B}, p))(s) = [0, 0]$ for all $s \in S_{\mathfrak{M}[\mathcal{P}]} \setminus \{(\tilde{B}, p, \text{lb}), (\tilde{B}, p, \text{rb})\}$.
- For each $(\tilde{B}, (l, g, \wp), \text{dir}) \in S_{\mathfrak{M}[\mathcal{P}]}^{\text{end}}$ and $(l', B') \in S_{\mathfrak{M}[\mathcal{P}]}^{\text{reg}}$, let:

$$\lambda_{(l', B')}^{(\tilde{B}, (l, g, \wp), \text{dir})} = \begin{cases} \wp[\text{dir}(\tilde{B})](\{x\}, l') + \wp[\text{dir}(\tilde{B})](\emptyset, l') & \text{if } B' = \tilde{B} = [0, 0] \\ \wp[\text{dir}(\tilde{B})](\emptyset, l') & \text{if } B' = \tilde{B} > [0, 0] \\ \wp[\text{dir}(\tilde{B})](\{x\}, l') & \text{if } B' = 0 \text{ and } \tilde{B} > 0 \\ 0 & \text{otherwise.} \end{cases}$$

Then $\mathfrak{D}_{\mathfrak{M}[\mathcal{P}]}((\tilde{B}, (l, g, \wp), \text{dir}), \tau)$ is the interval distribution such that, for all $s \in S_{\mathfrak{M}[\mathcal{P}]}$:

$$\mathfrak{D}_{\mathfrak{M}[\mathcal{P}]}((\tilde{B}, (l, g, \wp), \text{dir}), \tau)(s) = \begin{cases} [\lambda_s^{(\tilde{B}, (l, g, \wp), \text{dir})}, \lambda_s^{(\tilde{B}, (l, g, \wp), \text{dir})}] & \text{if } s \in S_{\mathfrak{M}[\mathcal{P}]}^{\text{reg}} \\ [0, 0] & \text{otherwise.} \end{cases}$$

Next, we consider the correctness of the construction of $\mathfrak{M}[\mathcal{P}]$, i.e., that $\mathfrak{M}[\mathcal{P}]$ can be used for solving quantitative and qualitative properties of the 1c-cdPTA \mathcal{P} . The proof relies on showing that transitions of the semantic MDP $\llbracket \mathcal{P} \rrbracket$ of \mathcal{P} can be mimicked by a sequence of *two* transitions of the semantic MDP $\llbracket \mathfrak{M}[\mathcal{P}] \rrbracket$ of $\mathfrak{M}[\mathcal{P}]$, and vice versa. Let $\llbracket \mathcal{P} \rrbracket = (S_{\mathcal{P}}, A_{\mathcal{P}}, \Delta_{\mathcal{P}})$ be the semantic MDP of \mathcal{P} . Given state $(l, v) \in S_{\mathcal{P}}$, we let $\text{reg}(l, v) = (l, B) \in S_{\mathfrak{M}[\mathcal{P}]}^{\text{reg}}$ be the unique region such that $v \in B$. In the following, we let $\llbracket \mathfrak{M}[\mathcal{P}] \rrbracket = (S_{\mathfrak{M}[\mathcal{P}]}, A_{\mathfrak{M}[\mathcal{P}]}, \Delta_{\mathfrak{M}[\mathcal{P}]})$ be the semantic MDP of $\mathfrak{M}[\mathcal{P}]$.

We now show that, for any scheduler of (the semantics of) the 1c-cdPTA \mathcal{P} , there exists a scheduler of (the semantics of) the IMDP $\mathfrak{M}[\mathcal{P}]$ such that the schedulers assign the same probability to reaching a certain set of locations from a given location with the value of the clock equal to 0. Let $\mathfrak{T}_F = \{(l, B) \in S_{\mathfrak{M}[\mathcal{P}]}^{\text{reg}} \mid l \in F\}$ be the set of regions with location component in F .

³ Note that $\mathfrak{A}_{\mathfrak{M}[\mathcal{P}]}(l, B) \neq \emptyset$ for each $(l, B) \in S_{\mathfrak{M}[\mathcal{P}]}^{\text{reg}}$, by the assumptions that we made on 1c-cdPTA in Section 3 (namely, that it is always possible to take a probabilistic edge, either immediately or after letting time elapse).

⁴ We recall that $\mathfrak{D}_{\mathfrak{M}[\mathcal{P}]}(s, \mathfrak{a}) = \perp$ for $s \in S_{\mathfrak{M}[\mathcal{P}]}$ and $\mathfrak{a} \in \mathfrak{A}_{\mathfrak{M}[\mathcal{P}]} \setminus \mathfrak{A}_{\mathfrak{M}[\mathcal{P}]}(s)$.

Lemma 2. *Let $l \in L$ be a location and let $F \subseteq L$ be a set of locations. Given a \mathbb{B} -minimal scheduler $\pi \in \Sigma_{\mathbb{B}}^{\llbracket \mathcal{P} \rrbracket}$, there exists scheduler $\hat{\pi} \in \Sigma^{\llbracket \mathfrak{M}[\mathcal{P}] \rrbracket}$ such that $\Pr_{(l,0)}^{\pi}(\diamond T_F) = \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(\diamond \mathfrak{F}_F)$.*

The proof of Lemma 2 (see [34]) is simplified by the fact that, by Lemma 1, it suffices to consider \mathbb{B} -minimal schedulers: for each finite path r of π , we can identify a set of finite paths of $\hat{\pi}$ of length twice that of r , that visit the same locations in order, choose the same probabilistic edges in order, and visit the same intervals in order, both regarding the clock valuations/intervals in states and in actions. In fact, finite paths of $\hat{\pi}$ that are associated with r differ *only* in terms of the lb and rb components used in endpoint indicators. Furthermore, $\hat{\pi}$ replicates exactly the choice of π made after r in terms of interval of $\mathcal{I}_{\mathbb{B}}$ chosen and probabilistic edge in all of its finite paths associated with r . Finally, $\hat{\pi}$ chooses assignments (over edges labelled with $(0,1)$) in order to represent exactly the choices of clock valuations made by π , in the manner described in Example 3 above: more precisely, the choice of action (\tilde{v}, p) by π , where \tilde{B} is the unique interval such that $\tilde{v} \in \tilde{B}$, is mimicked by $\hat{\pi}$ choosing the action $((\tilde{B}, p), \alpha)$ for which $\alpha(\tilde{B}, p, \text{lb}) = \frac{\text{rb}(\tilde{B}) - \tilde{v}}{\text{rb}(\tilde{B}) - \text{lb}(\tilde{B})}$, and $\alpha(\tilde{B}, p, \text{rb}) = 1 - \alpha(\tilde{B}, p, \text{lb}) = \frac{\tilde{v} - \text{lb}(\tilde{B})}{\text{rb}(\tilde{B}) - \text{lb}(\tilde{B})}$.

Example 4. Consider the 1c-cdPTA of Figure 1. Let $\pi \in \Sigma_{\mathbb{B}}^{\llbracket \mathcal{P} \rrbracket}$ be a scheduler such that $\pi(W, 0)$ assigns probability 1 to the action $(\frac{3}{2}, p_W)$. Then $\hat{\pi} \in \Sigma^{\llbracket \mathfrak{M}[\mathcal{P}] \rrbracket}$ is constructed such that $\hat{\pi}(W, [0, 0])$ assigns probability 1 to $((1, 3), p_W, \alpha)$, where $\alpha((1, 3), p_W, \text{lb}) = \frac{3}{4}$ and $\alpha((1, 3), p_W, \text{rb}) = \frac{1}{4}$ (observe that $\alpha((1, 3), p_W, \text{lb}) = \frac{3 - \frac{3}{2}}{2}$). Furthermore, $\hat{\pi}((W, [0, 0])((1, 3), p_W, \alpha))$ assigns probability 1 to τ . Now consider the finite path $r = (W, 0)(\frac{3}{2}, p_W)(F, \frac{3}{2})$ of π : then the corresponding set of finite paths of $\hat{\pi}$ is $r' = (W, [0, 0])((1, 3), p_W, \alpha)((1, 3), p_W, \text{lb})(F, (1, 3))$ and $r'' = (W, [0, 0])((1, 3), p_W, \alpha)((1, 3), p_W, \text{rb})(F, (1, 3))$. Now say that $\pi(r)$ assigns probability 1 to the action $(\frac{9}{2}, p_F)$: then both $\hat{\pi}(r')$ and $\hat{\pi}(r'')$ assign probability 1 to the action $((4, 5), p_F, \alpha')$, where $\alpha'((4, 5), p_F, \text{lb}) = \frac{1}{2}$ and $\alpha'((4, 5), p_F, \text{rb}) = \frac{1}{2}$ (note that $\alpha'((4, 5), p_F, \text{lb}) = 5 - \frac{9}{2}$). Hence, regardless of whether $((1, 3), p_W, \text{lb})$ or $((1, 3), p_W, \text{rb})$ was visited, scheduler $\hat{\pi}$ makes the *same* choice to mimic $\pi(r)$.

The following lemma considers the converse direction, namely that (starting from a given location with the clock equal to 0) any scheduler of $\llbracket \mathfrak{M}[\mathcal{P}] \rrbracket$ can be mimicked by a \mathbb{B} -minimal scheduler of $\llbracket \mathcal{P} \rrbracket$ such that the schedulers assign the same probability of reaching a certain set of locations.

Lemma 3. *Let $l \in L$ be a location and let $F \subseteq L$ be a set of locations. Given a scheduler $\hat{\pi} \in \Sigma^{\llbracket \mathfrak{M}[\mathcal{P}] \rrbracket}$, there exists a \mathbb{B} -minimal scheduler $\pi \in \Sigma_{\mathbb{B}}^{\llbracket \mathcal{P} \rrbracket}$, such that $\Pr_{(l,0)}^{\pi}(\diamond T_F) = \Pr_{\text{reg}(l,0)}^{\hat{\pi}}(\diamond \mathfrak{F}_F)$.*

We characterise the size of a 1c-cdPTA as the sum of the number of its locations, the size of the binary encoding of the clock constraints used in invariant conditions and guards, and the size of the binary encoding of the constants used in the distribution templates of the probabilistic edges (i.e., c_e^p and d_e^p for each $p \in \text{prob}$ and $e \in 2^{\{x\}} \times L$).

Theorem 1. *Given a 1c-cdPTA $\mathcal{P} = (L, \text{inv}, \text{prob})$, $l \in L$ and $F \subseteq L$, the quantitative and qualitative problems can be solved in polynomial time.*

The theorem follows from Lemma 1, Lemma 2, Lemma 3, Proposition 1 and the fact that quantitative and qualitative problems for IMDPs can be solved in polynomial time, given that there exist polynomial-time algorithms for analogous problems on IMCs with the semantics adopted in this paper [30, 11, 9, 33], and from the fact that the IMDP construction presented in this section gives an IMDP that is of size polynomial in the size of the 1c-cdPTA. We add that the quantitative problems for 1c-cdPTAs are PTIME-hard, following from the PTIME-hardness of reachability for MDPs [29], thus establishing PTIME-completeness for quantitative problems for 1c-cdPTAs.

5 Conclusion

We have presented a method for the transformation of a class of 1c-cdPTAs to IMDPs such that there is a precise relationship between the schedulers of the 1c-cdPTA and the IMDP, allowing us to use established polynomial-time algorithms for IMDPs to decide quantitative and qualitative reachability problems on the 1c-cdPTA. Overall, the results establish that such problems are in PTIME. The techniques rely on the initialisation requirement, which ensures that optimal choices for non-constant probabilistic edges correspond to the left or right endpoints of intervals that are derived from the syntactic description of the 1c-cdPTA. The initialisation requirement restricts dependencies between non-constant probabilistic edges: while this necessarily restricts the expressiveness of the formalism, the resulting model nevertheless retains the expressive power to represent basic situations in which the probability of certain events depends on the exact amount of time elapsed, such as those described in the introduction.

The IMDP construction can be simplified in a number of cases: for example, in the case in which at most two outcomes e_1, e_2 of every probabilistic edge p are non-constant, i.e., for which $d_{e_1}^p \neq 0$ and $d_{e_2}^p \neq 0$, endpoint indicators are unnecessary; instead, when a probabilistic edge is taken from an open interval \bar{B} , each of e_1 and e_2 are associated with (non-singleton) intervals (other outcomes are associated with singleton intervals), and the choice of probability to assign between the two intervals represents the choice of clock valuation in \bar{B} . This construction is also polynomial in the size of the 1c-cdPTA. Future work could consider lifting the initialisation requirement: we conjecture that this is particularly challenging for quantitative properties, in particular recalling that Figure 2 of [32] provides an example of a non-initialised 1c-cdPTA for which the maximum probability of reaching a certain location is attained by choosing a time delay corresponding to an irrational number. Solutions to the qualitative problem for non-initialised 1c-cdPTAs could potentially utilise connections with parametric MDPs [19, 35]. Furthermore, time-bounded reachability properties could also be considered in the context of 1c-cdPTAs.

References

1. S. Akshay, P. Bouyer, S. N. Krishna, L. Manasa, and A. Trivedi. Stochastic timed games revisited. In P. Faliszewski, A. Muscholl, and R. Niedermeier, editors, *Proc. MFCS 2016*, volume 58 of *LIPICs*, pages 8:1–8:14. Leibniz-Zentrum für Informatik, 2016.
2. R. Alur and D. L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, 1994.
3. C. Baier and J.-P. Katoen. *Principles of model checking*. MIT Press, 2008.
4. N. Bertrand, B. Bordais, L. Hélouët, T. Mari, J. Parreaux, and O. Sankur. Performance evaluation of metro regulations using probabilistic model-checking. In S. C. Dutilleul, T. Lecomte, and A. B. Romanovsky, editors, *Proc. RSSRail 2019*, volume 11495 of *LNCS*, pages 59–76. Springer, 2019.
5. N. Bertrand, P. Bouyer, T. Brihaye, Q. Menet, C. Baier, M. Größer, and M. Jurdzinski. Stochastic timed automata. *Logical Methods in Computer Science*, 10(4), 2014.
6. N. Bertrand, T. Brihaye, and B. Genest. Deciding the value 1 problem for reachability in 1-clock decision stochastic timed automata. In G. Norman and W. H. Sanders, editors, *Proc. QEST 2004*, volume 8657 of *LNCS*, pages 313–328. Springer, 2014.
7. P. Bouyer, U. Fahrenberg, K. G. Larsen, N. Markey, J. Ouaknine, and J. Worrell. Model checking real-time systems. In E. M. Clarke, T. A. Henzinger, H. Veith, and R. Bloem, editors, *Handbook of Model Checking*, pages 1001–1046. Springer, 2018.
8. P. Bouyer, K. G. Larsen, and N. Markey. Model checking one-clock priced timed automata. *Logical Methods in Computer Science*, 4(2), 2008.
9. S. Chakraborty and J.-P. Katoen. Model checking of open interval Markov chains. In M. Gribaudo, D. Manini, and A. Remke, editors, *Proc. ASMTA 2015*, volume 9081 of *LNCS*, pages 30–42. Springer, 2015.
10. K. Chatterjee, K. Sen, and T. A. Henzinger. Model-checking ω -regular properties of interval Markov chains. In R. Amadio, editor, *Proc. FOSSACS 2008*, volume 4962 of *LNCS*, pages 302–317. Springer, 2008.
11. T. Chen, T. Han, and M. Kwiatkowska. On the complexity of model checking interval-valued discrete time Markov chains. *Information Processing Letters*, 113(7):210–216, 2013.
12. E. M. Clarke, O. Grumberg, and D. A. Peled. *Model checking*. MIT Press, 2001.
13. L. Feng, C. Wiltsche, L. R. Humphrey, and U. Topcu. Synthesis of human-in-the-loop control protocols for autonomous systems. *IEEE Trans. Automation Science and Engineering*, 13(2):450–462, 2016.
14. V. Forejt, M. Kwiatkowska, G. Norman, and D. Parker. Automated verification techniques for probabilistic systems. In M. Bernardo and V. Issarny, editors, *Formal Methods for Eternal Networked Software Systems (SFM 2011)*, volume 6659 of *LNCS*, pages 53–113. Springer, 2011.
15. R. Givan, S. M. Leach, and T. L. Dean. Bounded-parameter Markov decision processes. *Artificial Intelligence*, 122(1-2):71–109, 2000.
16. H. Gregersen and H. E. Jensen. Formal design of reliable real time systems. Master’s thesis, Department of Mathematics and Computer Science, Aalborg University, 1995.
17. S. Haddad and B. Monmege. Interval iteration algorithm for MDPs and IMDPs. *Theoretical Computer Science*, 735:111–131, 2018.

18. E. M. Hahn. *Model checking stochastic hybrid systems*. PhD thesis, Universität des Saarlandes, 2013.
19. E. M. Hahn, T. Han, and L. Zhang. Synthesis for PCTL in parametric Markov decision processes. In M. G. Bobaru, K. Havelund, G. J. Holzmann, and R. Joshi, editors, *Proc. NFM 2011*, volume 6617 of *LNCS*, pages 146–161. Springer, 2011.
20. V. Hashemi, H. Hatefi, and J. Krcál. Probabilistic bisimulations for PCTL model checking of interval MDPs. In É. André and G. Frehse, editors, *Proc. SynCoP 2014*, volume 145 of *EPTCS*, pages 19–33, 2014.
21. T. A. Henzinger, P. W. Kopke, A. Puri, and P. Varaiya. What’s decidable about hybrid automata? *Journal of Computer and System Sciences*, 57(1):94–124, 1998.
22. B. Jonsson and K. G. Larsen. Specification and refinement of probabilistic processes. In *Proc. LICS 1991*, pages 266–277. IEEE Computer Society, 1991.
23. M. Jurdziński, F. Laroussinie, and J. Sproston. Model checking probabilistic timed automata with one or two clocks. *Logical Methods in Computer Science*, 4(3):1–28, 2008.
24. I. O. Kozine and L. V. Utkin. Interval-valued finite Markov chains. *Reliable Computing*, 8(2):97–113, 2002.
25. M. Kwiatkowska, G. Norman, R. Segala, and J. Sproston. Automatic verification of real-time systems with discrete probability distributions. *Theoretical Computer Science*, 286:101–150, 2002.
26. F. Laroussinie, N. Markey, and P. Schnoebelen. Model checking timed automata with one or two clocks. In P. Gardner and N. Yoshida, editors, *Proc. CONCUR 2004*, volume 3170 of *LNCS*, pages 387–401. Springer, 2004.
27. A. Nilim and L. El Ghaoui. Robust control of Markov decision processes with uncertain transition matrices. *Operations Research*, 53(5):780–798, 2005.
28. G. Norman, D. Parker, and J. Sproston. Model checking for probabilistic timed automata. *Formal Methods in System Design*, 43(2):164–190, 2013.
29. C. H. Papadimitriou and J. N. Tsitsiklis. The complexity of Markov decision processes. *Mathematics of Operations Research*, 12(3):441–450, 1987.
30. A. Puggelli, W. Li, A. L. Sangiovanni-Vincentelli, and S. A. Seshia. Polynomial-time verification of PCTL properties of MDPs with convex uncertainties. In N. Sharygina and H. Veith, editors, *Proc. CAV 2013*, volume 8044 of *LNCS*, pages 527–542. Springer, 2013.
31. K. Sen, M. Viswanathan, and G. Agha. Model-checking Markov chains in the presence of uncertainties. In H. Hermanns and J. Palsberg, editors, *Proc. TACAS 2006*, volume 3920 of *LNCS*, pages 394–410, 2006.
32. J. Sproston. Probabilistic timed automata with clock-dependent probabilities. In M. Hague and I. Potapov, editors, *Proc. RP 2017*, volume 10506 of *LNCS*, pages 144–159. Springer, 2017.
33. J. Sproston. Qualitative reachability for open interval Markov chains. In I. Potapov and P.-A. Reynier, editors, *Proc. RP 2018*, volume 11123 of *LNCS*, pages 146–160. Springer, 2018.
34. J. Sproston. Probabilistic timed automata with one clock and initialised clock-dependent probabilities. *CoRR*, 2020.
35. T. Winkler, S. Junges, G. A. Pérez, and J. Katoen. On the complexity of reachability in parametric Markov decision processes. In W. Fokkink and R. van Glabbeek, editors, *Proc. CONCUR 2019*, volume 140 of *LIPICs*, pages 14:1–14:17. Leibniz-Zentrum für Informatik, 2019.
36. D. Wu and X. D. Koutsoukos. Reachability analysis of uncertain systems using bounded-parameter Markov decision processes. *Artificial Intelligence*, 172(8-9):945–954, 2008.